A Wilson Yang–Mills Construction

A Natural-Language Mathematical Exposition

Jeremy Rodgers
Independent Researcher

https://www.everythingequation.com

DOI: 10.5281/zenodo.24173681

12 November 2026

Abstract

This paper gives a self-contained natural-language exposition of the expanded *Wilson Yang–Mills Construction: Mathematical Translation*. For a given closed matrix subgroup of $U(m)$ with nonabelian simple real Lie algebra, the construction supplies one calibrated four-dimensional Wilson family, full fixed-period mesh limits and a subsequential outer volume limit, a common hierarchy of local fields and phases, and a selected local Wightman theory with Poincaré covariance, forward-cone spectrum and a unique vacuum. We give the finite quadratic kernels and nonlinear geometry, marked integration, coupling reset and shooting histories, physical detector and equation quotient, phase and boundary reconstruction, separated and nested operator products, contact extensions, and the finite-to-continuum stress and trace calculation with their hypotheses and derivations. The centered trace obeys

$$\lim_{r\downarrow 0} r^8 \log^2(1/r) G_\Theta(r) = \frac{3 \dim_{\mathbb{R}} \mathfrak{g}}{\pi^4} > 0.$$

The associated exclusions retain their stated finite mass-shell and Maxwell comparison hypotheses. No positive mass gap is asserted. Two formula amendments are derived and documented with the conventions used throughout the argument.

Attribution. The underlying Yang–Mills Lean work is credited to OpenAI. Jeremy Rodgers prepared this mathematical exposition with AI assistance.

Funding. No external funding.

Contents

1 Introduction: the construction and its dependencies 1

2 Statement of the public conclusion 3

3 Notation and analytic conventions 5

4 Finite-lattice positivity 9

5 Quadratic kernels and their ultraviolet limit 11

6 Local nonlinear geometry and estimates 38

7 Marked integration and the coupling reset 111

8 Convergence of the interacting scale histories 149

9 Physical labels and their detection 238

10 Free shell charges and separation of the Ward kernel 254

11 The common continuum hierarchy 276

12 Phase decomposition and Wightman reconstruction 288

13 Separated and nested operator products 339

14 Extension across collision sets 360

15 The Gaussian short-distance comparison 372

16 The centered trace and its non-freeness test 385

17 Completion of the proof 451

18 Conclusion 453

A Editorial amendments and conventions 453

B Source, attribution and reproducibility 455

1 Introduction: the construction and its dependencies

The starting point is a finite probability law: group elements live on oriented lattice links, and the Wilson action measures each plaquette's departure from the identity. The problem addressed by the supplied construction is to control one family of these laws while the mesh tends to zero and the physical period subsequently tends to infinity. Inserting local gauge-invariant functions into that same family is meant to produce a compatible hierarchy, rather than unrelated limits for different observables. The public theorem in Section 2 states the exact common-family and subsequence quantifiers.

The derivation proceeds through a finite linear model, an actual nonlinear integration, and two successive continuum limits. The quadratic model provides local inverse kernels and stationary response maps. Nonlinear charts compare group links with that model, while positive energy bounds pay for exceptional configurations. Marked local activities keep products of insertions attached to the same represented density. An exact reset identity and finite coupling shooting select calibrated histories on one calendar. Detector estimates then separate physical label classes and identify the equation image before the common continuum hierarchy is formed. These constructions supply the estimates and compatibility statements used in the limits and reconstruction.

Reflection positivity and slab bounds provide Hilbert-space control. The ground compression algebra supplies the phase space and its measure; tube continuation and tempered boundary limits supply fields on one common dense domain. The final part returns to actual finite current and stress functions. Their local equation estimates, coefficient calibration and Gaussian comparison determine a conserved physical stress and the centered trace. Positivity upgrades comparison along arbitrary retained scaling sequences to the full positive-distance trace limit. That limit supports the precise comparison exclusions stated in the public theorem.

Source and the two-paper release. The mathematical source is the expanded project extraction, *Mathematical Translation* [1], prepared with AI assistance. It presents the Yang–Mills construction associated with the Lean work credited to OpenAI. This paper develops that source as a complete mathematical exposition. Its definitions, hypotheses, estimates and derivations are included here, together with the formula amendments in Appendix A.

The accompanying flagship paper, *Shadow Theory and the Wilson Yang–Mills Construction: A Source–Readout Reconstruction* [2], applies Jeremy Rodgers's Shadow Theory, source–readout and Source Lift programme to this construction [3]. Rodgers originated and directed that programme; AI assisted its execution and the preparation of the two expositions. The papers share the Zenodo release 10.5281/zenodo.24173681. Each contains its complete mathematical development: this paper presents the construction, and the flagship develops its source–readout realization and recovery. Their common subject does not require a reader of either paper to obtain an essential derivation from the other.

The given gauge group, faithful representation, four-dimensional lattice, Wilson action and calibration conventions are physical and mathematical inputs. The detailed construction supplies stress Ward identities, nested expansions and contact extensions on its common witnesses; the public theorem retains its original list of clauses.

How to read the argument. Read the finite setup and conventions first, especially the healthy-domain derivative and support budgets. The quadratic section explains the aliases, metric amendment and filled stationary predictors. The nonlinear and marked-integration sections use those objects to prove an actual density identity. The long scale-history section then buys every fixed finite request in a specified order, and distinguishes a full Cauchy limit from an auxiliary extraction.

The label, shell-charge and quotient sections are needed for reconstruction and stress detection. The OPE, contact, Gaussian and trace sections retain their own finite horizons, comparison devices and domains. The closing assembly returns to the theorem and identifies all its witnesses on one family and one refinement.

Background. The finite law is Wilson's lattice formulation [4]; compact group integration uses Haar measure [5]. Wick pairings [6], reflection positivity and Euclidean reconstruction [10, 11, 12], chessboard inequalities [13], positive-functional representations [7, 8], and the character-algebra approximation [9] form the mathematical background. Renormalization and cluster methods are described in [16, 17, 18, 14]; the field-theoretic framework is that of [15], and the conventional one-loop coefficient is associated with [19, 20]. These references credit the background; the actual hypotheses, constants and derivations used here are stated locally.

Notation	Role in this paper
$G \subseteq U(m)$	Given closed group and faithful matrix representation; $r = \dim_{\mathbb{R}} \mathfrak{g}$ is the color dimension.
$h_N = d_N^{-1}, L_p$	Physical fine mesh and physical period in the public family.
a_i, L_i, t_i	Intermediate spacing, integer block ratio, and $t_i = \log(a_i^{-1})$.
$\widehat{\beta}_i, \beta_i$	Prescribed reference clock and actual represented coupling; their comparison is proved in the running argument.
ϑ, θ	Healthy-field thresholds; the local text specifies which schedule or fixed chart is being used.
F_d, E_d	Finite physical label filtration and coefficient spaces; the relevant equation quotient is specified before use.
χ, S^{χ}, S	Character phase, its separated hierarchy, and the unselected barycenter hierarchy. A scalar cutoff also denoted χ is always a function with an explicit position argument.
T_{ab}, Θ	Conserved physical stress and its selected-vacuum centered trace. Finite candidates are defined separately.

2 Statement of the public conclusion

The assertion below is the supplied construction theorem in its public form. Its detailed proof occupies the subsequent sections. The Wilson family is fixed before a volume subsequence is requested, and the full mesh limit is obtained at each suitable fixed period before an outer volume subsequence is taken. An auxiliary compactness extraction is subsequently identified with that full inner limit. The concluding trace test concerns precisely the finite mass-shell class specified below.

Let G be a closed subgroup of $U(m)$ in its given faithful matrix representation. Its real Lie algebra is

$$\mathfrak{g} = \{X : \exp(tX) \in G \text{ for every } t \in \mathbb{R}\}.$$

Assume that $\mathfrak{g}$ is nonabelian and simple: every real linear subspace $V \subseteq \mathfrak{g}$ with $[\mathfrak{g}, V] \subseteq V$ is either zero or all of $\mathfrak{g}$. No connectedness hypothesis is added.

The common Wilson family. There is one Wilson family $\mathcal{C}$ for this G, with strictly increasing positive integer periods L_p and denominators d_N. Its mesh is $h_N = d_N^{-1}$ and its periodic lattice has side $L_p d_N$. Group elements are assigned to its oriented links. With normalized Haar measure ν and the independent link prior, the probability law is

$$d\mu_{N,p}(U) = Z_{N,p}^{-1} e^{-\beta_{N,p} S(U)} \prod_e d\nu(U_e),$$

$$S(U) = \sum_{x,i<j} \left[1 - \frac{1}{m} \operatorname{Re} \operatorname{tr}(U_{x,i} U_{x+e_i,j} U_{x+e_j,i}^{-1} U_{x,j}^{-1})\right].$$

The family includes a token space and continuous, gauge-invariant local insertions $I_{N,p,v,x}$ for every token v. Each v has a fixed finite set of integer offsets, independent of N, p, x, controlling its insertion: if two configurations agree on the links whose endpoints lie in the translated offset set, their insertion values agree. There is a positive slope s such that

$$\sup_p \left| \frac{\beta_{N,p}}{\log(h_N^{-1})} - s \right| \longrightarrow 0 \qquad (N \to \infty).$$

One hierarchy on a permitted refinement. For every strictly increasing sequence $p(k)$ there is a strictly increasing subsequence index $q(k)$ and one hierarchy $\mathcal{T}$ on $p(q(k))$ with the following properties. Its fields have finite-dimensional complex normed label spaces F_d, $d \geq 0$, with linear injections into the label space of one Wightman theory that jointly cover that space. Tokens project surjectively onto the disjoint union of these graded spaces. There is a compact metric character space $\mathfrak{X}$, a full-support Borel probability measure μ_{ph} on it, a selected character, and separated Euclidean correlations S_n^χ for each $\chi \in \mathfrak{X}$ and S_n for the unselected hierarchy. They are smooth away from collisions, invariant under simultaneous permutations of slots and points, continuous in χ on that separated domain, and obey

$$S_n = \int_{\mathfrak{X}} S_n^\chi \, d\mu_{\mathrm{ph}}(\chi).$$

The subscript on the phase measure distinguishes it from the finite Wilson laws; it does not change the measure in the supplied statement.

For a smooth compactly supported collision-free test t and a finite token word $v = (v_1, \ldots, v_n)$, the lattice smeared moment is

$$R_{N,p,v}(t) = h_N^{4n} \sum_{x_1,\ldots,x_n} t(x_1^{\mathrm{phys}}, \ldots, x_n^{\mathrm{phys}}) \int \prod_{i=1}^n I_{N,p,v_i,x_i}(U) \, d\mu_{N,p}(U),$$

where $x \mapsto x^{\mathrm{phys}}$ is the centered physical placement. At each period whose centered open box contains $\operatorname{supp} t$, there is a full mesh limit

$$r_{p,v}(t) = \lim_{N\to\infty} R_{N,p,v}(t).$$

Along $p(q(k))$ it converges to the integral of t against the unselected correlation for the projected word. Every such fixed test is suitable at all sufficiently large periods of that subsequence. Thus the inner mesh limit is full at each suitable fixed volume; the outer volume limit uses the permitted refinement.

The selected Wightman theory. The selected theory has a normalized vacuum Ω and one common dense complex invariant domain, the span of finite smeared field words applied to Ω. Fields depend complex-linearly on labels and Schwartz tests, and their values on each domain vector depend continuously on the test. They satisfy the smeared adjoint identity, spacelike local commutativity and Poincaré covariance. Translations act by a strongly continuous unitary group whose joint energy-momentum projection measure is supported in the closed forward cone. The invariant vacuum is unique up to scalar. Every nonzero label gives a nonzero smeared operator, and at least one field acts nontrivially. The public symmetry coordinates cover every proper orthochronous affine Poincaré transformation. The selected Euclidean correlations are the values of a holomorphic continuation in the tube of strictly ordered imaginary times. Its smeared boundary limits on product Schwartz tests give the Wightman vacuum moments.

Separated operator products. For a cluster $a = (a_i)$ of graded fields, write $|a| = \sum_i \deg(a_i)$. There are finitely many channels at each degree j, with fields $O_{j,\lambda}$ and coefficients $C_{a,j,\lambda}(x; o)$, in one system independent of phase and truncation. Given $M \geq 0$, $\eta > 0$ and $\varepsilon > 0$, every fixed channel satisfies, for all sufficiently small $r > 0$,

$$|C_{a,j,\lambda}(x; o)| \leq r^{j-|a|-\varepsilon}$$

whenever $|x_i - o| \leq Mr$ and $|x_i - x_{i'}| \geq \eta r$ for $i \neq i'$. For a fixed spectator word b, $c > 0$ and $D \geq 0$, the same cluster domain and the spectator separations $|o - y_l| \geq c$ and $|y_l - y_{l'}| \geq c$ give

$$\left| S^\chi(a(x), b(y)) - \sum_{j=0}^{D} \sum_{\lambda} C_{a,j,\lambda}(x; o) S^\chi(O_{j,\lambda}(o), b(y)) \right| \leq r^{D+1-|a|-\varepsilon}.$$

This holds for all sufficiently small r, uniformly in all phases and also for the unselected hierarchy. All terms belong to the same hierarchy. The threshold may depend on the fixed words and the stated parameters.

The centered trace and its comparison class. There is a centered degree-four field Θ. Let $G_\Theta(r)$ be its selected two-point correlation at spatial zero and Euclidean times $-r/2, r/2$. Then

$$\lim_{r\downarrow 0} r^8 \log^2(1/r) G_\Theta(r) = \frac{3 \dim_{\mathbb{R}} \mathfrak{g}}{\pi^4} > 0.$$

It cannot agree at every sufficiently small radius with a finite complex linear combination of finite products of the specified mass-shell integrals. Each factor is of the form

$$\int_{\mathbb{R}^3} \frac{E_\mu(p)^{t-1}}{2} P(p) e^{-rE_\mu(p)}\, d^3p, \qquad E_\mu(p) = \sqrt{\mu^2 + |p|^2},$$

where $\mu \geq 0$, $t \in \mathbb{Z}_{\geq 0}$ and $P : \mathbb{R}^3 \to \mathbb{R}$ is continuous and positively homogeneous of a nonnegative integer degree. Here μ denotes the comparison mass called m in the source, to distinguish it from the representation size. This is the public non-freeness test with its precise comparison class. No positive physical mass gap is asserted.

The public theorem is obtained by choosing blocking penalty $b = 1$ in the detailed construction, taking the refinement for the requested volume sequence, and projecting its single set of witnesses. This preserves the common hierarchy across all conclusions. The detailed construction also supplies stress Ward identities, nested expansions and contact extensions. They are developed below on its stronger common witnesses; they are neither extra assumptions nor additional clauses silently added to the public theorem.

3 Notation and analytic conventions

The same letters occur in different mathematical roles in this long construction. In particular, L_p is a physical period, whereas L_i is an integer block ratio; m is the dimension of the given matrix representation, whereas $m(q)$ below is a Fourier multiplier. The context and arguments distinguish these uses. The reference coupling $\widehat{\beta}_i$ is a prescribed clock, while an unadorned β_i is the actual calibrated coupling. They agree only to the extent established by the running estimates. None of the estimates below permits substituting an arbitrary covariance for the particular positive covariance supplied by its construction.

All lattices are four dimensional. The integer cover is $\mathbb{Z}^4$, and the periodic lattice of side n is $(\mathbb{Z}/n\mathbb{Z})^4$. A positive link is (x, i), joining x to $x + e_i$; reading it in reverse gives the inverse group element. A plaquette in the positive plane $i < j$ is the ordered product appearing in the theorem. Signed coordinate permutations act on links with this reversal convention. Lattice distance is the maximum of the centered coordinate distances. A support is a finite set of sites; its internal links have both endpoints in that set.

The Lie algebra carries the physical inner product

$$\langle X, Y \rangle_{\mathfrak{g}} = -\frac{1}{m} \operatorname{Re} \operatorname{tr}(XY).$$

It is positive definite on skew-Hermitian matrices and invariant under conjugation by G. Fix an orthonormal real color basis for this inner product, identifying the algebra with $\mathbb{R}^{\dim \mathfrak{g}}$. A color vector has matrix Frobenius norm $\sqrt{m}$ times its color norm. The group chord is $\operatorname{chord}(g) = \|g - 1\|_{\mathrm{F}}/\sqrt{m}$. Inversion and conjugation preserve it, and $\operatorname{chord}(gh) \leq \operatorname{chord}(g) + \operatorname{chord}(h)$. The Lie-algebra curvature of a holonomy g is the color projection of $(g - g^{-1})/2$. These normalizations are kept in the quadratic action, coupling slope and trace calculation.

Write $a_i > 0$ for the lattice spacing at scale i, and $L_i = a_i/a_{i+1}$ for its integer block ratio. The construction chooses odd ratios at least three. The scale time is $t_i = \log(a_i^{-1})$, so $t_{i+1} - t_i = \log L_i$. Block ratios can be increased stage by stage; subsequent lower bounds on the reference coupling are called delays. Such choices are always made before the terminal cutoff and physical period. The reference coupling is $\widehat{\beta}_i = \bar{\beta} + s t_i$, where

$$s = \frac{11 C_A}{24\pi^2}, \qquad C_A = -\frac{1}{\dim \mathfrak{g}} \sum_{\alpha} \operatorname{tr}(\operatorname{ad} T_\alpha)^2.$$

The simple nonabelian assumption gives $C_A > 0$. Intermediate represented couplings are distinct from this reference coupling and are compared with it explicitly.

For a finite support X, its radius-B halo is the union of the lattice balls of radius B about its points. The cover number $\operatorname{cov}_B(X)$ is the minimum number of such balls whose union is connected

by unit lattice steps and contains X. The centers need not lie in X. This connected-union definition, including its behavior under periodic wrapping, is the one used in all support estimates.

The numerical estimates often allow arbitrarily small power losses. For a fixed block ratio L, put

$$\beta_F(x, L) = x + s \log L, \qquad B(x) = \lceil (\log x)^2 \rceil, \qquad \mathrm{temp}(x) = x^{-1/2} e^{(\log\log x)^2}.$$

An envelope f has rate a if for every $b > a$,

$$|f(x)| \le \beta_F(x, L)^b \quad \text{for all sufficiently large } x.$$

The threshold may depend on the fixed stage and block ratio, but not on the period. Rate zero means subpower growth; it does not mean boundedness. The fine threshold is $\mathrm{fiery}(x) = \mathrm{temp}(\beta_F(x, L))$. Finite sums preserve the larger rate, products add rates, and a strictly positive exponent gap absorbs any fixed finite multiplicity or subpower factor. For instance, if f has rate a and g has rate b, choose slightly larger exponents for their separate bounds; their product is then bounded by every exponent strictly greater than $a + b$. This is the meaning of adding rates, rather than an assertion of an exact leading power. Where the later operator-product argument uses rates along a scale path, that path and its limiting logarithmic exponents are specified separately.

Weighted bounds on local terms include their derivative order, direction restrictions and exponential support weight. A bound on their sum always means an absolute bound, or equivalently a common bound on every finite partial sum of the nonnegative majorant. This convention permits countable dictionaries without assuming a finite number of terms. Terms such as a bill or a paid bound in the proof denote these explicit inequalities, not additional probabilistic assumptions.

3.1 Local activity domains, derivatives and absolute budgets

An activity is a local contribution to the represented density. Its budget simultaneously controls size, the number of supports meeting a site, and the cost of their spatial extent. A flag records failure of a plaquette small-field test. A regular derivative is taken on its declared healthy domain; it does not differentiate the discontinuous indicator at that domain's boundary. The distinction between an assigned gas support and the enlarged set of links actually inspected will be used in every locality claim involving flags.

How the bounds are used The local estimates below have three separate pieces of data: the set of configurations on which they hold, the derivative order and directions allowed there, and the weight used when supports are summed. A bound on one healthy configuration is not a bound on a neighborhood of that configuration. Likewise, a support budget is a sum over an indexed dictionary, so two terms with the same support still count twice. The subsequent integration keeps these distinctions when it differentiates and sums the actual insertion gas.

We fix the meaning of all activity estimates before using them. At a fixed finite torus $\Lambda_n = (\mathbb{Z}/n\mathbb{Z})^4$, let $\mathcal{C}_n = G^{\Lambda_n \times \{0,1,2,3\}}$ be the configuration space and $\mathcal{A}_n = \mathfrak{g}^{\Lambda_n \times \{0,1,2,3\}}$ its real space of left-link directions. A function is local on a site set X if its value depends only on links with both endpoints in X. Gauge invariance refers to $U_{x,i} \mapsto k_x U_{x,i} k_{x+e_i}^{-1}$.

For a plaquette p, write $V(p)$ for its four visited vertices and U_p for its holonomy. At threshold $0 < \vartheta \le 1$ define

$$\mathcal{F}_\vartheta(U) = \{x : \exists p,\ x \in V(p),\ \mathrm{chord}(U_p) > \vartheta\}, \qquad \mathcal{H}_\vartheta(X) = \{U : X \cap \mathcal{F}_\vartheta(U) = \varnothing\}. \tag{AB1}$$

Thus equality with the threshold is healthy. Each flag event is open, because it is a finite union of strict inequalities for continuous plaquette functions. Each healthy domain is closed and Borel. The mask is the literal function $\mathbf{1}_{\mathcal{H}_\vartheta(X)} F$; its boundary is not differentiated in any regular budget.

A coefficient which also reads the final flags in X is required to agree whenever both the links internal to X and $X \cap \mathcal{F}_\vartheta(U)$ agree. Its actual link read set is bounded by

$$X^{\mathrm{ins}} = X \cup \bigcup_{p : V(p) \cap X \neq \varnothing} V(p). \tag{AB2}$$

Agreement on internal links of this enlarged set fixes every plaquette incident to X, hence all flags inspected there. This is why flag-dependent payloads are assigned enlarged physical footprints even though compatibility in a gas is tested on the originally assigned supports.

For $A \in \mathcal{A}_n$, set

$$(\mathrm{inc}(U, tA))_e = \exp(tA_e)U_e, \qquad \nabla_A F(U) = \frac{d}{dt} F(\mathrm{inc}(U, tA))\Big|_{t=0}.$$

For an ordered list $\ell = (A_1, \dots, A_q)$, define recursively

$$\nabla_\varnothing F = F, \qquad \nabla_{A::\ell} F(U) = \frac{d}{dt} (\nabla_\ell F)(\mathrm{inc}(U, tA))\Big|_{t=0}. \tag{AB3}$$

The order matters; no commutation of left derivatives is assumed. These derivatives are required to exist where tested. A regular representative on X is a globally Borel, gauge-invariant local function F such that $v \mapsto F(\mathrm{inc}(U, v))$ is C^j near zero for every $U \in \mathcal{H}_\vartheta(X)$, and its recursively specified left derivatives have the derivative identities in (AB3) through order j. In particular the smooth germ is required also at a healthy configuration lying on the threshold boundary. Off-domain values of the representative remain available for locality and measurability, but do not enter its regular derivative price.

A budget point consists of integers $p \geq 2, j \geq 6, B \geq 1$, support exponent $s > 0$, and threshold $0 < \vartheta \leq 1$. We denote it by $\mathfrak{p} = (p, j, s, B, \vartheta)$. For anchor o, its allowed directions satisfy

$$|A_e| \leq \vartheta(1 + \mathrm{dist}(e^-, o))^p \quad \text{for every link } e. \tag{AB4}$$

Here e^- is the initial vertex of the stored positive link, and the distance is the periodic maximum-coordinate distance. A nonnegative number b is a derivative envelope for F on (o, X) if

$$|\nabla_\ell F(U)| \leq b \quad \text{whenever } U \in \mathcal{H}_\vartheta(X),\ |\ell| \leq j, \text{ and every } A \in \ell \text{ satisfies (AB4).} \tag{AB5}$$

Norms replace absolute values for vector-valued coefficients. The support price is

$$\mathrm{mass}_{\mathfrak{p}}(X, b) = e^{s\ \mathrm{cov}_B(X)} b. \tag{AB6}$$

The factor is at least one. Every sum of envelopes below means a nonnegative summable family together with its asserted bound on the sum. Equivalently, every finite partial sum obeys that bound; taking increasing finite partial sums proves summability and the same total bound for countable dictionaries.

A regular address is (o_i, X_i) with $o_i \in X_i$. A regular family E_i has anchored budget δ if it has envelopes $b_i \geq 0$ satisfying (AB5) and

$$\sup_{o \in \Lambda_n} \sum_{i : o_i = o} e^{s\,\mathrm{cov}_B(X_i)} b_i \leq \delta. \tag{AB7}$$

The representation used after the kinetic reset additionally has $E_i(1) = 0$. On a nonwrapping support its identity Hessian

$$D^2 [v \mapsto E_i(\mathrm{inc}(1, v))]_0 [A, C]$$

vanishes whenever A, C are affine link fields in the canonical lift about o_i. Explicitly an affine field has

$$A_{x,a} = a_a + \sum_b \mathrm{lift}(x - o_i)_b \, d_{ab}. \tag{AB8}$$

The reset chapter constructs this Hessian cancellation; it is not imposed on arbitrary primitive functions before projection.

A forced address consists of a support Y_i, a nonempty inventory $I_i \subseteq Y_i$, and a fixed extra predicate $P_i(U)$. Its coefficient K_i is globally Borel and gauge invariant and vanishes unless

$$Y_i \cap \mathcal{F}_\vartheta(U) = I_i, \qquad P_i(U). \tag{AB9}$$

The predicate is Borel and gauge invariant; both coefficient and predicate may read links and final flags on Y_i, hence have the physical read enlargement (AB2). A forced budget κ consists of numbers $c_i \geq 0$ such that $|K_i(U)| \leq c_i$ on every configuration and

$$\sup_{x \in \Lambda_n} \sum_{i : x \in Y_i} e^{s \, \mathrm{cov}_B(Y_i)} c_i \leq \kappa. \tag{AB10}$$

This is a support-hit budget, in contrast to the single-anchor sum (AB7). Coefficients equal to zero still retain their addresses and eligibility predicates.

For a finite set D of forced addresses, the gas term is zero unless their supports are pairwise disjoint and their inventories have union exactly $\mathcal{F}_\vartheta(U)$; otherwise it is $\prod_{i \in D} K_i(U)$. Summing these terms over finite D defines the forced gas. The regular exponent is the sum of $\mathbf{1}_{\mathcal{H}_\vartheta(X_i)} E_i$. At a fixed finite torus these series are absolutely convergent: assign each regular term to its anchor and each forced term to one site of its nonempty support. Equations (AB7)–(AB10) bound their total envelope sums by $|\Lambda_n|\delta$ and $|\Lambda_n|\kappa$. The sum of absolute products over all finite subsets is at most $\prod_i (1 + c_i) \leq e^{\sum_i c_i}$. Restricting to compatible subsets only decreases it. If there are no flags, only the empty forced set can contribute, since each inventory is nonempty; the gas is then exactly one.

Signed translations are implemented by fixed permutations of the whole regular and forced address lists, carrying anchors, supports, inventories, predicates and coefficient functions together. Thus comparing two activity values in a single chart always means comparing the same lists and symmetry permutations. It does not mean identifying two representations merely because their total densities agree.

An early source budget uses $s + 3$, rather than s, in (AB7). This is the pre-reset reserve. The three additional support-exponent units pay the local splitting and conversion operations explicitly described in the reset and signed-expansion chapters. Reducing the exponent from $s + 3$ to s decreases every nonnegative summand. The rotation of derivative directions costs the finite factor

$$\mathrm{turnFee}(\mathfrak{p}) = (2\,2^p)^j. \tag{AB11}$$

Indeed the anchor itself is transported exactly. A reversed link is stored at its other endpoint, at distance at most one from the transported base, so $1 + d + 1 \leq 2(1 + d)$ costs 2^p in (AB4). Its right increment can be converted to a left increment by the adjoint action of the background link, which preserves color norm. Expanding each pulled-back left derivative into its positive-link and negative-link contributions gives at most 2^q ordered lists at order q; each has length q and direction inflation at most 2^p. Multilinearity of the iterated derivatives bounds their sum by $(2\,2^p)^q b$, which is at most (AB11) times b for $q \leq j$. This expansion follows by induction on the ordered list, using the exact left/right increment identity and the disjoint link supports of the positive and negative

parts. Support enlargement is charged separately by its cover-number inequality; it is not included in this factor.

At level i put $x_i = \widehat{\beta}_i$ and $M_i = \exp((\log\log x_i)^2)$. The normalizations used in (LCN1) are the calendar's actual allowances

$$b_{P_i} = x_i^{-11/10}, \qquad k_{P_i} = \exp(-M_i^{1/5}). \tag{AB12}$$

Both are strictly positive. Thus $\|A\|_{P_i} \le t$ means exactly that its regular and forced components obey (AB7) and (AB10) with budgets tb_{P_i} and tk_{P_i}, respectively, in their fixed address layout. The retained physical states have $t \le 1/8$. Triangle inequalities give the norm properties used later: sums of admissible nonnegative envelopes remain envelopes, derivatives are linear, and scalar multiplication multiplies all envelopes by the scalar modulus. Least envelopes, followed by finite partial-sum limits, give attainment of an infimum whenever it is used in (LCN1). The separated completion removes only differences of zero price in these tested germs and forced values.

The source aliases `RG/Masks`, `RG/Activities`, `Quotes`, `Investor` and `FullVolume/Sym/Shorn` refer to the mask, activity, reset-budget and normalization definitions made explicit in (AB1)–(AB12). They will be used with the same domains and directions in every subsequent reference to a regular, anchored, early or forced-hit budget.

4 Finite-lattice positivity

Reflection positivity starts at the finite Wilson law. Cutting a torus into two reflected halves rewrites each seam contribution as a positive kernel, so that a future observable paired with its reflected conjugate has nonnegative expectation. The finite-dimensional statement gives the null-space quotient used later to build a Hilbert space. It is also the input to the chessboard bounds: positivity converts a product of separated chamber observables into a product of individual rooted bounds. The details below retain the seam and orientation conventions needed for both uses.

First the finite Wilson law is an actual probability measure. Every unitary matrix V has $-m \le \operatorname{Re}\operatorname{tr} V \le m$: each diagonal entry has modulus at most one. Thus every plaquette price lies in $[0, 2]$. If there are F plaquettes and $\beta \ge 0$, its continuous density relative to the product Haar probability measure obeys

$$e^{-2\beta F} \le e^{-\beta S(U)} \le 1, \qquad e^{-2\beta F} \le Z \le 1.$$

In particular normalization is legitimate and bounded measurable cylinders are integrable. A site gauge change multiplies a link on its two ends; by Haar invariance its change of variables preserves the product prior. Its plaquette holonomy is conjugated at the base site, so the trace and the action are unchanged. The normalized law is consequently gauge invariant. A signed lattice permutation merely reindexes plaquettes, possibly reversing their holonomies; inversion preserves their real traces. It too preserves both the prior and the law. The local representatives in (PF12), and their finite products, therefore have precisely the finite-lattice expectations used below.

Reflection positivity is the nonnegativity of the expectation of a future observable times its reflected complex conjugate. It is a property of this finite law, before any quotient or completion. The Wilson measures used throughout have that property at every nonnegative coupling. We give the finite-dimensional argument, since it also supplies the rooted bounds used in extraction.

First, the kernel

$$K(g, h) = \operatorname{Re}\operatorname{tr}(gh^*)$$

on G is positive semidefinite: it is the real Euclidean inner product of the real and imaginary matrix entries. Finite sums, pullbacks, nonnegative multiples and pointwise products of positive

semidefinite kernels are positive semidefinite. For the product assertion, express each finite positive semidefinite matrix as a Gram matrix and tensor its Gram vectors. The exponential of such a kernel is positive semidefinite as well, by its power series with nonnegative coefficients and entrywise limits of the resulting finite matrices.

Cut the periodic lattice perpendicular to a coordinate axis. A leaf X consists of the links tangent to one slice, and a thread consists of the links joining successive slices. With the thread fixed at the identity, the factor joining two leaves has the form

$$W_\beta(X, Y; 1) = e^{-\beta q}\, E_\beta(X) E_\beta(Y) \exp\left[\frac{\beta}{m} \sum_\ell \operatorname{Re}\operatorname{tr}(X_\ell Y_\ell^*)\right], \tag{RP1}$$

where q counts the crossing plaquettes and E_β is the positive factor from the allocated half of the tangential action. Thus (RP1) is a continuous positive semidefinite kernel for $\beta \geq 0$.

For general threads, use their endpoint gauge action on the leaf links. Haar invariance gives the identity

$$\int W_\beta(X, Y; u)\, du = \iint W_\beta(u \cdot X, v \cdot Y; 1)\, du\, dv. \tag{RP2}$$

Both integrations use the product normalized Haar measure on the thread variables. The identity follows by replacing one thread variable by its left translate and combining the two endpoint gauge actions. The right side is positive semidefinite: apply the positive-kernel quadratic inequality to finite linear combinations of the averaged feature vectors, or approximate the integrals by simple functions in the same Gram identity. All integrands are continuous and bounded on compact spaces, so this passage and Fubini are justified. The averaged kernel in (RP2) is continuous by dominated convergence.

Now split the lattice into the two halves of the reflection. For a reflection through sites, the two halves share the seam-site variables, and after fixing them the cross kernel is the constant kernel 1. For a reflection through links, there are two seam threads on the periodic lattice. They are distinct even at the shortest allowed half-length. Their joint Haar law is the product thread law. After integration, the cross kernel is the product of the two kernels (RP2), one with its leaf arguments reversed. It is again positive semidefinite.

To make the splitting precise, list the seam variables, the two seam threads when present, and the open variables on the two halves. Assemble the original link configuration by their coordinate assignments and the endpoint gauge multiplications along the cut direction. Each change of coordinates multiplies a Haar link variable on the left or right by fixed already listed group elements, or inverts that variable. Iterated Haar invariance shows that this assembly preserves the product prior. Reading each plaquette shows that the action splits into the two half factors and the seam factor just described. A probe F supported strictly in one half reads only that half and the shared seam variables; its gauge invariance permits the endpoint gauge changes. Its reflected conjugate reads the other half.

Consequently, conditional on the seam variables, the unnormalized reflected pairing has the form

$$\iint f(X)\overline{f(Y)}\, K_{\mathrm{seam}}(X, Y)\, dX\, dY,$$

with K_{seam} positive semidefinite. This integral is nonnegative by the integrated Gram inequality. The half-action factors have been included in f. Bounded measurable probes are allowed: the Wilson density is bounded on the compact finite configuration space, and their reflected products are bounded and measurable, so all integrations are absolutely integrable. Integrating the conditional nonnegative quantities and dividing by the strictly positive Wilson partition function proves

$$\langle F\, \overline{F \circ \vartheta} \rangle_\beta \geq 0. \tag{RP3}$$

Here ϑ is the actual geometric link reflection, including inverse links. In particular no crossing link has been replaced inside the probe.

We next derive the chessboard inequality on an admissible rectangular tiling. Let $v > 0$ be its number of cells. For a function F in its strict reference chamber, write $F^{(c)}$ for its translated, reflected and, where prescribed, conjugated copy in cell c. Define

$$\mathcal{R}(F) = \left| \left\langle \prod_c F^{(c)} \right\rangle_\beta \right|^{1/v}.$$

For different chamber functions F_c, let $V(F_\bullet) = |\langle \prod_c F_c^{(c)} \rangle_\beta|$. The Cauchy–Schwarz inequality from (RP3), applied across any translated cell seam, gives

$$V(A)^2 \le V(A_+)V(A_-), \tag{RP4}$$

where A_+ and A_- are the assignments obtained by reflecting the corresponding half-assignment into the other half.

For a finite alphabet of functions, choose a cell assignment maximizing V. Both folded values in (RP4) are at most this maximum. If the maximum is positive, (RP4) forces both folded values to equal it; if it is zero this is immediate. Hence the property of being a maximizing assignment is preserved by every fold. Such a fold-closed property has a constant assignment: in one coordinate start with a one-cell-wide strip of identical slices and fold at its preceding boundary. This extends the strip by one slice while preserving the property. Continue until a half-period is filled, then reflect it to fill that coordinate. A fold in another coordinate preserves equality along the coordinates already filled. Repeat in all four coordinates. This is a finite construction and needs no power-of-two restriction on the number of slices.

For arbitrary F_c, apply this maximum principle to the finite alphabet

$$\widetilde{F}_c = F_c/(\mathcal{R}(F_c) + \epsilon), \qquad \epsilon > 0.$$

Every constant assignment of a normalized function has value at most 1, because its value is $[\mathcal{R}(F_c)/(\mathcal{R}(F_c) + \epsilon)]^v$. The maximum principle bounds their mixed assignment by 1. Rescale each cell and then let $\epsilon \downarrow 0$ to obtain

$$\left| \left\langle \prod_c F_c^{(c)} \right\rangle_\beta \right| \le \prod_c \mathcal{R}(F_c). \tag{RP5}$$

The same root is a seminorm on the chamber functions. Complex homogeneity follows because a copied scalar is either z or $\bar{z}$, and the absolute value of their full product is $|z|^v$. To prove subadditivity, expand the full copied product of $F + H$ into its 2^v assignments of F, H. The triangle inequality and (RP5) bound its expectation by $(\mathcal{R}(F) + \mathcal{R}(H))^v$. Take the positive v-th root. This proves all the rooted homogeneity, subadditivity and mixed-product estimates used below directly for the Wilson law.

5 Quadratic kernels and their ultraviolet limit

The quadratic analysis constructs the linear response around the flat configuration. There are three separate issues. The evaluated quadratic form on link potentials must converge; its curvature representation must remain positive and exponentially local; and the stationary predictor must have a regular Fourier symbol at the apparent zero-momentum singularity. The auxiliary metric amendment changes only directions not detected by curls. Thus retaining the physical potential precision is a proved identity, not an inference from positivity.

5.1 Flat quadratic ultraviolet convergence

Reading the finite linear records This opening convergence argument states the estimates required from a finite linear history before constructing that history below. A history, also called a horizon in the source notation, is a finite sequence of metrics and adjacent stationary response maps ending at the fixed bare quadratic form. Its price table records analytic widths, lower coercivity bounds and upper kernel bounds at each stage. A smoothing train is the finite chain of averaging patches defined in Section 5.8; a transition tower is the collection of auxiliary rows used to realize one adjacent response and its Gram identities. These names describe finite numerical and operator records. The later metric, tower and pinning calculations construct their required identities, so the opening conditional estimate is not an extra assumption of the public theorem.

Write a_i for the positive lattice spacing at level i, with a_i tending to zero. For a reporting level k put

$$T(k,i) = 2 + Z(k) + |t_i - t_k|,$$

where Z is nonnegative. Let $Q_{n,i}$ denote the cover kernel of the quadratic action at level i constructed with finer cutoff n. The convention is the action Hessian convention; there is no additional factor one half. The kernel is obtained from the curvature kernel by composition on both sides with the lattice coboundary, so it describes the quadratic form on link potentials. It does not determine an arbitrary extension of the curvature metric away from curls.

For $s \geq 0$, let $R_s(K)$ be the weighted row bound of a cover kernel K: for every output link and each fixed input orientation, sum the absolute kernel entries over input sites with weight $e^{s\,\mathrm{dist}}$, then take the supremum over the output link and input orientation. This is the source's `RowCost` at side one. Exponential weights control spatial locality as well as absolute summability.

A finite quadratic history has, at each $i \leq n$, five real Laurent rows: an unamended two-form metric H, its inverse, an amended metric M, its inverse, and a residual Gram root F. Their weighted coefficient bounds have width $\alpha(J(i)) > 0$ and price $\Omega(J(i))$. At every torus period,

$$HH^{-1} = MM^{-1} = I, \qquad F^*F + cI = M, \qquad Md_1 = Hd_1, \qquad M(0) = I.$$

Here $c > 0$ is the direct clip coefficient; the full Gram factor is the stack $(\sqrt{c}\,I, F)$. Both metrics are symmetric with quadratic forms between $\eta(J(i))I$ and $\Omega(J(i))I$, where $\eta > 0$. The terminal unamended metric is I. Each transition has a fine potential response A_0, curvature response T, and mixed, normal and image rows satisfying the stationary and Gram identities derived below. The mixed row has bound Ω/L^2; the normal and image rows have bound Ω. A word of r carried fine differences of A_0 or T has bound ΩL^{-1-r} or ΩL^{-2-r}, respectively, through order $J(i+1) - 2$. A carried difference increments the fine residue, shifts the coarse displacement when that increment crosses its centered seam, and subtracts the original entry. Thus it is the actual fine difference, including the carry. These all-period conditions are constructed by the metric, positive-margin, auxiliary-tower and pinning arguments below; they are not additional assumptions of the public theorem.

Put $y_i = w(J(i))$. The spacing hypotheses are

$$\begin{gathered}
0 < w(j) \leq \min\{1, P.\alpha(j)\}, \\
y_i \leq \min\{1, P.\alpha(J(i+1)), P.\alpha(J(i))\} \quad (i \geq e), \\
\text{one-step price at } i \ \leq \mathrm{block}(i)^{-\theta}, \qquad 2y_i \leq \mathrm{block}(i) y_{i+1}, \\
P.\Omega(J(i)) \leq B\,T(k,i)^l \quad (k \text{ arbitrary},\ i \geq e),
\end{gathered}$$

for $\theta > 0$, $B \geq 0$ and a nonnegative integer l. Here $P.\alpha$ and $P.\Omega$ are the width and upper-bound columns of the fixed price table. Each earlier history also uses the specified destination response T, the compatibility condition called `Flat.meshed` in the source.

Under these hypotheses there exist $A \geq 0$, an integer $d \geq 0$, and $c > 0$ such that, for all k and $e \leq i \leq N \leq n$,

$$\begin{aligned} &R_{y_i}(Q_{n,i} - Q_{N,i}) \\ &\leq AT(k,i)^d (a_N/a_i)^c. \end{aligned} \tag{1}$$

The source proves this through curvature packets. At the terminal comparison level N subtract the two saddle curvature kernels. Their strip bounds give a packet whose cost is at most

$$\mathrm{topTab}(N) = 2|\mathrm{Faces}|P.\Omega(J(N)).$$

This packet represents the difference of the quadratic forms on every scheduled torus. Propagate it backwards from i+1 to i using the two histories' fine response maps. The exact crossing identity preserves its representation of the quadratic difference; neutrality and translation covariance are retained. The weighted packet cost is multiplied by $\mathrm{quote}(i, y_{i+1})$. The width-crossing hypothesis licenses this propagation at width y_i. In particular, this is an identity for the two actual histories rather than a comparison with a separately chosen limiting history.

If C_i is the propagated cost, the recursive majorization is

$$C_N = \mathrm{topTab}(N), C_i \leq \mathrm{block}(i)^{-\theta} C_{i+1}.$$

One may first define a nonnegative majorant using the reciprocal recursive weight, called `Stair.weight` in the source. Its recurrence is exact before the one-step factor is replaced by its upper bound. Applying the polynomial-age decay lemma of Section 5.2 gives

$$C_i \leq AT(k,i)^d (a_N/a_i)^c.$$

In this particular terminal-only recurrence the majorant can also be written without the auxiliary weight notation. With $t_j = \log(a_j^{-1})$, it is

$$\widehat{C}_i = \mathrm{topTab}(N) \prod_{j=i}^{N-1} L_j^{-\theta} = \mathrm{topTab}(N) e^{-\theta(t_N - t_i)}.$$

The empty product gives its terminal value, and each backward step has exactly the required contraction. Set $u = t_N - t_i \geq 0$. The triangle inequality gives $T(k,N) \leq T(k,i) + u$, whereas $T(k,i) \geq 2$. The polynomial cap on the terminal price therefore gives

$$\widehat{C}_i \leq CT(k,i)^l (1+u)^l e^{-\theta u} \leq C' T(k,i)^l e^{-\theta u/2}.$$

The last constant is finite because the continuous function $(1+u)^l e^{-\theta u/2}$ is bounded on $[0, \infty)$. Since $e^{-u} = a_N/a_i$, this is the asserted estimate with $c = \theta/2$. The more general recurrence lemma below also allows forcing at intermediate levels. Neither argument assumes continuum convergence.

To recover the cover quadratic kernel, let D be the lattice coboundary and set $W(K) = D^* K D$, with K interpreted as its two-point kernel. Finite support of D gives a width-one row and column bound $L \geq 0$ for this paragraph. Weighted convolution and the distance triangle inequality give

$$R_s(W(K)) \leq |\mathrm{Faces}|^2 L^2 C$$

whenever the packet has width $s \le 1$ and cost C. This constant is independent of the torus size.

The required uniqueness concerns periodization of absolutely summable cover kernels, not inversion of a finite matrix. In the source these operations are respectively named `Flat.resolve` and `FreeResolvent.resolve`. Suppose two summable cover kernels give the same matrices on tori whose side lengths tend to infinity. For a fixed pair of cover sites, express a periodized entry as the sum over sites with the same torus image. For each fixed summand, eventual injectivity of wrapping makes its indicator tend to the indicator of the selected cover site. The absolute kernel entries form a summable dominating family. Dominated convergence recovers that cover entry from the torus entries. Both kernels have the same limit and therefore agree entry by entry.

Apply this argument to $Q_{n,i} - Q_{N,i}$ and $W(K_i)$. Their finite-torus quadratic matrices agree by the propagated packet identity, and both have summable rows. Thus their cover kernels agree. The preceding convolution bound proves (1), absorbing $|\mathrm{Faces}|^2 L^2$ into A.

For each fixed $i \ge e$, (1) makes every entry of $Q_{n,i}$ a Cauchy sequence, since a_N tends to zero and $c > 0$. Let F_i be its entrywise limit. Fatou's lemma for the nonnegative weighted absolute entries gives

$$\begin{aligned} &R_{y_i}(Q_{N,i} - F_i) \\ &\le AT(k,i)^d (a_N/a_i)^c. \end{aligned}$$

The horizon strip bound and the polynomial cap on the upper-bound column $P.\Omega$ similarly give

$$R_{y_i}(F_i) \le B' T(k,i)^l.$$

The limits obtained with different reporting levels k coincide by uniqueness of entrywise limits. Consequently one family F_i works for every k. On levels before e it can be defined as zero because the conclusion imposes no requirement there.

5.2 Polynomial majorants and decay

Let t_i be increasing scale times, with $\Delta_i = t_{i+1} - t_i \ge \log 3$, and put $\tau_i = 2 + Z + |t_i - t_k|$, where $Z \ge 0$. Then $\tau_i \ge 2$ and $\tau_{i+1} \le \tau_i + \Delta_i$. Suppose

$$f_i \le e^{-\theta\Delta_i} f_{i+1} + g_i \quad (i < N), \qquad f_N \le g_N, \qquad 0 \le g_i \le A\tau_i^d,$$

with $f_i \ge 0$, $A \ge 0$, $d \in \mathbb{N}$, and $\theta > 0$, uniformly in the schedules and terminal indices. Define

$$K = \frac{2(d+1)}{\theta}, \qquad q = e^{-(\theta/2)\log 3} \in (0,1), \qquad B = \frac{A}{1-q}.$$

Since $K + \tau_i \ge K$, the inequality $1 + u \le e^u$ gives

$$K + \tau_{i+1} \le (K + \tau_i) e^{\Delta_i / K}.$$

Thus $d/K \le \theta/2$ implies

$$e^{-\theta\Delta_i}(K + \tau_{i+1})^d \le e^{-(\theta/2)\Delta_i}(K + \tau_i)^d \le q(K + \tau_i)^d.$$

Backward induction proves $f_i \le B(K + \tau_i)^d$. At N, use $A \le B$. At an earlier index,

$$f_i \le (Bq + A)(K + \tau_i)^d = B(K + \tau_i)^d.$$

Finally $K + \tau_i \le (K+1)\tau_i$, so $f_i \le B(K+1)^d \tau_i^d$. This includes $d = 0$ and $A = 0$.

For the decaying version, suppose $g_i e^{c(t_N - t_i)}$ has a uniform nonnegative polynomial majorant, with $c > 0$. Put $\mu = \min(c, \theta/2) > 0$, and multiply the recurrence by $e^{\mu(t_N - t_i)}$. The new unknown has contraction exponent $\theta - \mu > 0$; its forcing is polynomially bounded because $\mu \le c$. The preceding induction gives

$$f_i \le A' \tau_i^{d'} e^{-\mu(t_N - t_i)}$$

for uniform $A' \ge 0$, $d' \in \mathbb{N}$. For the lattice schedule $e^{-(t_N - t_i)} = a_N/a_i$. A forcing supported only at N automatically satisfies the weighted hypothesis, because its exponential factor is one on its support. This proves the decay estimate used in the flat overlap construction. The same argument extends to a well-founded collection of grades. At a grade j, permit polynomial forcing in finitely many strictly earlier grades, at the current and next scale. Sums, products, and integer powers preserve polynomial bounds. Next-scale ages are polynomially controlled when the block lengths are polynomially bounded in the current age, because $\tau_{i+1} \le \tau_i + \log(\mathrm{block}(i)) \le \tau_i + \mathrm{block}(i)$. Induction on the grade therefore supplies the forcing bound needed for the scalar recursion. For differences, include the absolute bounds from both histories, including the current grade, in the forcing polynomial; permit difference bounds only at strictly earlier grades. A finite sum of exponentially decaying polynomial bounds retains such a bound with the minimum positive exponent, and multiplication by polynomial absolute bounds preserves it. The scalar weighted recursion then closes the induction. This recurrence is applied to the one-step estimates for the actual histories constructed in the nonlinear and history-convergence chapters.

5.3 Independence of the evaluated quadratic kernel

Which quadratic object is independent The object compared here is the evaluated quadratic form on link potentials. The auxiliary two-form metric can change on directions not tested by the curl without changing that form. The word "shadow" in this particular calculation names the displayed modification of clip parameters in a second finite linear history. It does not introduce a source–readout framework into this companion paper. Equality is proved by the stationary matrix identities and backward induction from the common terminal form; no equality of the full nonlinear actions is inferred from it.

The horizon's terminal quadratic form is the same fixed bare quadratic form for every admissible horizon. Its one-step stationary action identity is, on each finite torus,

$$Q_{\mathrm{coarse}} = T^{\mathsf{T}} Q_{\mathrm{fine}} T + b(I - \mathrm{probe}\, T)^{\mathsf{T}}(I - \mathrm{probe}\, T),$$

where $b > 0$ is the blocking penalty and T is the extracted stationary response. Here is the exact comparison algebra. Write P for the common probe, Q,R for two symmetric fine precisions, A,B for their symmetric coarse precisions, and T,W for the responses. The primal and dual equations are

$$QT = P^{\mathsf{T}} A, RW = P^{\mathsf{T}} B, A = b(I - PT), B = b(I - PW).$$

Taking the transpose of $RW = P^{\mathsf{T}} B$ gives $W^{\mathsf{T}} R = BP$. Consequently

$$\begin{aligned} W^{\mathsf{T}}(R - Q)T &= BPT - (PW)^{\mathsf{T}} A \\ &= B(I - A/b) - (I - B/b)A = B - A. \end{aligned}$$

The mixed terms cancel in their displayed order; no commutation of A and B is required. Thus equality of the fine quadratic forms implies equality of the coarse forms even if the auxiliary factorizations differ. Backward induction from the common terminal form gives equality at every level

$i \le N$ on every scheduled torus. The periodization uniqueness argument above then gives equality of the summable cover kernels. This proves the quadratic-history comparison. The harmonic and scaled-exterior sections construct the responses and derive their primal and dual equations.

The comparison construction replaces each clip's two numerical factors by $\gamma(j)$ and $\gamma(j)/32$, with $\gamma(j) > 0$, retaining the relaxation, radius and slow selector. It retains the schedule, patches, smoothing indices and blocking penalty used by the block operator. The same terminal and backward comparison therefore applies to the original and modified horizons. Their evaluated quadratic cover kernels agree at every finite cutoff. This concerns the quadratic action on link potentials and does not identify higher-order horizon data.

Here is the entrance argument, called `entrance_gate` in the source. First choose the retained-support linear construction for the fixed kernel choice; it supplies a positive clip curve γ and a length demand n. Apply the linear horizon construction to the numerical schedule with its rearranged clip profile. For its price choice X, set $E(j) = \max\{\text{tickets}(X, j), n(j)\}$. The condition $E(J(i+1)) \le \text{block}(i)$ pays both demands, including auxiliary steps before the admitting tail. The delay at e supplies the spacing hypotheses proved next. Apply the flat convergence estimate to the modified horizon to obtain one family F_i, its polynomial absolute bound and its decaying difference bound. The exact comparison just proved, called `same_shadow`, identifies every original finite-cutoff quadratic output with its modified output. Equality of those sequences transfers their limits and difference bounds to every admissible original horizon. This proves the entrance statement once the retained-support construction is supplied.

5.4 Choosing the widths before the block lengths

For the fixed price table P, choose $w(j) = \min\{1, P.\alpha(j), P.\alpha(j+1)\} > 0$. Let $h(j,u)$ be the allotted bound for adjacent smoothing indices j, u, and put

$$B(j,u) = |\text{Faces}|^2[2/w(j)]h(j,u)^2,$$
$$E_P(j) = \left\lceil 2/w(j) + B(j,j)^2 + B(j,j-1)^2 \right\rceil.$$

The smoothing index stays fixed or increases by one at each forward step. Thus, if $j = J(i+1)$ and $u = J(i)$, then $u = j$ or $u = j-1$. Requiring $L = \text{block}(i) \ge E_P(j)$ gives $B(j,u)^2 \le L$ and $L \ge 2/w(j)$. The one-step price $B(j,u)/L$ is at most $L^{-1/2}$. Moreover $2w(u) \le 2 \le Lw(j)$. These give contraction and width crossing with $\theta = 1/2$. The same adjacency ensures that $w(u)$ is no greater than either analytic width required at the crossing. All thresholds depend on the previously chosen prices, rather than on the subsequently chosen L.

For the age bound, put $Z(k) = 3 + \log \widehat{\beta}_k$, where $\widehat{\beta}_k > 1$, and let $s > 0$ be the coupling slope. Its affine dependence on scale gives $\widehat{\beta}_i \le \widehat{\beta}_k + s|t_i - t_k|$. Divide by $\widehat{\beta}_k > 1$ and use $\log v \le v - 1$ to obtain

$$\begin{aligned}\log(\widehat{\beta}_i) &\le \log(\widehat{\beta}_k) + s|t_i - t_k| \\ &\le (1+s)[2 + Z(k) + |t_i - t_k|].\end{aligned}$$

The delay condition $P.\Omega(J(i)) \le \log(\widehat{\beta}_i)$ on $i \ge e$ now supplies a polynomial bound of degree one. Together with the previous inequalities this establishes every spacing condition used in the flat convergence argument.

5.5 Neutral packet propagation

The predictor difference estimate used here follows directly from the one-step bound included in its Atlas data. Suppose its weighted forward difference in each fine direction has row bound c/L^{d+1}. Join x to z by a nearest-neighbor path of length $a = |z - x|_1$ which never exceeds distance a from x. Such a path is obtained by changing each integer coordinate toward its target one unit at a time. The scaled-distance triangle inequality gives

$$\exp(\sigma\delta_L(x, y)) \le \exp(\sigma a/L) \exp(\sigma\delta_L(v, y))$$

for each intermediate site v. Thus each forward or backward difference along that path has weighted row sum, measured from x, at most $\exp(\sigma a/L)c/L^{d+1}$. Summing the a telescoping differences and using the triangle inequality yields

$$\sum_y |U((z, i), (y, j)) - U((x, i), (y, j))|$$
$$\exp(\sigma\delta_L(x, y))$$
$$\le a \exp(\sigma a/L)c/L^{d+1}.$$

The centered estimate below is this bound with $d = 2$, followed by spending half the packet's input exponential width on the factor a. The actual response and its one-step bound are constructed in the Atlas response and exterior-lifting arguments that follow.

For a two-point curvature packet let $K(x, v)$ be its kernel, where each argument includes a site and an orientation. Neutrality means $\sum_z K(x, (z, a)) = 0$ for each fixed x and orientation a separately. If T, U are the response kernels used to carry the packet from the fine level to the coarse level, its pushed kernel is

$$K'(y, z) = \sum_{x,v} T(x, y)K(x, v)U(v, z).$$

All rearrangements below use absolute convergence. It follows from the packet's weighted row bound and the row and column bounds on T,U: first sum the nonnegative absolute majorant in the output variables, then in the input variables. The resulting finite product of bounds dominates the full multiple series. Thus both iterated sums and regrouping by the finite orientation sets are justified.

Neutrality removes the spatially constant part of the right response. For fixed x, replace $U(v, z)$ by $U(v, z) - U((x.\text{site}, v.\text{orientation}), z)$. The removed sum vanishes separately for each orientation. If $a = \text{dist}(x.\text{site}, v.\text{site})$, the spatial difference estimate gains a/L relative to the ordinary row estimate. For positive input width τ and output width σ with $2\sigma \le L\tau$,

$$a \exp(\sigma a/L) \le (2/\tau) \exp(\tau a).$$

Indeed $\sigma/L \le \tau/2$, and $a \le (2/\tau) \exp(\tau a/2)$, the latter following from $u \le \exp(u)$ for $u \ge 0$. The centered right row therefore costs $|\beta|(2/\tau)c/L^3$ for a degree-two predictor of cost c. The held left column costs $|\alpha|qL^2$. Multiplication, followed by the input packet bound p, gives

$$\text{folded}_\sigma(K') \le [|\alpha|q|\beta|(2/\tau)c/L]p.$$

This is the precise neutralCoin factor used in the flat recursion. The gain L^{-1} comes from centering before taking absolute values. The secured pilot and wire bounds in this chapter provide the predictor spatial difference estimate.

To preserve neutrality after propagation, require that $\sum_z U((v,a),(z,b))$ be independent of the input site v, for every pair of orientations a, b. Summing $K'(y,(z,b))$ over z gives a linear combination of the zero orientation-wise sums of K. This is the source's constant-sum condition `travels`. Translation covariance follows by shifting both fine summation variables by L times the coarse shift and using the covariance of K, T, U. Finally, periodization of K' is exactly $T^{\mathsf{T}}_{\mathrm{torus}} K_{\mathrm{torus}} U_{\mathrm{torus}}$: absolute convergence permits splitting fine sites into residue classes, and covariance identifies each periodized factor. The mixed stationary identity then makes K' represent the coarse quadratic difference. This proves the algebra, regrouping and price multiplication called `toward`, once the response bounds and constant-curvature transport have been constructed.

5.6 Positive three-form correction

The linear flight construction uses a matrix symbol on the whole three-form coordinate space of the form

$$A = V + lI - mS^*S,$$

Here $m \geq 0$ is the fine amendment parameter and l the coarse parameter. On real momentum the symbols are Hermitian. The analytic estimates give $V \geq -a^2 CI$ and $S^*S \leq a^2 BL^{-2} I$, where a is a fixed orientation-count constant, $C, B > 0$ depend only on the smoothing pair and penalty, and L is the block length. Choose $l \geq a^2 C + 2$ and require $ma^2 B \leq L^2$. Then, for every vector x,

$$\langle x, Ax\rangle \geq [-a^2 C + l - ma^2 B/L^2]\|x\|^2 \geq \|x\|^2.$$

This is positivity on the full three-form space. No restriction to the image of the exterior derivative is used to obtain it. For the corrected form $R = F * AF$, a uniform inverse bound $\|x\|^2 \leq W\|Fx\|^2$ yields $R \geq W^{-1} I$. The source obtains W from a uniformly bounded right inverse of F, which in these square finite matrices is also a left inverse. An independent analytic row estimate supplies the upper bound. This explains why the coarse amendment threshold is chosen before the fine block length, while the subsequent length condition absorbs the already fixed fine amendment.

The rebased correction identity has the following exact content. Let X,Y be the fine and coarse amended inverse metrics, B the coarse amended metric, T the installed inverse response, and d_1, d_2 the coarse exterior derivatives. Then

$$B[Y - T^{\mathsf{T}} XT - b^{-1} d_1 d_1^{\mathsf{T}}]B = d_2^{\mathsf{T}} R d_2.$$

To derive this, write $Y = B_0 + lR_0^{\mathsf{T}} R_0$ and $X = b_0 + mF_0^{\mathsf{T}} F_0$. The unamended defect is $R_0^{\mathsf{T}} V R_0$ and the intertwining identity is $SR_0 = F_0 T$. Substitution makes the bracket $R_0^{\mathsf{T}}(V + lI - mS^{\mathsf{T}} S)R_0$. Since $R_0 B = F d_2$ and B is symmetric, the displayed formula follows with $R = F^{\mathsf{T}} AF$. The exterior lifts, regular symbols and analytic inverse estimates construct B_0, b_0, S, F with the bounds used here.

5.7 Roots for the bottom metric

Let the amended metric M have real-axis lower bound $b_0 I$, and choose the clip coefficient $0 < c \leq b_0/2$. Then $M - cI \geq (b_0/2)I$. Subtracting the constant row cI preserves the analytic strip bound. The positive analytic-root construction gives a real Laurent row F with $F^{\mathsf{T}} F = M - cI$ on every finite torus. Adjoin the constant block $\sqrt{c}\, I$. Its Gram matrix is exactly

$$(\sqrt{c}I, F)^{\mathsf{T}}(\sqrt{c}I, F) = cI + F^{\mathsf{T}}F = M.$$

The source chooses c's allowed upper cap as half a uniform positive lower metric bound, intersects the finitely many positive analytic widths, and increases the common cost to cover the metric, inverse, diagonal, dual diagonal, and root rows. The resulting bed satisfies the required finite-torus floor identities and bounds for every permitted train. This reduces the remaining root step to the analytic square-root construction, rather than assuming a factorization with the desired decay.

For completeness, the analytic square root can be constructed with one branch on a collar uniform over the entire family. Suppose a real Laurent matrix row Q has strip width $\eta > 0$ and entrywise weighted coefficient bound A, and its real symbols obey $vI \le Q(p) \le AI$ with $v > 0$. Increase A if necessary so $A \ge v$. On the real axis, unitary diagonalization gives

$$\|A^{-1}Q(p) - I\| \le 1 - v/A.$$

On the smaller strip $\eta/2$, differentiating the absolutely convergent Fourier series gives the entrywise bound for symbol differences with constant $D = A/(\eta - \eta/2)$. Let C be one plus the norm of the fixed identification from entrywise matrices to the matrix operator norm, and choose

$$\tau = \min(\eta/2, v/(2CD)) > 0.$$

For a complex momentum p in this collar, comparison with its real part gives $\|Q(p) - Q(\operatorname{Re} p)\| \le v/2$. Consequently $\|A^{-1}Q(p) - I\| \le 1 - v/(2A) < 1$ uniformly. The binomial series

$$P(p) = \sqrt{A} \sum_{n \ge 0} \binom{1/2}{n} (A^{-1}Q(p) - I)^n$$

defines a bounded holomorphic periodic matrix function on the collar. On the real axis, diagonalization reduces the series to the positive scalar square root, so $P(p)^* = P(p)$, $P(p)^2 = Q(p)$ and $\sqrt{v}\, I \le P(p) \le \sqrt{A}\, I$. The Laurent coefficients of Q are real, and conjugation at opposite real momenta preserves the selected branch. Fourier inversion, taking real parts of the coefficients, therefore recovers the same real-axis function.

Bounded holomorphy gives exponential Fourier decay by shifting each contour within the collar. Choose output width $e = \tau/3$, leaving another width e for summing coefficient moduli. In four dimensions the weighted sum is bounded by $K[2(1 + e)/e]^4$, where K bounds the entries of P on the collar. The constants depend on v, A, η and the fixed matrix size, not on the torus period. Discrete Fourier diagonalization gives $P_{\text{torus}}^{\mathsf{T}} = P_{\text{torus}}$ and $P_{\text{torus}}^{\mathsf{T}} P_{\text{torus}} = Q_{\text{torus}}$ on every torus. The invertible discrete Fourier transform transfers the real-axis symbol identity to the real-coordinate matrices. Hence factorization and analytic decay hold simultaneously at all periods, as required for the bottom metric.

5.8 Smoothing histories and the actual block average

A block average is an actual map on the fine lattice, with its carry across residue seams included. Its Fourier multiplier therefore has both a principal alias and nonzero aliases. The latter cannot be discarded before the stationary problem is solved. Their lower denominator bounds, proved next, are what make the filled principal formula and all the subsequent contour estimates uniform on the prescribed complex collar.

Smoothing mass and physical mass The mass of a smoothing train below is a positive integer: the product of its block lengths. It fixes the rescaling in its Fourier multiplier. It is unrelated to the positive regulator masses in the later resolvent calculation, and to a spectral gap above the vacuum. Its useful property is multiplicative growth, which makes the inverse-square contributions of earlier smoothing stages summable independently of the depth of the history.

For a positive odd integer M define the line average and its centered version by

$$R_M(z) = M^{-1} \sum_{u=0}^{M-1} \exp(iuz),$$
$$P_M(z) = e^{-i\lfloor M/2 \rfloor z} R_M(z).$$

Both filters equal one at $z = 0$. Summing the geometric series gives $MR_M(z)(e^{iz} - 1) = e^{iMz} - 1$, including the zeros of $e^{iz} - 1$. Writing an integer in $\{0, \dots, LM - 1\}$ as $Mu + v$ yields $R_L(Mz)R_M(z) = R_{LM}(z)$. For odd L, M,

$$M\lfloor L/2 \rfloor + \lfloor M/2 \rfloor = \lfloor LM/2 \rfloor,$$

so $P_L(Mz)P_M(z) = P_{LM}(z)$. Reflection of the centered sum gives $P_M(-z) = P_M(z)$, and conjugation gives $\overline{P_M(z)} = P_M(-\overline{z})$. All these identities hold for entire functions.

Set $A_M(p) = \prod_{a=1}^{4} P_M(p_a)$. A train is a finite chain of patches whose smoothing indices stay fixed or decrease by one at each backward step. Its mass is the product of its odd block lengths. A terminal train of index j has mass and smoothing function both equal to one. If v has index J and a patch takes it to index j with $j \le J \le j + 1$ and block L, the new mass is $L \,\mathrm{mass}(v)$ and its smoothing function is

$$\mathrm{smooth}_{\mathrm{new}}(p) = A_L(\mathrm{mass}(v)p)^J \,\mathrm{smooth}_v(p).$$

Equivalently,

$$\mathrm{smooth}_v = A^j_{\mathrm{mass}(v)} \,\mathrm{extra}_v, \qquad \mathrm{extra}_{\mathrm{new}} = A^{J-j}_{L\,\mathrm{mass}(v)} \,\mathrm{extra}_v \,.$$

This separates the fixed-stage power from the extra smoothing introduced when the stage changes.

The extra product is bounded independently of chain length. On $|\operatorname{Im} p_a| < \zeta/H$, direct estimation of the finite exponential sum gives $|A_M(p)| \le e^{4\zeta M/H}$. Successive masses increase by factors of at least three, and $J - j$ is zero or one. Induction, equivalently summing the resulting geometric series, gives

$$|\,\mathrm{extra}_v(p)| \le e^{8\zeta\,\mathrm{mass}(v)/H}, \qquad \mathrm{mass}(v) \le H.$$

On the real principal region, $|P_M(p_a/H)| \ge e^{-cM/H}$ gives

$$|\,\mathrm{extra}_v(p/H)| \ge e^{-8c\,\mathrm{mass}(v)/H}.$$

Consequently $|\,\mathrm{smooth}_v(p/\,\mathrm{mass}(v))|$ has a positive lower bound depending only on j, ζ, c. The normalized finite-sum proof is given in (QF1)–(QF4).

At $q_l = (p + 2\pi\,\mathrm{lift}(l))/\,\mathrm{mass}(v)$, the principal factor $A_{\mathrm{mass}(v)}$ has inverse-linear decay in each centered integer alias coordinate, while the extra factor has the preceding upper bound. Hence

$$|\,\mathrm{smooth}_v(q_l)| \le C_j \prod_a (1 + |\,\mathrm{lift}(l)_a|)^{-j}.$$

The equivalent source normalization uses $\mathrm{up}(n) = 1 + |n|$. This is the decay used in the high-alias sums. The nonzero-alias reciprocal estimate supplies the required bound for R_M.

These are the actual patch weights. One centered box average of Fourier modes is A_L; J independent box rolls give A_L^J. Different rolls can have the same total displacement, and the patch weight retains their multiplicities. A positive line path in direction a has L edges at displacements $0, \dots, L-1$, with Fourier sum $LR_L(q_a)$. Thus its directional patch mean has symbol

$$LA_L(q)^J R_L(q_a).$$

Applying the defined mean to every discrete Fourier basis vector proves the source identity `patch_service`. The factor L remains because the path integral counts edges. Wrapped paths on a small torus still count every edge occurrence; no distinct-visit assumption is used.

5.9 Uniform reciprocals at the nonzero aliases

A real lower bound is converted to a complex lower bound before inversion. Suppose f_ξ is holomorphic with $|f_\xi| \le Cw_\xi$ on a common complex cube of width η, and $|f_\xi| \ge cw_\xi$ on its real cube of width θ, where $C, c > 0$ and $w_\xi > 0$. Set

$$r = \min(\eta/4, \theta, c\eta/(8C)).$$

For p in the width-r cube, put $y = \operatorname{Re} p$. The ball of radius $\eta/2$ about y stays in the original cube. Cauchy's estimate along the line from y to p, using $|f_\xi(z) - f_\xi(y)| \le 2Cw_\xi$ on that ball, gives

$$|f_\xi(p) - f_\xi(y)| \le 4Cw_\xi/\eta|p - y| \le (c/2)w_\xi.$$

Thus $|f_\xi(p)| \ge (c/2)w_\xi$. Its reciprocal is holomorphic on this common cube and bounded by $2/(cw_\xi)$. The width r is independent of ξ. This is the required reciprocal estimate; real positivity alone would not prove complex nonvanishing.

In particular, define

$$\lambda_M(p) = \sum_a [M\chi(p_a/M)][M\chi(-p_a/M)].$$

On real momenta this is $\sum_a |M\chi(p_a/M)|^2$. The bounds for χ give

$$c_-^2|p|^2 \le \lambda_M(p) \le c_+^2|p|^2$$

on the real cube, uniformly in M. For block L and alias l, put

$$B_{M,L,l}(p) = L^2\lambda_M((p + 2\pi \operatorname{lift}(l))/L).$$

For $l \neq 0$, choose a coordinate attaining $\operatorname{grade}(l) = 1 + \max_a |\operatorname{lift}(l)_a|$. The lower alias estimate gives $B_{M,L,l} \ge (c_-\zeta)^2 \operatorname{grade}(l)^2$ on the real cube. The holomorphic upper bound is $C \operatorname{grade}(l)^2$. Apply the reciprocal estimate with $w_l = \operatorname{grade}(l)^2$ to obtain a fixed complex collar and $|B_{M,L,l}^{-1}| \le C' \operatorname{grade}(l)^{-2}$, uniformly in $M, L, l \neq 0$. The principal alias instead satisfies the exact identity $B_{M,L,0} = \lambda_{LM}$ and is not inverted on a collar containing zero.

Define

$$\operatorname{mix}(l) = \prod_a (1 + |\operatorname{lift}(l)_a|), \qquad \operatorname{decay}_j(l) = \operatorname{mix}(l)^{-j}.$$

Since all factors are at least one, $\operatorname{grade}(l) \le \operatorname{mix}(l)$. Therefore $\operatorname{grade}(l)^u \operatorname{decay}_{j+u}(l) \le \operatorname{decay}_j(l)$. Finally,

$$\sum_{l \bmod L} \operatorname{decay}_2(l)$$
$$\le \left[\sum_{n\in\mathbb{Z}} (1+|n|)^{-2}\right]^4 < \infty,$$

uniformly in L. The centered representatives are injective, so each one-dimensional finite sum is bounded by the full integer series. The finite four-dimensional sum factors into the product of those four finite sums. This proves the decay summation used in the earlier feed and high-alias estimates, without paying a factor equal to the number L^4 of residues.

For the residue-column estimates, a word u of fine differences contributes at most $CL^{-|u|}\operatorname{grade}(l)^{|u|}$. The scale factor contributes L^{-e}. The residue phase is uniformly bounded because residue coordinates are of size at most L and $|\operatorname{Im} q_a| \le \zeta/L$. The symbol contributes $\operatorname{decay}_j(l)$. Absorption and summation therefore bound the residue sum by $C'L^{-(e+|u|)}$ whenever $j - |u| \ge 2$. Reciprocal shifts permute aliases, giving periodicity; conjugation and reflection give real Fourier coefficients. Fourier estimates on a narrower collar yield the priced fine-difference rows. The coefficient and symmetry arguments below verify these symbol bounds for the actual potential and curvature responses, called pilot and wire in the source.

The scalar bounds use $\chi(z) = e^{iz} - 1$. Choose $c_-, c_+ > 0$ so that $c_-|z| \le |\chi(z)| \le c_+|z|$ on $|z| \le 3\pi/2$. Such constants exist because $\chi(z)/z$ extends continuously with nonzero value i at zero, and has no zero on that disk: $e^{iz} = 1$ requires a real multiple of 2π. Compactness bounds this nonvanishing quotient above and below.

If $M|z| \le 3\pi/2$, the geometric-sum identity gives $|R_M(z)| \ge c_-/c_+$ for $z \ne 0$, and $R_M(0) = 1$. Put $a_0 = \min\{1, c_-/c_+\}$. Centering leaves the modulus unchanged for real z, so $|P_M(z)| \ge a_0$. If also $H|z| \le 3\pi/2$ and $M \le H$, put $x = M/H$. Estimating each centered monomial minus one gives $|P_M(z) - 1| \le kx$, where $k = c_+(3\pi/2)$. For $y = |P_M(z)|$ we have $y \ge a_0$ and $1 \le y + kx$. Since $e^{kx/a_0} \ge 1 + kx/a_0$,

$$y \exp(kx/a_0) \ge y + kx \ge 1.$$

It follows that $|P_M(z)| \ge e^{-(k/a_0)M/H}$. The coefficient in front of this exponential is exactly one, which is essential when the estimate is composed along arbitrarily long smoothing trains.

For a complex alias $q = (p + 2\pi l)/L$ with centered nonzero integer l, the selected cube gives $L|q| \ge \zeta \operatorname{up}(l)$ and $L|\operatorname{Im} q| \le \zeta$, where $\zeta = \pi/4$ and $\operatorname{up}(l) = 1 + |l|$. These follow by separating $2\pi l$ from $|\operatorname{Re} p| \le \pi + \zeta$. Also $|q| \le 3\pi/2$ for odd $L \ge 3$. The geometric-sum identity and $\chi(p + 2\pi l) = \chi(p)$ give

$$L|R_L(q)|c_-|q| \le |\chi(p)| \le c_+(3\pi/2),$$

so $|R_L(q)| \le C/\operatorname{up}(l)$. For $l = 0$, the finite-sum estimate gives $|R_L(q)| \le e^{\zeta}$. Centering changes this by at most e^{ζ}, yielding the same inverse-alias bound for P_L. Coordinate products give the four-dimensional estimate above. The remaining alias estimates concern the Laplacian denominator and the sum of the decay weights.

To install a response kernel, split a fine site into its centered residue r modulo L and a coarse displacement. Advancing by e_a changes the residue to $r + e_a$ and adds a coarse displacement e_a exactly when $\operatorname{lift}(r)_a = \lfloor L/2 \rfloor$. Denote that added displacement by $c_a(r)$. Then

$$\operatorname{lift}(r + e_a) + Lc_a(r) = \operatorname{lift}(r) + e_a.$$

In the residue Fourier sum this identity gives a factor $\exp(ip \cdot c_a(r))$ on the advanced residue. Subtracting the original residue value is exactly multiplication of each alias term by $\exp(iq_a)$-1. Repeating this operation gives a product of those factors for a word of fine differences. Fourier inversion therefore identifies the actual carried differences of the cross row with the separately priced difference symbols. This identity follows from the real sums, rather than treating a residue wrap as an ordinary shift within the same copy.

For a fine-difference word u, the symbol estimates give a kernel bound $CL^{-(e+|u|)}$, with a positive width and C uniform over all transitions at fixed smoothing indices, whenever $|u| \le j-2$. There are finitely many such direction words. Intersect their widths and enlarge their costs to obtain one table for the row and all required differences. The coefficient estimates and the inverse-third-power estimate below supply the analytic symbol bound; the residue calculation identifies the actual carried differences to which it applies.

5.10 Laurent exterior division

Write $q_s(p) = e^{ip_s} - 1$ in the lattice Fourier convention, and let $\delta_r(p)$ be exterior multiplication by q on degree-r forms. We divide a cycle without assigning a singular quotient at $q_s = 0$. For a lattice monomial e_x, $x \in \mathbb{Z}^4$, let E_s replace its s-coordinate by zero. Define $J_s e_x$ by summing monomials with s-coordinates $0, \dots, x_s - 1$ when $x_s \ge 0$, and minus the sum over $x_s, \dots, -1$ when $x_s < 0$, keeping other coordinates fixed. Telescoping gives

$$q_s(p)(J_s e_x)(p) = e_x(p) - e_x(p|_{p_s=0}).$$

If $\sum_x |a_x| e^{\eta|x|_1} \le C$, then E_s preserves that bound. The sum defining J_s has $|x_s|$ terms, each of length at most $|x|_1$. For $\theta = \eta/2$,

$$|x_s| \exp(\theta|x|_1) \le \theta^{-1} \exp(\eta|x|_1).$$

Therefore J_s maps the width-η bound C to the width-$\eta/2$ bound $2\mathrm{C}/\eta$. This proves absolute convergence of the transformed rows and licenses termwise telescoping. In particular $q_s J_s A = A - E_s A$ as an identity of real-momentum symbols, including the hyperplane $q_s = 0$.

Suppose A maps degree-$(r+1)$ forms to a fixed coordinate space and $A\delta_r = 0$. Let $L_r(s)$ be exterior multiplication by the s-th coordinate vector. Define the relief row $A_s = J_s A L_{r+1}(s)^{\mathsf{T}}$. Obtain the residual R_s by evaluating A at $p_s = 0$ and deleting columns whose orientation contains s. Equivalently,

$$R_s = (E_s A)[I - L_r(s) L_r(s)^{\mathsf{T}}].$$

The exterior contraction identity

$$L_{r+1}(s)^{\mathsf{T}} \delta_{r+1} = q_s I - \delta_r L_r(s)^{\mathsf{T}}$$

and $\delta_r(p) = \delta_r(p|_{p_s=0}) + q_s L_r(s)$ imply

$$A_s \delta_{r+1} + R_s = A.$$

First prove this where $q_s \neq 0$ by multiplying by q_s, using $q_s J_s A = A - E_s A$, and cancelling that nonzero scalar. The symbols are continuous, so perturbing p_s through nonzero real values proves the equality also at $p_s = 0$. Furthermore $R_s \delta_r = 0$ because $\delta_{r+1}\delta_r = 0$. Thus the residual remains a cycle and one may repeat the operation.

Apply the four axes successively to the preceding residual and add the relief rows. Once an orientation contains a processed axis its column is zero and stays zero. In positive degree every orientation contains an axis. After four operations the residual is zero and the sum B satisfies $B\delta_{r+1} = A$. The explicit finite width losses preserve a positive analytic width, so B retains exponential coefficient decay uniformly for a common input bound. Fourier evaluation yields $B_{\mathrm{torus}} d_{r+1} = A_{\mathrm{torus}}$ at every period. Conversely, $A_{\mathrm{torus}} d_r = 0$ at all periods implies the symbol cycle identity by Fourier uniqueness. That reverse implication uses all periods.

Transposing gives the corresponding left lift. A two-sided lift requires a correction because an arbitrary left lift need not remain a right cycle. Suppose A is a two-form row with $Ad_1 = 0$ and $d_1^{\mathsf{T}} A = 0$ on all periods. Choose W by a left lift so $d_2^{\mathsf{T}} W = A$. Put $F = W d_1$; then $d_2^{\mathsf{T}} F = 0$. Lift again on the left to obtain a degree-four row T with $d_3^{\mathsf{T}} T = F$. Now $T d_0 = 0$: its left derivative vanishes since $d_3^{\mathsf{T}} T d_0 = W d_1 d_0 = 0$, and a left cycle in top degree four vanishes by the same lifting theorem, whose next form space in dimension four is zero. Right-lift T to obtain X with $X d_1 = T$. Define $K = W - d_3^{\mathsf{T}} X$. Then $K d_1 = 0$ and $d_2^{\mathsf{T}} K = A$, using $d_3 d_2 = 0$. Right-lift K to obtain B with $B d_2 = K$. Thus $d_2^{\mathsf{T}} B d_2 = A$. Every step is a finite composition of the preceding analytic lifts, exterior derivatives, sums and transposes, so the same familywise exponential decay statement holds. This proves the exterior factorization used for the fund; the refined block-length power in the sale estimate is additional quantitative information.

5.11 The regular two-form symbols

Fix a smoothing train v with scale parameter M. Set $z_a(p) = \mathrm{seedQ}_M(p, a)\,\mathrm{seedQ}_M(-p, a)$ and $\lambda_M(p) = \sum_a z_a(p)$. These z_a are nonnegative at real momentum. The positive smoothing factors are $h_0 = \mathrm{halo}(v)$ and $l_a = \mathrm{lid}(v, a)$. Let $r_a = \mathrm{drain}(b, v, a)$ be the regular remainder and put

$$g_a = l_a + \lambda_M r_a,$$
$$h = h_0^{-1} - \sum_a \frac{z_a^{(1)} r_a}{l_a g_a}.$$

Here $z_a^{(1)}$ is the analogous factor for $M = 1$. The exact identity $z_a^{(1)} h_0 = z_a l_a$, after taking common denominators, gives

$$\lambda_M h = \sum_a z_a^{(1)} / g_a. \tag{2}$$

The two regular symbols are diagonal on the oriented two-forms. At orientation $\{a, c\}$, the symbol called vault has entry $h g_a g_c$ and the symbol called pulse has its reciprocal; off-diagonal entries vanish. At $p = 0$, every g_a and h equals one. These are the values of the same analytic expressions at zero, with no separately substituted harmonic value.

The positivity proof uses $l_a > 0$ and $r_a \geq 0$ on the real cube, so $g_a \geq l_a > 0$. For $p \neq 0$, (2) gives $h \geq 0$. More quantitatively, if $g_a \leq C$ and the comparison constants for λ_M and λ_1 with $|p|^2$ are c_-^2 and c_+^2, then (2) gives

$$h \geq \frac{c_-^2}{C c_+^2}.$$

At $p = 0$, $h = 1$, giving the joint lower bound $\min(1, c_-^2/(C c_+^2))$. Hence h never vanishes. Pulse and vault are therefore positive, real diagonal mutual inverses throughout the real cube. Uniform bounds and positive lower bounds for g_a, h give analytic narrower-collar bounds for their reciprocals

and their finite products. Taking Fourier coefficients gives the exponentially decaying metric and inverse rows used above.

Away from zero, set $D_a = g_a/\lambda_M$ and $\rho = \sum_a z_a^{(1)}/D_a$. Equation (2) gives $\rho = \lambda_M^2 h$. Thus

$$\text{vault} = \rho\,\text{mask}_2(D_a), \qquad \text{pulse} = \rho^{-1}\,\text{mask}_2(D_a^{-1}),$$

where mask_2 takes the pairwise products of diagonal entries on two-form orientations. Periodicity of D_a and ρ proves agreement under reciprocal-lattice shifts on overlapping real cubes. The regular g_a, h expressions give the values through the apparent pole. At the terminal train, $D_a = \lambda_1^{-1}$ and $\rho = \lambda_1^2$, so both two-form symbols are the identity, as required for the bare Wilson quadratic form.

The centered-average, nonzero-alias and cumulative-filter estimates in this chapter provide the smoothing bounds, regular remainder bounds and scale-uniform comparison of λ_M used in (2).

The depth-independent analytic estimate works as follows. Each earlier train contributes with weight mass^{-2}. A hop of block length L multiplies its mass by L, dividing the sum of old weights by L^2; the new list adds one train of mass one. Hence, if $t(v)$ is the sum of inverse-square masses over that list,

$$t(\text{top}) = 0, \qquad t(\text{hop}(v)) = 1 + L^{-2}t(v).$$

Since $L \geq 3$, induction gives $t(v) \leq 2$, uniformly in depth. Each train's actual contribution has analytic bound $C\,\text{mass}^{-2}$, so the finite list has bound $2C$. This controls the function called `levelNoise`; an unweighted contribution is not accumulated at every renormalization step.

At high aliases the smoothing factor has decay indexed by j, while the inverse nonzero-alias divisor costs an inverse squared grade. Their product is dominated by the summable reference decay of order two when $j \geq 4$. Its sum is independent of the alias side length. The same argument gives the inverse-square-mass bound for each contribution. Combining the high-alias sum with b^{-1} levelNoise bounds drain. The alias lower bounds, inverse-divisor bound and reference decay sum used here were proved above; their scalar assembly is also displayed in the cumulative-filter construction.

The factor lid is the product of its smoothing expression at p/M and $-p/M$. A uniform positive smoothing lower bound, together with the positive lower line-filter bound, makes lid uniformly positive. The same holds for halo. The bound on drain and the bound on λ_M bound every g_a above. Equation (2) then bounds h below. Only halo, lid and g_a occur in the rational denominators defining h. Their positive lower bounds permit a common smaller holomorphic collar and analytic bounds for h. No bound for the singular quotient $1/\lambda_M$ at zero is required.

5.12 Amendment and preservation of the potential precision

Let H be the unamended two-form metric and $B = H^{-1}$. Both are symmetric positive matrices on each finite torus. Put $X = d_2 B$ and, for $l \geq 0$, define

$$A = B + lX^{\mathsf{T}}X, \qquad G = A^{-1}, \qquad R = I + ld_2Bd_2^{\mathsf{T}}, \qquad D = R^{-1}.$$

The source names these rows pledge, amend, bail and treble, respectively. Positivity gives $A > 0$ and $R \geq I$, so their analytic inverses exist on uniform narrower strips.

Since $d_2 d_1 = 0$, we have $XHd_1 = d_2BHd_1 = 0$. Hence $AHd_1 = d_1$ and $Gd_1 = Hd_1$, so

$$d_1^{\mathsf{T}}Gd_1 = d_1^{\mathsf{T}}Hd_1.$$

The amendment changes the two-form extension while preserving the potential precision exactly. Also $RX = d_2 A$; multiply by D on the left and G on the right to obtain $XG = Dd_2$. These are the rebasing identities used above. At zero momentum $d_2 = 0$, and $B(0) = I$ gives $G(0) = I$.

Uniform inversion uses the fixed orientation dimension, rather than the torus site count. If the real symbols satisfy $vI \le A(p) \le CI$, then diagonalization gives $|\det A(p)| \ge v^s$ in dimension s and $C^{-1}I \le A(p)^{-1} \le v^{-1}I$. On a common narrower collar the determinant stays separated from zero: analytic coefficient bounds control its change from the real slice, and the collar can be reduced to lose at most half the positive bound. Adjugate divided by determinant gives uniformly bounded holomorphic inverse symbols. Periodicity and real-coefficient symmetry survive. Fourier coefficients decay exponentially, and discrete Fourier evaluation gives the inverse matrices at every period. This proves the inversion needed for G and D.

The same Fourier calculation transfers all real-axis spectral inequalities to finite tori. The discrete Fourier matrix has Gram matrix $n^4 I$; that positive scalar multiplies both the quadratic form and the squared norm and cancels. Thus a bound uniform in momentum gives the identical bound on every finite torus, including side one.

5.13 Constant-curvature transport

For a translation-invariant Laurent row, its total spatial sum is its zero-momentum symbol and is independent of the starting site. For an installed cross row, the starting site also has a residue modulo the block length. The required independence is therefore a separate assertion. For the specified duo response, its zero-momentum entry at each residue z is $\mathrm{toll}(L, 2)$ times the identity in face orientations. In the alias formula all nonzero aliases vanish at the origin; the zero alias equals the normalized identity. Consequently the spatial sum of each orientation column is independent of that residue as well.

The property survives convolution. If the spatial sums of S, T are the orientation matrices s, t, absolute convergence gives total sum st for their convolution. The fine inverse and coarse metric are translation-invariant rows, and the intermediate curvature response has the constant-residue property just proved. Their sandwich therefore has constant spatial sums. This is the condition called `travels_route`, required to preserve packet neutrality. The regular curvature-response formulas and harmonic values establish the alias vanishing and normalization used here.

5.14 Exterior factorization of the unamended defect

Exterior factorization separates the part of a two-form kernel seen by curls from its complementary directions. The calculation is performed in Laurent symbols before evaluating on a torus, so the same algebraic identity holds at all permitted periods. The zero Fourier mode is then restored by its explicit harmonic equation. This order avoids claiming that division on the punctured momentum domain alone determines the harmonic fibre.

From symbols to a finite history The exterior factorization supplies the auxiliary rows needed by the adjacent stationary problem. Those rows are then evaluated at every discrete torus momentum. Invertibility of the finite Fourier transform turns the symbol identities into actual matrix identities, including the separately checked zero mode. The subsequent assembly labels such as "Present" and "H.Suffices" record precisely this collection of identities and purchased bounds for a scheduled horizon; they are not invocations of an additional theorem in an unavailable file.

On a coarse torus let H_1 be the original metric and $B_1 = H_1^{-1}$ its inverse, with fine counterparts H_2, B_2. For the installed response T and penalty $b > 0$, define

$$V = B_1 - T^{\mathsf{T}} B_2 T - b^{-1} d_1 d_1^{\mathsf{T}},$$
$$M = H_1 V H_1.$$

The stationary identities of the transition are $TH_1 d_1 = H_2 e_1 A_0$, $T^{\mathsf{T}} e_1 = d_1 Q$ and $QA_0 = I - b^{-1} d_1^{\mathsf{T}} H_1 d_1$. Consequently,

$$T^{\mathsf{T}} B_2 T H_1 d_1 = d_1 [I - b^{-1} d_1^{\mathsf{T}} H_1 d_1],$$

so $VH_1 d_1 = 0$ and $Md_1 = 0$. Symmetry gives $d_1^{\mathsf{T}} M = 0$. The two-sided exterior lift therefore gives a three-form row F with $d_2^{\mathsf{T}} F d_2 = M$ on every torus. Replace F by $(F + F^{\mathsf{T}})/2$; this preserves the factorization and makes its real symbols Hermitian. Multiplying by B_1 on both sides gives

$$V = (d_2 B_1)^{\mathsf{T}} F (d_2 B_1).$$

This is the fund used in the positive correction. Analytic bounds are retained through the two exterior lifts, their symmetrization and the surrounding bounded rows. The exterior-division construction below supplies the analytic lift and its bounds.

The fine curvature charge $e_2 B_2 T H_1$ annihilates d_1, because $TH_1 d_1 = H_2 e_1 A_0$ and $e_2 e_1 = 0$. A right exterior lift produces S with $S d_2 = e_2 B_2 T H_1$. Multiplying by B_1 gives $S d_2 B_1 = e_2 B_2 T$, the rebasing intertwining identity. The refined estimate takes a fine difference before lifting and keeps L^{-3} per residue; the residue sum then gives the L^{-2} bound on $S^* S$ used in three-form absorption. The inverse-third-power argument below proves this quantitative bound.

The retained-support construction makes stagewise choices before the eventual block lengths. For a legal pair (j, t), where $j, t \geq 4$ and $t = j$ or $t + 1 = j$, the linear transition estimate supplies a positive coarse-floor threshold $\delta(j, t)$. Choose

$$l(i) = \max(\delta(i, i), \delta(i + 1, i)).$$

This pays the coarse threshold for both possible incoming transitions at stage i. With fine and coarse parameters $l(j)$,$l(t)$, choose the positive analytic width, positive coercivity parameter, cost, and integer block threshold supplied by each legal pair. At each stage also choose the positive clip cap and the bottom metric estimates. Take the minimum of the finitely many widths and coercivity parameters for the bottom and two outgoing possibilities, the maximum of their costs, and the maximum of their length thresholds. Weakening an analytic width or lower coercivity bound and increasing an upper cost preserves each estimate. Thus one table works for every allowed adjacent transition. This assembles the common price table from the metric, positive-margin and flight estimates established in this chapter.

For a fixed level schedule A, construct the linear trajectory backwards. Start with the top train at level N, and at each coarser level take the landing train of the flight determined by the block and adjacent smoothing indices. Denote the resulting train at i by $R_{N,i}$. At every pair (N,i), the bottom estimate chooses a metric bed defined by that train and $l(J(i))$. Choose as its transition tower the flight's specified duo with the fixed penalty. The identity $R_{N,i} = \mathrm{land}(\mathrm{destination}(N, i))$ aligns the bottom metric with this tower. The flight's tangent equations give the primal and dual stationary identities on each scheduled finite torus; its rebased identity supplies the associated inverse and error data. The collected bottom bounds, transition bounds, and terminal normalization form the required linear horizon.

The pinning step retains these metric beds and the indicated tower components while arranging the additional horizon conditions for the given sheet profile. Since retention preserves the beds and the relevant components, their defined train metrics and destination response equalities survive.

This supplies Present for the pinned horizon as well as the admissibility statement H.Suffices. The bottom metric, flight and pinning estimates used in this assembly are proved in the harmonic, positive-margin and auxiliary-tower arguments in this chapter.

5.15 Harmonic values and finite-torus stationary equations

At $p = 0$ every exterior signal vanishes; the regular factors bulb and bend equal one, and the principal patch multipliers are identities. Thus the filled potential response (pilot) and curvature response (wire) have principal value I. At a nonprincipal zero alias choose a coordinate with nonzero residue. Its line multiplier vanishes because $LR_L(q_a)\chi(q_a) = \chi(0) = 0$ and $\chi(q_a) \neq 0$. The box mass filter vanishes, since its smoothing exponent is positive. Every nonprincipal pilot and wire value is zero. The pilot's leading term contains $S = 0$ and its correction contains the zero coarse longitudinal row. The wire's first term contains the zero coarse adjoint derivative, and its second contains the vanishing flux. These explicit values establish the stationary symbol equations on the harmonic fibre; no harmonic projection is deleted.

For clarity, the finite Fourier transfer is as follows. Suppose a real matrix A has representation $AG = Fa$, with Fourier matrices satisfying $FF^* = F^*F = vI$ and $GG^* = G^*G = uI$. Taking adjoints and multiplying by G, F yields $A^{\mathsf{T}}F = G[(v/u)a^*]$. Two real matrices with the same symbol and Fourier matrices agree, since their difference annihilates the invertible matrix G. Products and differences of represented matrices become products and differences of symbols when the intervening Fourier matrices coincide.

The coarse Fourier Gram scalar is c^4 and the fine one is f^4, with $f = cL$. The transpose of the block mean therefore has alias column $L^{-3}V_1(q_l)$: the mean carries L and $c^4/f^4 = L^{-4}$. The installed A_0 and T columns are $L^{-1}\,\mathrm{pilot}_l$ and $L^{-2}\,\mathrm{wire}_l$. The fine exterior derivative is $L^{-1}d^L$ on the scaled alias, and its precision is $L^{-2}K_l$. Restoring these factors gives the exact finite real-matrix identities

$$\begin{aligned} Q_f A_0 &= \mathrm{mean}^{\mathsf{T}}\, Q_c, \\ \mathrm{mean}\, A_0 &= I - b^{-1} Q_c, \\ T^{\mathsf{T}} d_f &= d_c\, \mathrm{mean}, \\ T H_c d_c &= H_f d_f A_0. \end{aligned}$$

The installed cross rows have these symbol representations because their residue Fourier sum is the alias Fourier matrix times the stacked alias column. Evaluation at every coarse discrete frequency then gives the represented identities above. Invertibility of the finite Fourier matrices proves the real matrix identities at every period. Together with the exact amended-metric preservation and pinning arguments, this supplies the finite linear horizons used in the flat entrance estimate. The cumulative-filter section gives the scalar covariance feedback formulas and the train identities used in this construction.

5.16 Scaled exterior identities for the stationary columns

For scale S, set

$$q_a^S(p) = S(e^{ip_a} - 1), \qquad d_r^S = q^S \wedge, \qquad \delta_r^S(p) = d_r^S(-p)^{\mathsf{T}}.$$

On real momentum, δ_r^S is the Hermitian adjoint. The patch identities are

$$F_{r+1}(q_l) d_r^L(q_l) = d_r(p) F_r(q_l), \qquad R_L(q_{l,a}) q_a^L(q_l) = q_a(p).$$

The second proves the first coordinatewise. Transpose at $-p$ to obtain the dual transport identity. Also $d^S_{r+1}d^S_r = 0$ identically.

At regular momentum let $D = \operatorname{diag}(S^{-2}D_a)$, $I_D = D^{-1}$ and $\alpha = S^4\rho$. The two-form metric is $\alpha^{-1}\operatorname{mask}_2((I_D)_{aa})$, with the pairwise-product convention above. The weighted exterior contraction identity makes its potential precision

$$K = I_D - \alpha^{-1} I_D d^S_0 \delta^S_0 I_D.$$

Indeed expanding the two-form wedge product yields the diagonal weighted scalar $\sum_a q^S_a(-p)(I_D)_{aa}q^S_a(p) = \alpha$; the off-diagonal terms are exactly the rank-one product subtracted above. Thus, with $P = \alpha^{-1}\delta^S_0 I_D$ and $C = \alpha^{-1} I_D d^S_0$,

$$K = I_D - I_D d^S_0 P,$$
$$KD = I - C\delta^S_0, DK = I - d^S_0 P.$$

Also $Kd^S_0 = 0$ and $\delta^S_0 K = 0$ directly from the exterior chain identities. These are the rotor identities used in matching pilot and wire.

Write D_l for the fine diagonal covariance at q_l, H_l for its two-form pulse and K_l for its potential precision. Use D_c, H_c, K_c for the coarse quantities, F_r, V_r for the patch multiplier and its dual, and P_c for the coarse longitudinal row. Define

$$E_l = \frac{m(-q_l)\operatorname{tr}(D_l)}{Z}, \qquad Z = \sum_l |m(q_l)|^2 \operatorname{tr}(D_l).$$

On the punctured real cube all summands in Z are nonnegative and the principal one is positive. Hence $Z \neq 0$ and $\sum_l E_l m(q_l) = 1$. Covariance recursion gives

$$\sum_l F_1(q_l) D_l V_1(q_l) = D_c - b^{-1} I.$$

The regular potential response is

$$A_l = D_l V_1(q_l) K_c + E_l d^L_0(q_l) P_c.$$

Multiply by $F_1(q_l)$ and sum over l. The patch identity turns the second term into $\sum_l E_l m(q_l) d_0(p) P_c$. The preceding identities and $D_c K_c = I - d_0 P_c$ therefore give

$$\sum_l F_1(q_l) A_l = I - b^{-1} K_c.$$

Similarly $K_l A_l = V_1(q_l) K_c$. The longitudinal term is killed by $K_l d^L_0 = 0$, and the remaining correction in $K_l D_l$ is killed by $\delta^L_0 V_1 \delta_1 = 0$. This proves the coarse averaging and primal stationary equations on the punctured cube.

Set $C = H_c d_1 D_c \delta_1$. The weighted identities give $\delta_1(I - C) = 0$ and $(I - C)H_c d_1 = 0$. The regular curvature response is

$$T_l = H_l d^L_1 D_l V_1 \delta_1 + y_l V_2 (I - C),$$

where y_l is the original chain scalar and all fine factors are evaluated at q_l. Left multiplication by δ^L_1 uses $K_l D_l V_1 \delta_1 = V_1 \delta_1$ and $\delta^L_1 V_2 = V_1 \delta_1$, giving $\delta^L_1 T_l = V_1 \delta_1$. Right multiplication by $H_c d_1$ kills the complementary term. The potential response's longitudinal term vanishes under $d^L_1 d^L_0 = 0$. Hence

$$T_l H_c d_1 = H_l d_1^L A_l.$$

The filled pilot and wire equal these expressions away from the puncture. Their continuity extends the identities to $p = 0$ once the Fourier and scale normalization factors are restored. The finite Fourier matrices use the coarse volume and fine volume separately when taking adjoints; their ratio is L^{-4}. Accordingly the installed columns use toll L^{-1} for A_l and toll L^{-2} for T_l, while the adjoint averaging column carries toll L^{-3}. With these factors the equations are exactly the primal, dual and transport equations on every compatible pair of tori. The Fourier transfer in "Harmonic values and finite-torus stationary equations" restores these factors and proves the asserted real matrix identities at every period.

5.17 Extending the three-form positive margin

Suppose the analytic Hermitian three-form symbol Q obeys

$$\langle d_2(p)a, Q(p)d_2(p)a\rangle \geq v\|d_2(p)a\|^2$$

for every real p, with $v > 0$. First this controls every vector at zero. For a complex three-form x in four dimensions, the real-linear map $n \mapsto (in)\wedge x$ from $\mathbb{R}^4$ to the complex one-dimensional top-form space has a nonzero real kernel vector n. The exterior contraction identity gives $x = (in) \wedge a$ for some two-form a. Since $q(tn)/t \to in$, the vectors $d_2(tn)(a/t)$ tend to x. Apply the image bound and continuity of Q to obtain $Q(0) \geq vI$.

Let $C \geq 0$ be its analytic coefficient bound and N the fixed orientation-count constant. When all $|q_a(p)| \leq \delta$, the symbol departure estimate bounds every entry of $Q(p) - Q(0)$ by $\delta C/\eta$. Increase C to a positive value and choose $\delta = \eta v/(2N^2C)$. Then $Q(p) \geq (v/2)I$ in that region.

Elsewhere choose k with $|q_k(p)| > \delta$. Cartan decomposition gives $x = y + z$, with

$$y = d_2(p)[q_k(p)^{-1}\iota_k x], \qquad z = q_k(p)^{-1}\iota_k d_3(p)x.$$

Thus $\|z\|^2 \leq N\delta^{-2}\|d_3(p)x\|^2$. Put $M = N^2C$. The quadratic-form and cross-term bounds imply

$$\begin{aligned}\langle x, Qx\rangle &\geq v\|y\|^2 - 2M\|y\|\|z\| - M\|z\|^2\\ &\geq (v/2)\|y\|^2 - (M + 2M^2/v)\|z\|^2.\end{aligned}$$

Since $\|x\|^2 \leq 2\|y\|^2 + 2\|z\|^2$, it follows that

$$\begin{aligned}\langle x, Qx\rangle + [M + 2M^2/v + v/2]N\delta^{-2}\|d_3(p)x\|^2\\ \geq (v/4)\|x\|^2.\end{aligned}$$

Adding $r\, d_3^* d_3$, where $r = [M + 2M^2/v + v/2]N\delta^{-2}$, makes Q uniformly positive. Near zero it adds only a nonnegative form. This correction has finite-support Laurent coefficients and a common analytic bound. It preserves the physical curvature contraction because $d_3 d_2 = 0$. Enlarging the upper cost therefore gives a uniformly positive full-space row with exactly the original $d_2^{\mathsf{T}} Q d_2$ at every period. This is the positive extension used in pinning.

5.18 The auxiliary tower Gram identity

Weighted convolution gives common analytic bounds at width $\gamma = \min\{\sigma_f, \sigma_c, \nu\} > 0$, where ν is the three-form-root width. The full Gram row J stacks $\sqrt{\text{clip.direct}}\, I$ with the bottom root row, and is bounded by $\max\{\sqrt{\text{clip.direct}}, \text{rootCost}\}$. Each exterior derivative has finite support: a bound

valid at all periods is the finite sum of $e^{\gamma|z|_1}$ times absolute stencil coefficients over orientations and sites. Transposition preserves that bound. Composing rows with an intermediate orientation space of size n costs at most n times their entrywise weighted bounds, by the distance triangle inequality and absolute convergence.

Apply this first to JR, then to composition with the fine response T, then to multiplication by J'^{T}. The bound for T retains its factor L^{-2} throughout, so the mixed row costs $C_f L^{-2}$. The normal row $s^{-1}d'^{\mathsf{T}}J'^{\mathsf{T}}$ has an L-independent bound, as do $I - J'R'J'^{\mathsf{T}}$ and $PeR'J'^{\mathsf{T}}$. Stacking takes their maximum. Choose one upper price covering these and the original input prices. Keep A_0, T unchanged; every carried fine difference keeps its L^{-1-r} or L^{-2-r} bound at the reduced width and enlarged price. This proves the estimates named `F08`; positivity and the Gram identity are separate algebraic statements proved here.

On a fixed pair of tori put $W = J^{\mathsf{T}}J$ and $W' = J'^{\mathsf{T}}J'$, the fine and coarse metrics from the full Gram rows. Let $R = W^{-1}$ and $R' = W'^{-1}$. Retain the inverse response T; let d, d' be the fine and coarse first exterior derivatives, e the coarse second derivative, and $s = \sqrt{b}$. The factorized inverse defect is

$$W'ZW' = e^{\mathsf{T}}P^{\mathsf{T}}Pe,$$
$$Z = R' - T^{\mathsf{T}}RT - b^{-1}d'd'^{\mathsf{T}}.$$

Define $Q = J'R'J'^{\mathsf{T}}$ and $X = PeR'J'^{\mathsf{T}}$. Then $Q^{\mathsf{T}} = Q$, $Q^2 = Q$ and $QJ' = J'$. For the mixed, normal and image rows

$$M = JRTJ'^{\mathsf{T}}, N = s^{-1}d'^{\mathsf{T}}J'^{\mathsf{T}},$$
$$O = \operatorname{stack}(I - Q, X),$$

direct multiplication gives

$$M^{\mathsf{T}}M = J'T^{\mathsf{T}}RTJ'^{\mathsf{T}},$$
$$N^{\mathsf{T}}N = J'b^{-1}d'd'^{\mathsf{T}}J'^{\mathsf{T}},$$
$$O^{\mathsf{T}}O = I - Q + J'ZJ'^{\mathsf{T}}.$$

Substitution of Z, Q makes their sum I. This holds on the entire source of the stack, including harmonic coordinates. The transport identities

$$TW'd' = WdA_0, \qquad T^{\mathsf{T}}d = d' \operatorname{mean}$$

give

$$MJ'd' = JdA_0, \qquad M^{\mathsf{T}}Jd + N^{\mathsf{T}}(-s \operatorname{mean}) = 0.$$

Also $sN^{\mathsf{T}} = J'd'$, $NJ'd' = s^{-1}d'^{\mathsf{T}}W'd'$ and $OJ'd' = 0$. The last follows from $(I - Q)J' = 0$ and $ed' = 0$. These are the exact identities called `F09` in the source. They use the same A_0, T and defect, without an additional spectral estimate.

5.19 Pinning the auxiliary tower without changing the quadratic history

The remaining assembly changes only auxiliary tower rows. Suppose a transition has metric factors H_f, H_c with the stated analytic and positive lower bounds, response rows A_0, T with their carried fine-difference bounds, and inverse defect $d_2^{\mathsf{T}}Cd_2$, where C is analytic and positive on $\operatorname{im} d_2$. Extend C to a uniformly positive full three-form row without changing $d_2^{\mathsf{T}}Cd_2$, take its analytic square root

P, construct the auxiliary rows below, and apply their Gram identity. The preceding positive-margin, root and Gram arguments justify each operation.

Let F_f, F_c be the full two-form Gram factors and B_f, B_c the inverse metrics. Keep A_0, T unchanged. On the finite torus use mixed row $F_f B_f T F_c^{\mathsf{T}}$, normal row $b^{-1/2} d_1^{\mathsf{T}} F_c^{\mathsf{T}}$, and the image row obtained by stacking

$$I - F_c B_c F_c^{\mathsf{T}}$$
$$P d_2 B_c F_c^{\mathsf{T}}.$$

These formulas are supplied by common cover rows at all torus periods. The analytic auxiliary-row construction gives a positive width and common costs for them, including the required fine differences. Its inputs are fixed analytic estimates, clip profiles, and penalty; its constants are chosen before a particular block length. The exact Gram identity, using the bed floors and defect factorization, proves the required tower roll equations.

For each smoothing index first apply the positive extension and root operations, then choose the auxiliary-row prices for each of the two adjacent-stage possibilities. Intersect the finitely many widths with the original width and increase the cost to their maximum. Choosing at each transition the corresponding constructed tower gives a horizon with the identical beds and first two response rows and with the full horizon conditions at those new prices. Terminal normalization, crossing equations and bed bounds are inherited because these components were retained. The positive extension, auxiliary-row estimates and exact Gram identity therefore complete the pinning construction used in the retained horizon.

5.20 Completing the regular coefficient bounds and symmetries

All potential- and curvature-response bounds are made on one output momentum cube before choosing L. The mass filter is bounded by $C\,\mathrm{decay}_j(l)$ and its principal value has a positive real lower bound from the centered-average estimate. The line-filter factors are uniformly bounded, with positive principal lower bounds. At $l \neq 0$, $D_{l,a}$ and R_l^{-1} have respective bounds $C\,\mathrm{grade}(l)^{-2}$ and $C\,\mathrm{grade}(l)^{-4}$, by the common-collar reciprocal estimate and the bulb and bend bounds. The seed expression bounds S uniformly.

The sums H, T_a therefore have common analytic bounds. Their two smoothing factors contribute decay_{2j}, and each D contributes inverse squared grade. For $j \geq 4$ these are bounded by the uniformly summable decay_2. Also $U = |m(q_0)|^2 \sum_a f_a + SH$ is bounded above. Every real summand is nonnegative, and the principal filter and one f_a have positive lower bounds. Thus U has an L-independent positive lower bound. The reciprocal estimate bounds U^{-1} on a narrower collar. These observations bound all factors in the regular correction coefficients of the potential response; every matrix has a fixed face-orientation dimension.

For the curvature response, Λ is bounded because its only inverse factors are the principal mass filter and line filters. The other principal factors B, M, y are bounded products and sums of controlled quantities. At $l \neq 0$, the first branch has a signal of grade order one, D_l of order $\mathrm{grade}(l)^{-2}$, and a smoothing flux with $\mathrm{decay}_j(l)$. The second has two smoothing factors, the inverse-fourth-grade bound on R_l, a bounded coarse factor and a further smoothing flux. Both products are bounded by $C\,\mathrm{decay}_j(l)$, with constants fixed before L.

For real overlap identities, first use the regular stationary expressions. Every covariance D_a and its scalar sum ρ is periodic, and exterior signals and patch multipliers have the corresponding reciprocal periodicity. Shifting coarse momentum by $2\pi x$ therefore permutes l to $l + \mathrm{wrap}_L(x)$ in both responses. If $x \neq 0$ and both momenta lie in the enlarged real cube, neither is zero, so the

regular matching applies to both. For $x = 0$ the identity is immediate. At nonzero momentum, scalar real symbols are even, Fourier rows conjugate at $-p$, and exterior matrices transform by entrywise conjugation. Substitution gives

$$\overline{\text{pilot}(p, l)} = \text{pilot}(-p, -l), \qquad \overline{\text{wire}(p, l)} = \text{wire}(-p, -l).$$

At zero, the principal formulas have their normalized values and all other branches vanish. Reflection therefore holds there too. These shift and reflection identities permit the common-collar Fourier export of the residue sums.

5.21 The inverse third power before residue summation

The response kernel called buckle is obtained from the symbol $\text{vault}(\text{wing}, q_l)\,\text{wire}(p, l)$ with scale factor L^{-2}. Vault and its overlap identities preserve the common analytic estimates, so one actual fine difference of buckle has coefficient bound CL^{-3}. Its fine curl e_2 sums the carried forward differences over four directions and the fixed two-form orientations, with incidence signs of modulus at most one. Hence its bound is $4|\text{Faces}_2|CL^{-3}$. Only orientation counts enter; no residue count is paid at this step.

Multiplying the curl by the bounded coarse metric gives the charge kernel with the same L^{-3} scale. Apply the analytic right exterior lift on the fixed coarse orientation columns; it retains the prefactor, so the lifted kernel called sale is bounded by $C'L^{-3}$ per fine residue. Its Gram row sums over L^4 residues and the fixed fine orientations. The bound is therefore

$$C''L^4(L^{-3})^2 = C''L^{-2}.$$

This is the inverse-square scale used in three-form absorption. Estimating before taking the fine difference would lose the power; the actual proof takes the fine curl before lifting.

5.22 Filling the stationary predictor at the principal alias

Predictor convention and amendment A1. The regular principal coefficient uses the parenthesized expression below. The extraction placed $m(q_0)^{-1}$ on its first term only. To see the required scope of that multiplier, retain the directional rail in the covariance identity. Put $m_+ = m(q_0)$, $m_- = m(-q_0)$ and, for this calculation only, $a = |\,\text{pullLine}_a(q_0)|^2 f_a/(m_+m_-)$, $F = \sum_c f_c$, $g = m_+m_-a + ST$ and $U = m_+m_-F + SH$. Thus a is a reduced principal coefficient, rather than an unannounced replacement of the source's directional multiplier by m_+m_-. The corresponding weights are $E_0 = m_-F/U$, $r_0 = m_-a/g$, and

$$E_0 - r_0 = \frac{m_-S(FT - aH)}{Ug} = \frac{S}{m_+}\left(\frac{T}{g} - \frac{H}{U}\right).$$

In the real even reduced convention $m_+ = m_- = m$ this is the shorter m^2 calculation; on the complex collar the two factors retain the source's opposite-momentum convention. This is the cancellation needed before extracting Fourier coefficients. The amendment changes the parentheses, and retains all nonprincipal branches, the collar hypotheses and the harmonic restoration.

For one transition, use

$$q_l = \frac{p + 2\pi\,\text{lift}(l)}{L}, \qquad q_0 = p/L, \qquad S = \lambda_{\text{mass(land)}}(p).$$

Put $g_a = \text{bulb}(\text{land}, p, a)$, $h = \text{bend}(\text{land}, p)$, $m(q) = \text{massFilter}(q)$ and $f_a = \text{bulb}(\text{wing}, q_0, a)$. The land and wing are the coarse and fine smoothing trains of this transition. For $l \neq 0$, define

$$D_{l,a} = \frac{\operatorname{bulb}(\operatorname{wing}, q_l, a)}{\operatorname{blow}(\operatorname{mass}(\operatorname{wing}), p, l)}.$$

Set $T_a = \sum_{l\neq 0} |\operatorname{pullLine}_a(q_l)|^2 D_{l,a} + b^{-1}$. On real momentum the notation is an absolute square; on complex momentum it means the product at q_l and $-q_l$. The alias covariance recursion gives

$$g_a = |\operatorname{pullLine}_a(q_0)|^2 f_a + S T_a. \tag{3}$$

Put $H = \sum_{l\neq 0} |m(q_l)|^2 \sum_a D_{l,a}$ and

$$U = |m(q_0)|^2 \sum_a f_a + SH.$$

Away from $p = 0$ this equals S times the original sum of weighted covariance entries. The first term has a positive real lower bound; the other term is nonnegative, so U has a positive lower bound. The earlier alias estimates bound H, T_a and U on a common collar, and the reciprocal lemma bounds U^{-1} there. Thus U is the regular denominator used in place of the original singular covariance sum.

Define the regular correction coefficients

$$\begin{aligned} G_{0,a} &= m(q_0)^{-1}(T_a/g_a - H/U), \\ G_{l,a} &= m(-q_l)[(\sum_c D_{l,c})/U \\ &\quad - D_{l,a}|R_L(q_{l,a})|^2/g_a], l \neq 0. \end{aligned}$$

Every denominator is regular: g_a, $m(q_0)$, U and the nonzero-alias divisor. The estimates above bound G_0 uniformly and G_l by $C \operatorname{decay}_j(l) \operatorname{grade}(l)^{-2}$ for $l \neq 0$. The principal formula follows from (3). For the original normalized covariance weights E_l and directional rebates $r_{l,a}$, the identity is $E_l - r_{l,a} = SG_{l,a}$. The original longitudinal factor is S^{-1} times the regular row

$$h^{-1} d_0(p)^* \operatorname{diag}(g_a^{-1}).$$

Thus S cancels before Fourier coefficients are extracted. The pilot is its regular transverse leading term plus

$$\begin{aligned} &\operatorname{diag}(G_{l,a}) d_0^L(q_l) \\ &h^{-1} d_0(p)^* \operatorname{diag}(g_a^{-1}). \end{aligned}$$

The transverse leading terms are

$$\begin{aligned} l = 0: &\quad \operatorname{diag}(f_a) \operatorname{flux}_1(q_0) \operatorname{diag}(g_a^{-1}), \\ l \neq 0: &\quad S \operatorname{diag}(D_{l,a}) \operatorname{flux}_1(q_l) \operatorname{diag}(g_a^{-1}). \end{aligned}$$

Here flux_1 is the transpose of the directional one-form patch multiplier, and d_0^L has the specified scaled normalization. The scalar cancellation and these regular formulas supply the value at zero. This is the filled potential response; the original expression with an isolated longitudinal denominator is not separately asserted to extend holomorphically.

The potential-response bound follows term by term. Its leading term has $\operatorname{decay}_j(l)$. Its correction has $\operatorname{decay}_j(l) \operatorname{grade}(l)^{-2}$ times a signal bounded by a constant times $\operatorname{grade}(l)$; every other factor is bounded. Since the grade is at least one, the same decay bounds the correction. Section 5.24 proves the covariance recursion; Section 5.16 gives the stationary equations; Section 5.15

restores the finite-torus normalizations; and Section 5.20 gives reciprocal-shift and reflection compatibility. Together these prove the analytic predictor and its actual stationary response, including the harmonic fibre.

5.23 The regular dual predictor on two-forms

Continue the filled-potential-response notation. Let $H_f(q)$ and $H_c(p)$ be the fine and coarse pulse metrics, $d = d_1(p)$ and $\delta = d_1^L(q_0)$. The adjoint d^* on real momentum is continued analytically using $-p$. Write $F_r(q) = \text{hinges}_r(q)$ for the patch multiplier on r-forms, and $V_r(q) = \text{flux}_r(q)$ for its transposed symbol at $-q$. Define

$$\begin{aligned}\Lambda &= m(q_0)^{-1}\,\text{mask}_2(R_L(q_{0,a})^{-1}),\\ y &= |m(q_0)|^2 h_c/h_f(q_0),\\ x_a &= T_a/g_a,\\ B &= H_f(q_0)\big[\text{diag}_{\text{pairs}}(x_a + x_c) - S\,\text{mask}_2(x)\big],\\ M &= \text{diag}(g_a).\end{aligned}$$

Here $\text{diag}_{\text{pairs}}$ sums the two coordinate entries in each two-form orientation. The principal wire is the regular expression

$$\begin{aligned}W_0 = yV_2(q_0)\\ + \big[B\Lambda dM - H_f(q_0)\Lambda d\,\text{diag}(T_a)\big]d^*.\end{aligned}$$

For a vector $x = (x_a)$, let $\text{mask}_2(x)$ have diagonal entries $x_a x_c$ at orientation $\{a, c\}$, and let $\text{diag}_{\text{pairs}}(x_a + x_c)$ have entries $x_a + x_c$. The entrywise identity

$$I - \text{mask}_2(1 - Sx) = S\big[\text{diag}_{\text{pairs}}(x_a + x_c) - S\,\text{mask}_2(x)\big]$$

and covariance identity (3) prove

$$SB = H_f(q_0) - yV_2(q_0)H_c(p)F_2(q_0),$$

by substitution in each two-form diagonal entry. Together with $F_2(q_0)\Lambda = I$ and $F_2(q_0)\delta = dF_1(q_0)$, this identifies W_0 with the original stationary dual expression for $p \neq 0$. Every factor of the displayed regular expression is uniformly analytic, so it extends the response through zero.

For $l \neq 0$, set $R_l = \text{blow}_l^2\, h_f(q_l)$ and $y_l = |m(q_l)|^2 h_c/R_l$, and put

$$N = S^2 I - SH_c(p)dMd^*.$$

The remaining wire branches are

$$\begin{aligned}W_l = H_f(q_l)d_1^L(q_l)\,\text{diag}(D_{l,a})V_1(q_l)d^*\\ + y_l V_2(q_l)N.\end{aligned}$$

These are the nonprincipal dual branches, called stake in the source. The identities portD = scaleD, portR = scaleR and $N = S^2(I - \text{cloud})$ on the regular set identify them with the original dual response. Their nonzero-alias divisor is separated, so no principal denominator is inverted. The lower bound gives $R_l^{-1} = O(\text{grade}(l)^{-4})$ and $D_l = O(\text{grade}(l)^{-2})$ with the common analytic bounds. Signal and smoothing-flux factors then give $C\,\text{decay}_j(l)$, as for the bounded principal branch. Section 5.20 gives the reciprocal-shift and reflection identities needed for Fourier export. The algebra here identifies the regular symbols themselves.

5.24 Cumulative filters and propagation of the quadratic covariance

The cumulative-filter calculation closes the quadratic argument. Each one-step identity is composed along the same scale history, and the weighted alias sums remain absolutely convergent. The resulting covariance is therefore attached to the constructed history, rather than to an independently chosen continuum kernel. These identities are later used as the fixed linear comparison in the nonlinear predictor and in the free hierarchy.

A finite linear history begins at smoothing order j. It is either empty, or consists of a patch with odd block side $L \geq 3$, new order J with $j \leq J \leq j+1$, followed by a history beginning at J. Let M_v be the product of its block sides, with $M_\varnothing = 1$. For the centered one-dimensional box filter p_L, write

$$A_L(p) = \prod_{\nu=0}^{3} p_L(p_\nu).$$

Its exact composition identity for odd sides is

$$A_{LM}(p) = A_L(Mp)A_M(p).$$

Define an extra filter recursively by $E_\varnothing = 1$ and

$$E_{(L,J;v)}(p) = A_{LM_v}(p)^{J-j}E_v(p), \qquad S_v(p) = A_{M_v}(p)^j E_v(p).$$

The composition identity gives

$$S_{(L,J;v)}(p) = A_L(M_v p)^J S_v(p). \tag{QF1}$$

The exponent $J-j$ is zero or one. Thus extra smoothing factors enter at the successive cumulative lengths, while the principal filter has exactly the original exponent j.

The filters are entire, periodic under reciprocal integer translations, invariant under signed coordinate permutations, equal to one at zero, and satisfy $\overline{S_v(p)} = S_v(-\overline{p})$. These statements follow first for the finite box sums, then for A_L, E_v, S_v by finite products and induction.

There is a bound independent of both the number of steps and the total dilation. On the strip $|\operatorname{Im} p_\nu| \leq z/L_0$, a box filter satisfies

$$|A_M(p)| \leq e^{4zM/L_0}.$$

When $M \leq L_0$, at real principal momenta p/L_0 in the fixed domain one also has

$$|A_M(p/L_0)| \geq e^{-4c_{\mathrm{box}}M/L_0}.$$

For a history of total mass $M_v \leq L_0$, induction using each block side at least three gives

$$|E_v(p)| \leq e^{8zM_v/L_0}, \qquad |E_v(p/L_0)| \geq e^{-8c_{\mathrm{box}}M_v/L_0}. \tag{QF2}$$

To check the induction, a new extra factor costs $4zM_v/L_0$ in the upper exponent, and the preceding bound costs at most $8zM_v/(3L_0)$. Their sum is at most $8zM_v/L_0$. If no extra factor occurs, monotonicity of the mass suffices. The lower bound has exactly the same calculation with negative exponents and c_{box} in place of z.

At an alias indexed by l, the principal box estimate is

$$|A_M(\mathrm{alias}_l(p))| \leq B^4 \operatorname{mix}(l)^{-1}.$$

Combining this with (QF2), with $L_0 = M_v$, proves

$$|S_v(\mathrm{alias}_l(p))| \le B^{4j} e^{8z} \,\mathrm{mix}(l)^{-j}. \tag{QF3}$$

At the real principal alias the lower bound is

$$|S_v(p/M_v)| \ge e^{-4jc_{\mathrm{box}}-8c_{\mathrm{box}}} > 0. \tag{QF4}$$

These bounds are uniform over all finite histories, with constants depending on the initial order only. In particular all alias-filter functions have a common holomorphic neighborhood and the indicated uniform alias decay.

Include now the longitudinal edge box filter $r_M(p_a)$, and set

$$D_{v,a}(p) = S_v(p) r_{M_v}(p_a), \qquad P_{L,J,a}(p) = A_L(p)^J r_L(p_a), \qquad Y(f)(p) = f(p)f(-p).$$

The exact rail composition and (QF1) imply

$$D_{(L,J;v),a}(p) = P_{L,J,a}(M_v p) D_{v,a}(p). \tag{QF5}$$

On real momenta $D_{v,a}(-p) = \overline{D_{v,a}(p)}$, so $Y(D_{v,a}) = |D_{v,a}|^2 \ge 0$.

For any scalar symbol f, define its propagated alias sum by

$$\mathcal{F}_{v,a} f(p) = M_v^{-2} \sum_{l \in (\mathbb{Z}/M_v\mathbb{Z})^4} Y(D_{v,a})(\mathrm{alias}_l(p)) f(\mathrm{alias}_l(p)). \tag{QF6}$$

The normalization is M_v^{-2}. Define $\mathcal{L}_{L,J,a}$ by the same formula for one patch. The radix bijection between aliases of LM_v and pairs of aliases of L and M_v, together with (QF5), gives the exact identity

$$\mathcal{F}_{(L,J;v),a} f = \mathcal{L}_{L,J,a}(\mathcal{F}_{v,a} f). \tag{QF7}$$

Indeed substitute the nested alias formula in the sum, factor out the outer $Y(P_{L,J,a})$ by periodicity, and use $(LM_v)^{-2} = L^{-2} M_v^{-2}$. The empty history has $\mathcal{F}_{\varnothing,a} f = f$.

Every renormalization step adds an independent local covariance k^{-1}. Let $\mathcal{N}_v$ be the sum of the covariances of these noises propagated through their respective prefixes. More explicitly the empty history has no prefixes; extending a history introduces the empty prefix and extends each old prefix by the new patch. Thus

$$\mathcal{N}_\varnothing = 0, \qquad \mathcal{N}_{(L,J;v),a} = 1 + \mathcal{L}_{L,J,a} \mathcal{N}_{v,a}.$$

If $\lambda(p)$ is the initial Wilson quadratic divisor, the cumulative diagonal covariance is

$$\mathcal{D}_{v,a} = \mathcal{F}_{v,a}(\lambda^{-1}) + k^{-1} \mathcal{N}_{v,a}.$$

Linearity and (QF7) prove the feedback equation

$$\mathcal{D}_{(L,J;v),a} = \mathcal{L}_{L,J,a} \mathcal{D}_{v,a} + k^{-1}. \tag{QF8}$$

This formula accounts for every earlier noise; retaining only the final noise would give a different covariance. All sums are finite. Real nonnegativity follows term by term when $k > 0$, and reciprocal periodicity, evenness and covariance under signed coordinate permutations follow by the corresponding permutation of aliases.

Separating the principal alias in (QF6) yields the pole decomposition

$$\mathcal{D}_{v,a}(p) = \frac{Y(D_{v,a})(p/M_v)}{\lambda_{M_v}(p)} + \mathcal{H}_{v,a}(p) + k^{-1}\mathcal{N}_{v,a}(p), \tag{QF9}$$

where λ_{M_v} includes the M_v^2 scaling of the principal divisor and

$$\mathcal{H}_{v,a}(p) = \sum_{l\neq 0} \frac{Y(D_{v,a})(\mathrm{alias}_l(p))}{M_v^2\lambda(\mathrm{alias}_l(p))}.$$

This is an algebraic identity on the regular set of the denominators. The auxiliary divided symbols are not asserted to be continuous extensions at harmonic momenta. Their regular parts and their harmonic reconstruction are handled by the pole-removal construction already given above. Both $\mathcal{H}$ and the propagated noise are nonnegative on real regular momenta, and (QF4) supplies the uniform nonvanishing principal numerator required there.

The one-dimensional filters in this construction are literal normalized finite sums:

$$r_M(z) = \frac{1}{M}\sum_{u=0}^{M-1} e^{\mathrm{i}uz}, \qquad p_M(z) = \frac{1}{M}\sum_{u=0}^{M-1} e^{\mathrm{i}(u-\lfloor M/2\rfloor)z}.$$

Partitioning the indices $0,\dots,LM-1$ as $Mu+v$ proves $r_L(Mz)r_M(z) = r_{LM}(z)$. For odd sides the centers satisfy $M\lfloor L/2\rfloor + \lfloor M/2\rfloor = \lfloor LM/2\rfloor$, proving the corresponding identity for p.

The principal lower bound used in (QF2) has coefficient one in front of its exponential, which is essential for arbitrarily long histories. To derive it, choose a disk on which

$$c_-|z| \le |e^{\mathrm{i}z} - 1| \le c_+|z|.$$

Such positive constants exist because the ratio extends continuously with nonzero value at zero, and the disk excludes the other zeros. The geometric-sum identity

$$Mr_M(z)(e^{\mathrm{i}z} - 1) = e^{\mathrm{i}Mz} - 1$$

gives $|p_M(z)| = |r_M(z)| \ge c_0$ for real z with $M|z|$ in that disk, where $c_0 = \min(1, c_-/c_+) > 0$. If $M \le L_0$ and $L_0|z| \le R_0$, the centered finite sum also gives

$$|p_M(z) - 1| \le c_+R_0\, M/L_0.$$

Put $x = M/L_0$, $k_0 = c_+R_0$, and $u = |p_M(z)|$. Then $u \ge c_0$ and $u + k_0x \ge 1$. Since $e^y \ge 1 + y$,

$$ue^{k_0x/c_0} \ge u + uk_0x/c_0 \ge u + k_0x \ge 1.$$

Thus $u \ge e^{-(k_0/c_0)x}$. This proves the stated principal estimate with $c_{\mathrm{box}} = k_0/c_0$, without multiplying a fixed prefactor less than one at every step. The complex-strip upper bound follows simply by bounding every exponential in the centered average by $e^{M|\operatorname{Im} z|}$.

6 Local nonlinear geometry and estimates

The nonlinear argument replaces a flat Lie-algebra field by actual group-valued links while controlling the difference from the quadratic model. It first constructs the chart and its Haar density, then obtains local inverse and predictor estimates, and finally pays for the exceptional configurations through packet and support budgets. Each estimate records whether it is a pointwise bound, a derivative bound on a healthy domain, or a summable bound over supports. These three forms of control have different uses in the primitive integration and cannot be substituted for one another.

6.1 Constructing the exponential chart and its Haar ratio

The chart is constructed from the closed matrix image of the compact gauge group. Its real Lie algebra consists of matrices X such that $\exp(tX)$ belongs to that image for every real t. Let N be the orthogonal complement of this Lie algebra in the ambient real matrix space. First suppose $X_j \neq 0$, $X_j \to 0$, $\exp(X_j) \in G$, and $X_j/\|X_j\| \to u$. For $t \geq 0$, put $k_j = \lfloor t/\|X_j\| \rfloor$. Then $k_j\|X_j\| \to t$ and

$$\exp(k_j X_j) = (\exp X_j)^{k_j} \in G, \qquad \exp(k_j X_j) \longrightarrow \exp(tu).$$

Closedness gives $\exp(tu) \in G$; negative t follow by inversion. Thus u belongs to the Lie algebra.

Consequently, in some neighborhood of zero in N, the condition $\exp(n) \in G$ forces $n = 0$. Otherwise choose nonzero such $n_j \to 0$ and extract a convergent subsequence of their unit directions. Its limit has norm one, belongs to N, and by the preceding argument belongs to the Lie algebra, a contradiction.

Identify the color space linearly with the Lie algebra by C, and write $\exp_G(a) = \exp(C(a))$. The map

$$(a, n) \longmapsto \exp(C(a)) \exp(n)$$

has derivative $(a, n) \mapsto C(a) + n$ at zero, an isomorphism by the orthogonal decomposition. Its local smooth inverse has a color component $\log_G$. In particular $\log_G(\exp(C(a))) = a$ for small a. If a nearby $g \in G$ has inverse coordinates (a, n), then $\exp(n) = \exp(-C(a))g \in G$. The isolation just proved gives $n = 0$, so $\exp(C(\log_G g)) = g$. Shrinking the color ball, its exponential image is open in G: continuity of the inverse sends a group neighborhood of each image point back into the ball. No global logarithm or connectedness assumption is used.

For $u \in G$, define $F_u(a) = \log_G(u \exp(C(a)))$ where the translated point lies in the chart. Put

$$A(a) = DF_{\exp(-C(a))}(a), \qquad J(a) = \det A(a).$$

Smoothness of the ambient inverse makes J smooth, and F_1 is the identity germ, so $J(0) = 1$. If $b = F_u(a)$, the identity of germs near a

$$F_{\exp(-C(a))} = F_{\exp(-C(b))} \circ F_u$$

gives, by differentiation and determinants,

$$J(b) \det DF_u(a) = J(a).$$

Shrink the ball S so $1/2 < J < 2$ there and put $U = \exp_G(S)$. Push $\mathbf{1}_S J(a)\, da$ forward by $\exp_G$, obtaining a finite nonzero measure v supported on U. For measurable $A \subseteq U$, the change-of-variables formula and the overlap law give

$$v(u^{-1}A) = v(A \cap uU).$$

Indeed F_u is injective on the relevant chart pieces, and positivity of J makes its density identity $|\det DF_u(a)| J(F_u(a)) = J(a)$.

Average over normalized Haar measure. Tonelli and right invariance make the left integral $v(G)\mu_G(A)$. For fixed $g \in A$, the set of u with $g \in uU$ has measure $\mu_G(U)$, by inversion and right translation, so the right integral is $\mu_G(U)v(A)$. Both constants are positive and finite: U is nonempty and open, and J is bounded above and below on the finite-radius ball. Hence $c = \mu_G(U)/v(G) > 0$ satisfies

$$(\exp_G)_*(\mathbf{1}_S c J(a)\, da) = \mu_G|_U.$$

Intersecting an arbitrary measurable test set with U proves the full local measure identity. This gives the required normalization, positivity, smoothness and value $J(0) = 1$.

6.2 Local Haar charts in the primitive integral

Active and dormant coordinates; amendment A2. The chart integral contains rh active real coordinates and d_v dormant Gaussian coordinates. Integration of the latter gives one, including when $d_v = 0$; it does not authorize dropping their normalization. Put $K = \sqrt{2\pi}$ and $m_g = c(Kg)^r$. The $f - h$ frozen free links cancel as $m_g^f e^{-(f-h)\log m_g} = m_g^h$. Consequently

$$K^{-(rh+d_v)} m_g^h = K^{-(rh+d_v)} c^h K^{rh} g^{rh} = K^{-d_v} c^h g^{rh}.$$

The extracted exponent $-rh + d_v$ would instead produce $K^{d_v} c^h g^{rh}$, an extra factor K^{2d_v}. We use the reconciled exponent $-(rh + d_v)$ below. The determinant of the reference form remains a separate factor.

The local change of variables in the primitive expansion uses the chart ratio of the same compact gauge group. If B_ρ is its chart ball and the local Haar ratio is $cJ(a)$, then

$$a \longmapsto \exp_G(a)V$$

sends $\mathbf{1}_{B_\rho}(a) cJ(a)\, da$ to Haar measure restricted to $\exp_G(B_\rho)V$, for every fixed $V \in G$. This follows by the defining chart measure identity for the exponential followed by right Haar invariance. On the finite set of active links, take the product of these maps. Product measure and multiplication of densities give the exact joint density $\prod_i cJ(a_i)$ on the product chart ball, with the product of the restricted Haar measures as its image. The inactive links are held fixed.

The active chart is sufficient because the sector selectors force every active link of a nonzero integrand to lie in that chart relative to its replica background. More precisely those links have the form $\exp_G(a_i)V_i$ with $|a_i| \le a_0 t$ inside the chart ball. Consequently restricting integration to the product chart does not discard a nonzero contribution. The transformed integrand's powers of c and its logarithmic Jacobian terms come from this product density.

Before this change, the tree-gauge configuration measure decomposes into active-link and exterior-link coordinates by a measure-preserving splitting. Fubini permits integration first over the active coordinates and then the exterior coordinates, even after the team's bounded local ingredients are inserted. Regrouping the integrand after this substitution produces its quadratic Gaussian density and the remaining explicit vertex factors. The chart identity and product-measure operations are exact. The local Haar-chart construction, grafted-signal argument and Gaussian exponent rearrangement in this chapter supply the chart ratio, selector implication and vertex decomposition used here.

The selector-to-chart implication has the following numerical form. On an active link let $z = U_i V_i^{-1}$, with V_i its replica background. The actual sector and tree-slice conditions give $\mathrm{chord}(z) \le Kt$, where K is the selected shell width and t the small-field scale. The local logarithm estimate gives $\|\log_G z\| \le \gamma\, \mathrm{chord}(z)$ and $\exp_G(\log_G z) = z$ on the wider shell $\mathrm{chord}(z) \le 2Kt$. Choose a with $\gamma K \le a$ and $at < \rho$, the chart radius. Then $v = \log_G z$ satisfies $\|v\| \le at < \rho$ and $U_i = \exp_G(v)V_i$. This is the chart representation of every active link of a nonzero sector integrand. Its substantive geometric input is the shell-access estimate deriving $\mathrm{chord}(z) \le Kt$ from the actual sector conditions.

The inserted cutoff functions can be handled at all color arguments. Write $\chi_a(v)$ for the inner-share cutoff at radius at. It lies between zero and one, equals one on $\|v\| \le at$, and can be nonzero only on $\|v\| \le 2at$. Therefore $\chi_{2a} = 1$ wherever χ_a is nonzero. Define the extended Jacobian logarithm by $g_a(v) = \chi_{2a}(v) \log J(v)$. On the support of χ_a, positivity of J in the chart gives $\exp(g_a(v)) = J(v)$, provided $2at < \rho$. This identity yields the exact product Jacobian in

the integral. Outside the chart the cutoff makes the transformed integrand vanish, so no exterior logarithm value is used in the change of variables.

Since $J(0) = 1$ and J is smooth and positive near zero, log J is locally Lipschitz with $|\log J(v)| \leq C\|v\|$. Wherever χ_{2a} is nonzero, $\|v\| \leq 4at$. Hence $|g_a(v)| \leq 4Cat$ for all v, after making 4at smaller than the fixed local radius. The zero-cutoff case is immediate. This proves the uniform small Jacobian-log bound used in the primitive vertex envelopes.

The later rescaling from active chart variables to fluctuation variables is likewise exact. Let $g > 0$ be the coupling scale, p the real number of active coordinates, and separate fluctuation coordinates into active and vacant parts. The active affine dilation used in the source sends g^p times Lebesgue measure to active-chart Lebesgue measure. Its coordinate reassembly is a permutation and preserves measure. The integrand depends on the active part through the original chart expression; the vacant part contributes a separate normalized Gaussian comfort factor with integral one. Thus, in these coordinates, its integral factors as

$$\int Y(x)g^p C(y)dxdy = g^p \int Y(x)dx, \int C(y)dy = 1.$$

The dormant quadratic energy identity gives precisely this factorization of the transformed ribbons, cutoff and Jacobian terms. Fubini is licensed by the primitive integrable envelopes. This accounts for the active-coordinate Jacobian and the unused Gaussian coordinates in the exact fluctuation integral, rather than discarding them by a dimension convention. The affine dilation, dormant energy and normalization are derived next; the packet and non-packet estimates below supply the vertex envelopes.

The coordinate map in that rescaling is simply affine dilation followed by reindexing. Index active real coordinates by pairs (active link, color coordinate), let p be their cardinality, and put $D(x) = \mu + gx$. Its derivative is g times the identity, so $\det(D_x D) = g^p > 0$. The finite-dimensional change-of-variables formula says D pushes g^p dx to Lebesgue measure. Reindex the pairs into one color vector per active link; this is a coordinate permutation followed by regrouping product coordinates and preserves Lebesgue measure on measurable rectangles, hence on all measurable sets. The active coordinate of the original lifted fluctuation $\mu + gz$ is exactly this $D(x)$. The configuration, active cutoff product and Jacobian-log sum depend only on these active coordinates.

On the vacant coordinates the source's dormant matrix is the identity and on active coordinates it is zero. Its quadratic energy is therefore one half the sum of the squared vacant real coordinates. If d_v is their number, its comfort density is

$$C(y) = (\sqrt{2\pi})^{-d_v} \exp(-\|y\|^2/2).$$

The product of the one-dimensional Gaussian integrals gives $\int C = 1$, also for $d_v = 0$. This proves the dormant factor and its normalization used in the previous integral identity.

The remaining volume constants can be checked by actual finite set partitions. If r is the color dimension, h is the number of active free links, and f is the total number of free tree-gauge links, then $p = rh$, the number of dormant coordinates is d_v, and the ambient fluctuation dimension is $p + d_v$. The number of frozen free links is f-h. Define $m_g = c(\sqrt{2\pi}\, g)^r$, where c is the Haar chart scale. The chart/rescaling prefactor satisfies the exact identity

$$\begin{gathered}(\sqrt{2\pi})^{-d_v} g^{rh} c^h \\ = (\sqrt{2\pi})^{-(rh+d_v)} \\ m_g^f \exp(-(f-h)\log m_g).\end{gathered}$$

It follows by cancelling m_g^{f-h} and substituting its definition. The determinant of the reference blank quadratic form contributes the separate factor $\exp(-\frac{1}{2}\log\det Q_{\text{blank}})$; it is not included in a dimension count. Together these are the deducted Gaussian normalization used by the local integral. Its reference covariance has density

$$(\sqrt{2\pi})^{-d}\exp(-\log\det Q/2 - z^{\mathsf{T}}Q^{-1}z/2)$$

in ambient dimension d, with positive definite Q. Thus all determinant and Gaussian-volume factors are explicit before the remainder vertices are estimated.

There is also a useful exact cancellation in the raw density before this transformation. Splitting the action into a retained crown and its reduced part gives

$$\begin{aligned}&\beta\,\mathrm{cut}_R(U) + E_R(U)\\ &= v_R|\Lambda| - \mathrm{reducedAction}(U) - \mathrm{additiveActivity}(U).\end{aligned}$$

The block-kernel correction has whole gap $\mathrm{cut}_R(U)$-coarseCrown(V). Its exponential cancels the retained crown already moved into the raw density, since

$$\begin{aligned}&\exp(-\beta\,\mathrm{coarseCrown})\exp(\beta\,\mathrm{cut}_R)\\ &\exp(-\beta(\mathrm{cut}_R - \mathrm{coarseCrown})) = 1.\end{aligned}$$

Consequently the sector sum of the raw marked density equals the tagged block kernel times the original marked density pointwise. A summable envelope for the original activity gas and bounded sector factors licenses integration and summation. This is the initial action cancellation in the primitive Mark formula. The local vertex decomposition and primitive-integral estimates in the marked-integration chapter complete this formula.

6.3 The quadratic floor before choosing the local covariance

The reference precision is a Gram matrix with an identity completion on dormant coordinates. Let L u be the active part of a link direction, let D u be the physical stacked motion, and let P_d be the diagonal projection onto dormant real coordinates. The plate is $Q = F^{\mathsf{T}}F + P_d$, where F is the scaled motion matrix. The scale σ satisfies

$$u^{\mathsf{T}}F^{\mathsf{T}}Fu = \sigma^2\sum|Du|^2.$$

To prove a volume-uniform positive floor, the source first obtains at each link the local inequality

$$|Lu_i| \le ZR_i(u) + MC_i(u).$$

Here R_i is the square root of the sum of squared physical motions within a fixed native radius d of i, while C_i is the square root of the sum of squared active link directions within radius 3B. Their definitions use actual torus max-distance and the actual link census. Summing squared inequalities and charging how often any term occurs gives

$$W \le 2Z^2N_dV + 2M^2N_{3B}W,$$

where $W = \sum|L\ u_i|^2$, $V = \sum|D\ u|^2$ and $N_r = (2r+1)^4\ L_block^4\ |\mathrm{Axis}|$. The local geometric driving estimate supplies Z fixed at the chosen block and stage, while M has reference-coupling order $-1/2$ up to the displayed width factors. Since N_{3B} has only polynomial width growth, eventually $2M^2N_{3B} \le 1/2$. Absorption gives $W \le 4Z^2N_d$ V. Enlarge the constant to $a > 0$ if needed.

Every real direction splits orthogonally into its active and dormant components:

$$\|u\|^2 = \|Lu\|^2 + u^{\mathsf{T}} P_d u, u^{\mathsf{T}} P_d u \geq 0.$$

Taking $k = \min(\sigma^2/a, 1) > 0$ proves

$$u^{\mathsf{T}} Q u \geq k\|u\|^2.$$

This floor uses the physical motion estimate, not an assumed lower bound for the full matrix. The local driving and flattening arguments are developed explicitly in (NF1)–(NF4).

The localized inverse-column construction also treats dormant coordinates exactly. The plate has identity row and column on each dormant coordinate, and its linear source there is zero. The same is true in every principal restriction. Its inverse therefore has identity row and column at those coordinates, by multiplying the inverse equation by that row. Padding a chosen principal inverse column into the full space preserves this property; off its footprint it is zero, and the diagonal belongs to the footprint. Symmetrizing these padded columns makes the reference covariance identity on the dormant rows and columns. The mean, defined as minus the transpose of the padded inverse times the linear source, is zero on dormant coordinates. These identities justify the dormant Gaussian factor used earlier. Positivity and the uniform upper bound of the full symmetrized padded covariance additionally require the localized inverse-window estimates.

6.4 Exponentially local inverses and a positive finite-window covariance

The covariance has its own construction An exponentially decaying inverse and a positive covariance on a finite window are separate objects. Truncating an inverse kernel need not preserve positivity. The argument below therefore constructs the window covariance with a positive margin and retains the comparison error in the norm used by the Gaussian integral. The later coordinates, mean and determinant refer to this particular covariance. They do not permit substitution of an arbitrary matrix with similar entrywise decay.

Let A be a real symmetric matrix indexed by coordinates placed at sites a_i. Suppose $A \geq rI$, $r > 0$, and its weighted absolute row and column sums at width $\eta \geq 0$ are at most $X > 0$. All principal restrictions inherit these statements: extend a vector by zero for the quadratic floor, and restrict the nonnegative row sums for the price. Choose $0 \leq \theta \leq 1$ with $\theta X \leq r/2$.

For $u \geq 0$ put $A_u = A + uI$. It is positive definite and has an actual two-sided inverse. Fix a column j, put $s_i = \exp(\theta\eta\,\mathrm{dist}(a_i, a_j))$, and set $y = A_u^{-1} e_j$, $w_i = s_i\, y_i$. Conjugation by the diagonal s gives

$$(A_u + E)w = e_j,$$
$$E_{ik} = (A_u)_{ik}(s_i/s_k - 1).$$

The triangle inequality for distance gives

$$|s_i/s_k - 1| \leq \theta \exp(\eta\,\mathrm{dist}(a_i, a_k)).$$

For completeness, the logarithm of s_i/s_k lies between $-\theta D$ and θD, where $D = \eta\,\mathrm{dist}(a_i, a_k) \geq 0$. Convexity gives $e^{\theta D} - 1 \leq \theta(e^D - 1) \leq \theta e^D$. Also $1 - e^v \leq e^{-v} - 1$ when $v \leq 0$. These two bounds prove the displayed inequality. The diagonal contribution of uI vanishes from E, so its absolute row and column sums are at most $\theta X \leq r/2$.

The elementary Schur quadratic estimate is $|v^{\mathsf{T}} E v| \leq (r/2)\|v\|^2$: expand the double sum and use $2|v_i v_k| \leq v_i^2 + v_k^2$, then apply the two sum bounds. Taking the inner product of the conjugated equation with w yields

$$(r/2+u)\|w\|^2 \le w_j.$$

With $c = \min(1, r/2)$, this implies $c(1+u)\|w\|^2 \le w_j \le \|w\|$, hence $\|w\| \le [c(1+u)]^{-1}$; the zero-vector case is immediate. Therefore

$$\begin{aligned}&\exp(\theta\eta\,\mathrm{dist}(a_i,a_j))|(A+uI)^{-1}_{ij}|\\ &\le [\min(1,r/2)(1+u)]^{-1}.\end{aligned}$$

This is uniform in matrix dimension and in all principal restrictions.

For the positive-width conclusion now take $\eta > 0$, set $\theta = \min(1, r/(2X))$, $b = \theta\eta$ and $k = b/3 > 0$. The preceding entrywise estimate also holds when $\eta = 0$; the following exponential site sum uses $k > 0$. If at most m coordinate labels occupy each site, the entrywise inverse estimate converts to row and column estimates at width 2k. Indeed the sum of $\exp(-(b-2k)\,\mathrm{dist})$ over torus sites is bounded by the corresponding absolutely convergent lattice sum, independently of the period. Multiply that sum by $m[\min(1, r/2)]^{-1}$; call the resulting price P. Thus both the full inverse and every zero-padded principal inverse column have price at most P at width 2k.

For each column j choose a principal window W_j containing all coordinates within distance B of a_j and contained within distance B+2. Its padded inverse column is

$$K_{ij} = \begin{cases} \left[(A|_{W_j})^{-1}\right]_{ij}, & i \in W_j,\\ 0, & i \notin W_j. \end{cases}$$

The error $E = AK - I$ vanishes exactly when $i \in W_j$, by the principal inverse equation. In particular it vanishes for $\mathrm{dist}(a_i, a_j) \le B$. Product price bounds at width 2k give $\mathrm{price}(E) \le XP+1$ there. On its remaining entries, losing one width unit k supplies $\exp(-kB)$, so at width k its price is at most

$$\varepsilon_B = \exp(-kB)(XP+1).$$

The resolvent identity $K - A^{-1} = A^{-1}E$ then gives width-k price at most $P\varepsilon_B$. These are bounds on actual matrices, not a declaration that the padded inverse is the full inverse.

Symmetrize $C = (K + K^{\mathsf{T}})/2$. Since A^{-1} is symmetric, averaging the two error bounds gives $\mathrm{price}(C - A^{-1}) \le P\varepsilon_B$, and $\mathrm{price}(C) \le P$. The Schur bound turns the error price into $|v^{\mathsf{T}}(C - A^{-1})v| \le P\varepsilon_B\|v\|^2$. Also $A \le XI$ implies $A^{-1} \ge X^{-1}I$. For example, put $w = A^{-1}v$ and use nonnegativity of $(w - X^{-1}v)^{\mathsf{T}}A(w - X^{-1}v)$, together with $A \le XI$. Expanding and using $Aw = v$ gives

$$0 \le (w - X^{-1}v)^{\mathsf{T}}A(w - X^{-1}v) \le v^{\mathsf{T}}A^{-1}v - X^{-1}\|v\|^2,$$

which proves $v^{\mathsf{T}}A^{-1}v \ge X^{-1}\|v\|^2$. Consequently, once $P\varepsilon_B \le 1/(2X)$,

$$(2X)^{-1}I \le C, |v^{\mathsf{T}}Cv| \le P\|v\|^2.$$

Thus C is the positive covariance required for the reference Gaussian, with the uniform upper bound $M = P$ used in the warning estimate. Its entries vanish beyond distance B+2. The residual AC−I has width-k price at most $XP\varepsilon_B$, by multiplying $A(C - A^{-1})$. All these choices are uniform in the torus volume and preserve the original endpoint-defined windows.

For the physical plate, the quadratic floor proved above supplies r, while the stencil bounds supply its width-η row price X and the coordinate census supplies m. At fixed block and stage, increasing B makes the exponential error small enough to meet the displayed requirement. The physical operator bounds and local driving estimate of this chapter establish these inputs for the plate.

6.5 Axial gauge and bounded logarithmic extension

Axial gauge is used locally to make a controlled logarithm available on a clean hull. Its chosen tree is an auxiliary coordinate device. The proof must therefore bound the resulting logarithm and preserve the observables when it extends links outside the hull. Subsequent comparison estimates depend on the size of the fixed inspected hull, not on a simultaneous global gauge choice for every query.

For a finite subset X of the integer cover, let $e_X(U)$ be the sum of the chords of its internal positive links and $f_X(U)$ the sum of the chords of its wholly internal elementary plaquettes. Boundary links and incomplete plaquettes contribute zero. Both sums are nonnegative, and f_X is gauge invariant.

An ordered list l of distinct coordinate axes defines a slab: coordinates in l range over the integers −r,...,r, while the other coordinates are pinned at o. From a point p in the slab, travel first in the first listed axis to the corresponding coordinate of x, then in the next axis, and so on. Let $H_p(x)$ be the ordered product of links along this route. Its endpoint is x because the remaining coordinates are pinned. For l followed by a new axis j, the gauge H_p makes every internal j-link equal to 1. Indeed the routes to x and $x+e_j$ have the same initial part; their final straight j-segments differ by precisely the link (x,j), so the gauge-transformed link cancels.

When these j-links equal 1, the plaquette relation gives

$$\begin{aligned}&|\operatorname{chord} U(x+e_j,i)-\operatorname{chord} U(x,i)|\\&\le \operatorname{chord}\operatorname{holonomy}(\operatorname{square}(i,j),x).\end{aligned}$$

This follows from the triangle inequality applied to a b^{-1} and its inverse, where a,b are the two horizontal links. For each site the right side is at most the sum of its internal plaquette chords; reversing the ordering of the two axes only inverts the holonomy and leaves its chord unchanged. Summing over the four positive axes and over a slice gives a bound of four times the plaquette sum in the adjacent layer. Telescoping between two slices therefore yields

$$|e_{\text{slice}}(z)U - e_{\text{slice}}(u)U| \le 4 f_{\text{slab}}(U).$$

All layers used by the telescoping sum lie inside the slab; adding the unused nonnegative layers proves this bound even when the two slices are far apart.

We prove by induction on $k = \text{length}(l)$ that some route root p satisfies

$$e_{\text{slab}}(\text{gauge}\, H_p U) \le 4k(2r+1) f_{\text{slab}}(U).$$

For $k = 0$ the slab is a single site with no internal links. For the induction step choose, among the 2r+1 slices normal to the new axis, one whose internal face sum is minimal. Its face sum is at most $f_{\text{slab}}(U)/(2r+1)$, because the slice face sets are disjoint subsets of the full slab face set. Use the induction hypothesis to choose a root on that slice, then extend its routes in the new axis. The edge sum on the chosen slice is unchanged, the new-axis links are 1, and the preceding slice comparison bounds every other slice by the chosen slice edge sum plus 4 $f_{\text{slab}}(U)$. Summing all 2r+1 slices gives the required bound with k+1. This also proves existence of the root; no curvature smallness was required.

For all four axes the slab is the box $[-r,r]^4$, so the bound is 16(2r+1) $f_{\text{box}}(U)$. If every plaquette chord in a periodic radius-r patch is at most q, then $f_{\text{box}} \le (2r+1)^4$ |PlaqAxes|q. Lift the patch to that box and apply the preceding construction. When $2r < n$, the lift is injective on the box, so H_p descends to a gauge on the corresponding torus sites, extended arbitrarily elsewhere. If $d+1 \le r$,

both endpoints of every radius-d bond lie in the box. Each such link chord is at most the full nonnegative edge sum, giving

$$\mathrm{chord}(\mathrm{gauge}\, zU)b \le 16(2r+1)^5 |PlaqAxes| q.$$

The logarithmic extension used after this gauge choice has an especially simple construction. Analyticity of the local logarithm at the identity, $\log(1) = 0$, gives $\|\log M\| \le D\|M-1\|$ on a small matrix ball. Shrink the ball so that $\exp(\log U) = U$ for group elements in it. Since

$$\|\,\mathrm{mat}(U) - 1\| = \mathrm{colorScale} \cdot \mathrm{chord}(U),$$

there are $C, \delta > 0$ with $\exp(\log U) = U$ and $\|\log U\| \le C\,\mathrm{chord}(U)$ whenever $\mathrm{chord}(U) \le \delta$. Define the logarithmic field linkwise by log U on the prescribed interior patch and by zero outside it. This field is globally bounded by Ch and exponentiates to U on that patch whenever its links have chords at most $h \le \delta$. No continuity across the boundary of the finite patch is needed.

We will also use two consequences of the same analytic Taylor argument. For an analytic map f at 0 with directional derivative L, the analytic series identifies L with the bounded derivative $Df(0)$. Its first two terms and the uniform series remainder give

$$\|La\| \le K\|a\|, \|f(a) - f(0) - La\| \le K\|a\|^2$$

on one fixed small ball. Applied to the matrix exponential this bounds chord(exp a) by $K\|a\|$. Applied to the curvature of exp a_0 exp a_1 exp$(-a_2)$ exp$(-a_3)$, its derivative is $a_0 + a_1 - a_2 - a_3$. Hence the difference between the actual plaquette curvature of exp a and its discrete curl is at most $K\|a\|^2$, uniformly in the plaquette and lattice size. Repeated edges cause no change: each of the four coordinates is bounded by the same field norm.

Finally,

$$\|Ad(U)a - a\| \le 2\,\mathrm{colorScale} \cdot \mathrm{chord}(U)\|a\|.$$

In matrix coordinates write $UXU^* - X = (U-1)XU^* + X(U^*-1)$, use unitary invariance of the Frobenius norm, and use the chord identity above together with $\|\,\mathrm{colorMat}(a)\| = \mathrm{colorScale}\,\|a\|$. Induction along a walk likewise gives $\mathrm{chord}(\mathrm{travel}(U,w)) \le \mathrm{length}(w)h$ if every link read by the walk has chord at most h; inverses have the same chord.

6.6 Quantitative nonlinear and flat row comparison

Put $\mathrm{widen}(R,h,\tau) = (1+R)h + \exp(-\tau R)$. All cutoffs in this paragraph use actual path length, including winding, rather than the distance between the two torus endpoints.

First consider the plotted curvature field. A path slot consists of a walk from its reading site to a plaquette, followed by adjoint transport of that plaquette curvature. If $\|a\| \le h$ is in the chart ball, the preceding estimates give link chords $\le B_{\mathrm{ch}}$ h, curl norm $\le A_{\mathrm{ch}}$ h, and curvature-minus-curl norm $\le A_{\mathrm{ch}}\, h^2$. For a retained slot of physical length at most R, transport changes the curl by at most

$$2\,\mathrm{colorScale}\, RB_{\mathrm{ch}}h \cdot A_{\mathrm{ch}}h.$$

Together with the quadratic curvature error this bounds slot-minus-curl by

$$Jh\,\mathrm{widen}(R,h,\tau), J = A_{\mathrm{ch}}(1 + 2\,\mathrm{colorScale}\, B_{\mathrm{ch}}).$$

For an omitted slot of length ℓ with $d\ell > R$ and stage width $\sigma > 0$,

$$1 \le \exp(\sigma\ell)\exp(-\sigma R/d) \le \exp(\sigma\ell)\,\mathrm{widen}(R, h, \sigma/d).$$

This pays its whole linear curl with the exponentially weighted path budget. If the path coefficients u_i obey $\sum_i |u_i|\exp(\sigma\ell_i) \le paths$, absolute summation of retained errors and omitted linear terms proves

$$\|x_{\mathrm{lin}}(a)_j\| \le Ch,$$
$$\|\,\mathrm{plotted}(\exp a)_j - x_{\mathrm{lin}}(a)_j\| \le Ch\,\mathrm{widen}(R, h, \sigma/d),$$

where $C = \sqrt{\mathrm{direct}}A_{\mathrm{ch}} + \mathrm{paths}\,J$. The direct-curvature coordinate uses just the quadratic face estimate and $h \le$ widen. The path coordinate uses the same summable envelope $|u_i|\exp(\sigma\ell_i)$; thus subtraction and summation are justified by absolute convergence, including countably many slots. Constants and the small chart radius depend on the fixed stage and d, not on torus size or R.

For the kernel rows, let $K_{ab}(u)$ be a cover coefficient with

$$\sum_u \exp(\eta|u|)|K_{\mathrm{ab}}(u)| \le A$$

for each pair of finite residue types a,b. A routed term is its scalar multiple of adjoint transport along its designated walk, retained when length $\cdot d \le R$, and included only in the endpoint congruence class. Suppose length $\cdot d \le D(1 + |u|)$. The full flat term is $K_{ab}(u)$ times the identity on color space, without the cutoff. Unitary adjoint action bounds the operator norm of each routed term by $|K_{ab}(u)|$.

If link chords are $\le C_{\mathrm{ch}}$ h, a retained transport differs from the identity by at most 2 colorScale C_{ch} R h. For an omitted term the length inequality implies

$$\eta R/D \le \eta + \eta|u|,$$
$$1 \le \exp\eta \cdot \exp(\eta|u|)\exp(-\eta R/D).$$

Thus both kinds of difference are bounded by

$$J\,\mathrm{widen}(R, h, \eta/D)\exp(\eta|u|)|K_{\mathrm{ab}}(u)|,$$
$$J = 1 + 2\,\mathrm{colorScale}\,C_{\mathrm{ch}} + \exp\eta.$$

Sum first on the cover. Each u contributes to exactly its endpoint congruence class; summing over endpoints therefore costs only the finite residue census, even on short tori with many windings. Absolute convergence allows the order of these nonnegative bounds to be interchanged. If there are n_a source types and n_b target types, the convenient common row and column bound is

$$C_K = (n_a + n_b)AJ.$$

Both the cut nonlinear matrix and the full flat matrix have row and column sums at most C_K, and their difference has row and column sums at most C_K widen$(R, h, \eta/D)$.

For the same-level lanes, the ordered coordinate route to u has length |u| and $d = L$. For the fine-to-coarse row, its route goes from a residue lift to Lu; its length is at most L|u|+4L. Both cases satisfy the preceding inequality with $D = 5L$. The replica background has either identity links or original coarse links, so it preserves the required chord bound. Apply this calculation to the F, I and N rows with

$$A = \max(\max(0, E), E/L^2), \tau = \eta/(5L).$$

Take the maximum of 1 and the three corresponding C_K values, and the common exponential-chart radius. This proves all the kernel-chart estimates used below, uniformly in the period and cutoff, with their row and transpose-row bounds. All constants are chosen after the block length L has been fixed.

6.7 Clean-hull guide comparison

The guide is the pair of coarse response fields Mv,Mw, while its linear comparison is obtained by applying the flat F and I rows to the linearized field $x_{\mathrm{lin}}(a)$. There is also an N row. Its linear comparison vanishes identically: the supplied stationary identity is flatN $\cdot D_{\mathrm{full}} \cdot d = 0$, and x_{lin} has precisely the coordinate form $D_{\mathrm{full}} \cdot d$ a. This includes harmonic coordinates; no scalar-mode deletion is used.

On a clean radius-B/4 hull agreeing with exp a, the clipped vision field equals the unclipped plotted field of exp a at every coordinate read by a guide row. The row's reach is at most $R =$ till(x); the plotted field needs only R+4 further sites, still within the same hull by the room inequalities. Cleanliness makes the retained indicator equal to one and makes the clipping factor $\min(1, \mathrm{temp} \,/\, \mathrm{toll})$ equal to one. The row itself can also be replaced by its value at exp a by its local dependence on this hull.

Here is the resulting quantitative calculation. Suppose the nonlinear row Q has row bound X and differs from its flat row D by row bound XW. Suppose also

$$\|x_{\mathrm{lin}}(a)_j\| \le Z, \|\,\mathrm{plotted}(\exp a)_j - x_{\mathrm{lin}}(a)_j\| \le U.$$

Writing $B = Q(\exp a)$, $f = \mathrm{plotted}(\exp a)$, and $g = x_{\mathrm{lin}}(a)$, the exact decomposition

$$Bf - Dg = B(f - g) + (B - D)g$$

and the row-sum inequality give $\|Bf - Dg\| \le XU + XWZ$. All equalities before this bound are restricted to the same clean hull.

In the source the uniform kernel-chart and plotted-chart estimates supply positive constants X, Z_0, s, t and

$$\begin{aligned} M(x) &= Z_0 H(x), \\ U(x) &= M(x)\,\mathrm{widen}(\mathrm{till}(x), H(x), t), \\ e(x) &= XU(x) + X\,\mathrm{widen}(\mathrm{till}(x), H(x), s)M(x). \end{aligned}$$

Here t is the lower-stage width divided by the positive block length. If H has rate $-1/2$ and is eventually nonnegative, broadening by widen preserves that rate: its bound is a radius factor times H plus an exponentially small cutoff term. Products in the displayed formula therefore have rate -1. The same e bounds the two guide differences and the nonlinear N guide itself, since its linear part is zero. The kernel-chart and plotted-field estimates of this chapter give these constants.

6.8 Stationary cancellation in every signed orientation

The stationary comparison uses a single supplied pair of quadratic beds and their common diagonal and roll identities. Their attached path realizations, positive floors, and invariance under signed coordinate changes are chosen together. Let W be the installed stationary map, U its extracted fine-field matrix for the chosen forest, and $D_0 = I -$ probe $\cdot U$. The diagonal action identity reads

$$K_{\mathrm{coarse}} = U^{\mathsf{T}} K_{\mathrm{fine}} U + penalty \cdot D_0^{\mathsf{T}} D_0.$$

For coarse logarithmic data a, the shifted fine field has color-coordinate vectors U a, while its residual has vectors D_0 a. Pairing this matrix identity with each color coordinate and summing gives

$$\begin{aligned}&\text{fineEnergy}(shifteda) + penalty\|\,\text{residual}(a, shifteda)\|^2\\&= \text{coarseEnergy}(a).\end{aligned}$$

This is an identity of quadratic forms, including zero harmonic modes.

The same identity holds after any signed coordinate change. To justify the noise part, apply the exact signed covariance of the actual nonlinear predictor and noise energy to the exponentiated fields exp(t a), exp(t shifted a), divide by t^2, and pass to their already established quadratic limits. The invariant fine and coarse bed metrics give the other two terms. Thus no isotropy of individual matrix entries is assumed.

Write w for the coarse linear moment field of the transformed input, u for the fine linear moment field of its transformed shifted field, and v for its linear noise field. The attached path realization and floor identify moment-square sums with the corresponding bed energies. The preceding energy identity is therefore

$$\|u\|^2 + \|v\|^2 = \|w\|^2,$$

where the norms here are finite counting sums of squared color norms. Let F,I,N be the full flat rows, and put $U_1 = Fw$, $V_1 = Iw$. The N row annihilates w. The diagonal roll and cross roll identities then give

$$w = F^{\mathsf{T}}U_1 + I^{\mathsf{T}}V_1, w = F^{\mathsf{T}}u + I^{\mathsf{T}}v.$$

The transpose here means the adjoint for the counting inner products. Pairing the first equality with w gives $\|U_1\|^2 + \|V_1\|^2 = \|w\|^2$. Pairing the second with w gives $\langle U_1, u\rangle + \langle V_1, v\rangle = \|w\|^2$. Expand squared differences and use these three equalities to obtain

$$\|U_1 - u\|^2 + \|V_1 - v\|^2 = 0.$$

Every summand is nonnegative, so $u = U_1$ and $v = V_1$ at every row. This proves equality of linearCost and linearGuide in every signed orientation. Subtracting the linear copy then gives precisely the uncut spring cancellation used below.

6.9 Nearest group points and the analytic focus

All matrix norms in this paragraph are Frobenius norms, with the real inner product. Compactness of the closed unitary matrix group gives a nearest group element to every matrix X. Left or right multiplication by a group element preserves this minimization property, by unitary invariance of the norm.

If 1 is nearest to X, differentiate the squared distance $\|\exp(tY) - X\|^2$ at its minimum $t = 0$ for every Lie algebra matrix Y. The derivative is zero, so X−1 is orthogonal to the Lie algebra. In general, if u is nearest to X, left multiplication by u^{-1} gives

$$u * X - 1 \in \text{lie}(G)^{\perp}.$$

Here u* denotes the adjoint matrix of u.

Consider the tubular map

$$T(a,n) = \exp(\text{colorMat}(a))(1+n), n \in \text{lie}(G)^{\perp}.$$

It is analytic and its derivative at (0,0) is colorMat(a)+n. This is an isomorphism by the orthogonal tangent/normal decomposition. The inverse function theorem therefore gives a local analytic inverse near $X = 1$. Define focus(X) as the exponential of the tangent coordinate of that inverse.

Every nearest point must agree with focus sufficiently near 1. To prove this without assuming continuity of a minimizing selection, choose any nearest $b(X)$. Its minimizing property gives $\|b(X)-X\| \le \|1-X\|$, hence $\|b(X)-1\| \le 2\|X-1\|$. Thus $b(X) \to 1$. Its normal coordinate $n(X) = b(X)$*X−1 tends to zero and is orthogonal to the Lie algebra by the preceding variational identity. Its tangent coordinate $a(X) = \log b(X)$ also tends to zero. Moreover $T(a(X), n(X)) = X$ exactly. Both coordinates therefore lie in the inverse-function neighborhood and equal its unique inverse coordinates, proving $b(X) = \text{focus}(X)$. If there were exceptional nearest points arbitrarily near 1, choosing one at each exceptional X would contradict this conclusion. Thus every nearest point equals focus there.

This proves local uniqueness simultaneously over all components of the group. Translating X by a nearby group element extends the same uniqueness radius to a tube about the whole group. Differentiating the tubular inverse gives

$$D(\text{mat} \circ \text{focus})(1) = orthogonalProjection_{\text{lie}}(G).$$

Indeed the inverse derivative sends colorMat(a)+n to (a,n), and differentiating the exponential discards n. Analyticity makes this a strict derivative. Consequently, for every $\varepsilon > 0$, after shrinking the neighborhood of 1, nearest points u of X and v of Y satisfy

$$\|\,\text{mat}(u) - \text{mat}(v) - \text{projection}_{\text{lie}}(G)(X-Y)\| \le \varepsilon \|X-Y\|.$$

These prove the focus properties used in the frame and predictor derivatives from the local group chart and the ordinary inverse function theorem.

6.10 Analyticity of the normalized frame

This is a fixed-period assertion; no volume-independent analytic radius is claimed here. Let v be the logarithmic input and H_p its actual frame matrices at the patch points, with $H_{p_0} = 1$ at the distinguished center. At $v = 0$ all path products equal 1. The identity frame minimizes the nonnegative frame score with score zero; uniqueness of the minimizer identifies the selected frame with it. The minimum branch has an open neighborhood of this input because its cutoff is positive. Continuity of the unique selected minimizer therefore gives continuity of all H_p at zero. On that neighborhood all frame tags select this same frame.

The frame self-consistency condition says that H_p is a nearest group element to

$$cloud_p(v,H) = \text{average}_q \,\text{average}_o \, H_q \cdot \text{travel}_v(p \to qinordero) * .$$

This finite expression is analytic in v and the matrix array H: its path products are finite products of exponentials and adjoints. At $(v,H) = (0,1)$, every cloud is 1. Continuity then puts every cloud in the neighborhood where its nearest group element equals the analytic focus map. Thus the actual frame solves the analytic equations

$$F_p(v,H) = H_{p_0} - 1 \quad \text{if } p = p_0,$$
$$F_p(v,H) = H_p - \text{focus}(cloud_p(v,H)) otherwise.$$

Here focus is viewed in matrix coordinates. To apply the implicit function theorem, compute the derivative in the H variable at (0,1). If an array h belongs to its kernel, then $h_{p_0} = 0$ and

$$h_p = Aforp \neq p_0, A = \text{projection}(\text{average}_q\, h_q).$$

Let N be the positive number of patch points. Averaging gives average $h = (N-1)A/N$. Since projection$(A) = A$, its defining equation yields $A = (N-1)A/N$, hence $A = 0$ and $h = 0$. The derivative is an injective endomorphism of a finite-dimensional matrix-array space, so it is invertible.

For precision, an analytic equation $F(v, H(v)) = 0$ and a continuous solution $H(v)$ with invertible H-derivative identify that solution with the analytic implicit branch: apply the inverse function theorem to $(v, H) \mapsto (v, F(v, H))$. Continuity puts the given solution inside its local inverse domain, where uniqueness forces equality. This proves analyticity of the actual frame. It supplies the analytic frame input used in the following predictor calculation. The nearest-point, self-consistency and selected-minimizer arguments in this chapter establish the geometric properties used here.

6.11 First derivative of the actual rooted predictor

For coarse root y and patch points p,q, let $\text{fee}_y(p, q, o; a)$ be the tangent link sum along the actual ordered route from p to q in order o, embedded at y. Define

$$\text{tap}_y(p; a) = -\,\text{average}_q\, \text{average}_o\, fee_y(p, q, o; a).$$

Each color coordinate of tap is the corresponding scalar path functional $\tau_y p$ applied to the color-coordinate input. The derivative of a route product at the identity is colorMat of its link sum, by differentiating its finite product of exponentials; reversed links contribute with a minus sign.

Let $H_y(p; a)$ be the actual normalized frame matrix near zero, and let d_p be its derivative in direction a. The local frame equation says that $H_y(p)$ is the nearest-group focus of its averaged cloud. Differentiating that cloud gives $\text{average}_q\, d_q + \text{colorMat}(\text{tap}_y(p; a))$. The derivative of focus at 1 is the orthogonal projection onto the matrix Lie algebra. Hence

$$d_p = \text{projection}(\text{average}_q\, d_q) + \text{colorMat}(\text{tap}_y(p; a)).$$

The frame at the distinguished patch center p_0 is identically 1, so $d_{p_0} = 0$. Subtracting its equation eliminates the common projected average:

$$d_p = \text{colorMat}(\text{tap}_y(p; a) - \text{tap}_y(p_0; a)).$$

The local analytic frame equation and its normalization are the inputs to this differentiation; the subtraction shows explicitly why the derivative is independent of the frame tag.

For a coarse bond (y,i), a connector consists of the left frame, the actual length-L bond route, and the inverse right frame. Its derivative is therefore colorMat of

$$\begin{aligned} \text{bill}_p(a) &= \text{tap}_y(p; a) - \text{tap}_y(p_0; a) \\ &+ linkSum(bondRoutefromp; a) \\ &- \text{tap}_{(y+e_i)}(p; a) + \text{tap}_{(y+e_i)}(p_0; a). \end{aligned}$$

Near zero every connector is near 1. The best-predictor branch is open and is used at zero because its cost there is zero and its cutoff is positive. Thus the actual predictor near zero is the nearest-group focus of the weighted connector average, irrespective of its fallback tag. The connector average equals 1 at zero. It is analytic because the frame and finite route maps are analytic, and focus is analytic near 1; consequently the actual predictor is analytic near zero.

Differentiating the weighted average gives colorMat($\sum_p weight_p \, \mathrm{bill}_p(a)$). This lies in the Lie algebra, so the focus derivative leaves it unchanged. Expanding the definitions of tap, the rooted path functionals, and the probe matrix shows coordinate by coordinate that

$$\sum_p weight_p \, \mathrm{bill}_p(a) = \mathrm{probe}(a).$$

The patch weights sum to one; this retains both frame-center subtraction terms. Thus the predictor equals 1 at zero and its literal first derivative is colorMat(probe(a)), exactly the derivative used in the noise Taylor estimate. The remaining foundation of this calculation is the local analytic frame construction and focus derivative, rather than any limiting or averaged predictor.

6.12 Uniform predictor linearization and entry bounds

For a fixed patch and kernel, the predictor at one coarse bond reads a finite cover stencil. Include all vertices on every ordered route between patch points, the translated copy of this set at the other bond endpoint, and every straight length-L connector from a patch point. Call their union $S_p red$. If two cover fields agree on links rooted in $S_p red$, every route product, axial frame, via-frame value and score agrees. Their minimizing frame also agrees by the supplied frame uniqueness; in the fallback branch their axial frame agrees directly. Thus their actual predictor values agree with the same endpoint tags and connector tag. Taking the maximum route distance of these finitely many roots defines predictionReach. Agreement on that parent-distance hull therefore suffices for predictor equality.

At fixed volume and fixed tags define

$$Z_k(a,b) = \sqrt{\mathrm{penalty}} \cdot \mathrm{curvature}(\exp b \cdot \mathrm{prediction}(\exp a, k)^{-1})$$

at the specified coarse bond. This is analytic near zero by the local analytic predictor theorem, matrix multiplication and adjoint, and the linear matrix formula for curvature. At zero the predictor is 1 and $Z_k = 0$. Its derivative is

$$z_{\mathrm{lin}}(a,b) = \sqrt{\mathrm{penalty}} \cdot (b - \mathrm{probe}(a))$$

in the physical color coordinates: differentiate the predictor to obtain colorMat(probe(a)), differentiate exp b, and use skew adjointness to differentiate the inverse. The curvature derivative recovers the difference. Taylor's theorem gives a linear bound and a quadratic remainder. Since there are finitely many tags and bond positions on a fixed torus, a maximum of their constants and a positive minimum of their radii makes the bounds simultaneous there.

To remove volume dependence, choose one reference torus with coarse side $c_0 > 1$ so large that every point of $S_p red$ lifts uniquely from its fine torus. One explicit choice is $c_0 = 2M + 2$, where M is the largest absolute coordinate of a stencil point. For any input torus and any bond rooted at x, fill the reference fine logarithmic field by pulling back the actual input values under

$$y \mapsto wrap_f(lift(y - embed(0))) + embed(x).$$

Fill its coarse logarithmic field constantly with the actual coarse value at the bond being read. Neither operation increases the supremum norm. Put the original left and right frame tags at 0 and e_i of the reference coarse torus; these sites are distinct because $c_0 > 1$. Put the original connector tag at the reference bond. All predictor stencil reads then agree exactly for every scalar multiple $t(a,b)$, even if the original torus identifies different stencil points. Thus the nonlinear Z values agree

at $t = 1$ and their linear terms agree by uniqueness of the derivative at $t = 0$. The reference-torus Taylor bounds consequently give constants C_z, r_z independent of the input volume and all its tags:

$$\|z_{\mathrm{lin}}(a,b)\| \le C_z \max(\|a\|,\|b\|),$$
$$\|Z_k(a,b) - z_{\mathrm{lin}}(a,b)\| \le C_z \max(\|a\|,\|b\|)^2$$

when the maximum is at most r_z. The local analytic predictor theorem and its derivative formula are the inputs to this Taylor argument; the transplantation adds no new branch choice.

The entry field is the direct sum of the fine plotted moment and Z. Its linear counterpart is the direct sum of x_{lin} and $z_l in$. For logarithmic inputs bounded by H, combine the plotted comparison already proved with the last inequality. If C_x is its plotted constant, set

$$M = (C_x + C_z)H,$$
$$W = C_x H \,\mathrm{widen}(R, H, width_{\mathrm{top}}) + C_z H^2.$$

Then each linear entry has norm at most M and each entry-minus-linear-entry has norm at most W, simultaneously in the two kinds of row. For $R = \mathrm{till}(x)$ and H of rate $-1/2$ these have rates $-1/2$ and -1. The entry is gauge covariant by predictor covariance and curvature conjugation. It is local on the protected hull: the plotted component uses its cutoff stencil and the Z component uses predictionReach together with the coarse link at the row root. These establish the entry-near and entry-locality inputs used next.

For completeness, the full and cutoff linear rows have uniform stencil bounds at every cutoff, not merely at till(x). A discrete curl has norm at most $4\|v\|$. Summing a path row against its weighted absolute coefficient budget gives $\|ellLine(v)\| \le 4\ paths\|v\|$. Terms omitted beyond R cost at most $4\ paths\|v\| \exp(-width_t opR)$. The direct-curvature term has bound $4\sqrt{\mathrm{direct}}\|v\|$; the noise term has bound $C_z\|v\|$. Thus, writing

$$U_{\mathrm{st}} = 4(paths + \sqrt{\mathrm{direct}}) + C_z,$$
$$K_{\mathrm{st}} = U_{\mathrm{st}} + 4paths,$$

both full and cutoff row norms are at most $K_{\mathrm{st}}\|v\|$ and their difference is at most $K_{\mathrm{st}}\|v\| \exp(-width_t opR)$. The cutoff bound follows by comparing to the full row and using $\exp(-width_t opR) \le 1$. Absolute summability justifies every linear path sum and its subtraction.

These bounds also give the spatial impulse estimate used in the plate prices. If a single-link input of color a is at parent distance g from a row root, the cutoff row at s is zero whenever $s +$ predictionReach $+4 < g$. For $g >$ predictionReach $+5$ choose $s = g -$ predictionReach -5. Compare both cutoffs R and s to the full row and apply their tail bounds; if R is too short the row is already zero. The result, also covering the nearby case by the uniform cutoff bound, is

$$\exp(width_{\mathrm{top}} g)\|barField_R(single_{\mathrm{link}} a)\|$$
$$\le 2K_{\mathrm{st}} \exp(width_{\mathrm{top}}(\mathrm{predictionReach}\ +6))\|a\|.$$

6.13 From a shell comparison to the tree-slice link estimate

For a differentiable group-valued curve $U(t)$ with $U(0) = 1$ and matrix derivative colorMat(a), the norm of its difference quotient gives

$$\mathrm{chord}(U(t))/|t| \to \|a\| ast \to 0, t \ne 0.$$

Indeed chord is the Frobenius distance from 1 divided by colorScale, and $\|\operatorname{colorMat}(a)\| = \operatorname{colorScale}\|a\|$. This preserves the physical normalization of the color norm.

Let w be a fine logarithmic field vanishing on the trunks of the given forest. Then $\exp(tw)$ lies on that exact tree slice for every t. The actual frame minimum branch is active near $t = 0$ at every coarse site; its frame mass is positive, so the chosen tags are usable there. Apply the nonlinear shell link comparison at the link e to $U(t) = \exp(tw)$, coarse comparison field 1, and its actual predictor $Q(t)$. It gives

$$\begin{aligned}\operatorname{chord}(U(t)_e) \le C\operatorname{port}(\operatorname{parent}(e), U(t))\\ + \operatorname{chord}\big(Q(t)_{(\operatorname{parent}(e),\operatorname{axis}(e))}\big).\end{aligned}$$

This upper bound is valid in the internal-block case as well, where adding the last nonnegative term is harmless. The port sum is bounded by the full macro-box sum of plaquette chords, including repetitions on a periodic lattice. Divide by |t| and pass to the limit. Each plaquette derivative is its discrete curl, and the actual predictor derivative was computed above. The finite macro sum commutes with this limit. We obtain

$$\begin{aligned}\|w_e\| \le C \sum_{z\in\operatorname{box}(\operatorname{macroR})} \sum_p \|\operatorname{curl}(w)_{(\operatorname{wrap} z+\operatorname{embed}(\operatorname{parent}(e)),p)}\|\\ +\|\operatorname{probe}(w)_{(\operatorname{parent}(e),\operatorname{axis}(e))}\|.\end{aligned}$$

This proves the local link estimate from the nonlinear shell comparison without a separate linear gauge choice.

The same calculation gives a volume-independent counting bound. Put $N = |\operatorname{box}(\operatorname{macroR})|$, $m = |PlaqAxes|$ and $q = L^4$. For any plaquette function f, Cauchy–Schwarz on the macro sum costs Nm. For each fixed offset z, the map (coarse site,plaquette *type*) $\mapsto$ (*wrap* z+embed(site),plaquette type) is injective, so summing its squares costs at most the full fine plaquette sum. Summing all N offsets adds the factor N. Each coarse site has q fine children and four outgoing fine-axis choices. Thus squaring the local link inequality with $(a+b)^2 \le 2a^2 + 2b^2$ gives

$$\begin{aligned}\sum_e \|w_e\|^2 \le \beta\sum_p \|\operatorname{curl}(w)_p\|^2 + \gamma\sum_b \|\operatorname{probe}(w)_b\|^2,\\ \beta = 2C^2\cdot q\cdot 4\cdot N^2 m, \gamma = 2q.\end{aligned}$$

This counting uses injectivity in the coarse site for each fixed offset, not injectivity of different offsets; repeated periodic plaquettes remain allowed.

If the fine quadratic bed has curl floor $a > 0$ and the predictor penalty is $\kappa > 0$, its vertical Hessian dominates $a\sum\|\operatorname{curl} w\|^2 + \kappa\sum\|\operatorname{probe} w\|^2$. The restriction to the vertical tree chart has the same Hessian as the full germ on directions (0,v), by the linear inclusion chain rule; the collected vertical field is v off trunks and zero on trunks. Therefore, with

$$\eta = \min(a,\kappa)/(1+\beta+\gamma) > 0,$$

the vertical Hessian is at least $\eta\sum_{\text{free}}\|v_e\|^2$, since $\eta\beta \le a$ and $\eta\gamma \le \kappa$. This also unfolds the fixed-block counting-coordinate Hessian floor. The shell comparison supplying C is described next.

6.14 Uniform branch cutoffs and continuity of selected minima

The finite-pair uniqueness result gives a positive d depending only on G. Choose the frame cutoff at most d/colorScale. For any patch satisfying the all-minimum criterion, use its center-based axial

frame a and define pair transports $W_{xyk} = a_x \ \mathrm{travel}(x \to y, k) a_y^{-1}$. The criterion bounds their chords by the frame cutoff, so their matrix distances from 1 are at most d. For an arbitrary frame b, the transformed array $b_x \ a_x^{-1}$ has pair cost exactly the original frame score. A normalized global minimizer of the frame score is a global minimizer of this pair cost: every unnormalized competitor can be multiplied on the left to fix its center, without changing its score. Pair uniqueness then gives uniqueness up to a left multiple, and the center value 1 makes that multiple equal to 1. Existence follows by compactness of the finite product of G with its center fixed. This proves the uniform frame cutoff independently of patch order and size.

The predictor cutoff follows from a related weighted-centroid calculation. For real weights w_i summing to one and group points f_i, set $m = \sum w_i \ \mathrm{mat}(f_i)$. Expansion of the real inner product gives

$$2repDim \sum_i w_i \, \mathrm{price}(f_i u^{-1})$$
$$= \sum_i w_i \| \mathrm{mat}(f_i) \|^2 - \|m\|^2 + \|m - \mathrm{mat}(u)\|^2.$$

Thus minimizing the weighted price is equivalent to being nearest to m. For nonnegative weights, the squared norm inequality for a weighted average shows

$$\|m - \mathrm{mat}(u)\|^2 \le 2repDim \sum_i w_i \, \mathrm{price}(f_i u^{-1}).$$

If the right side is below the square of the nearest-point tube radius r, m lies within r of the group and has a unique nearest point. Choosing predictor cutoff at most $r^2/(2 \, \mathrm{repDim})$ proves the uniform predictor uniqueness condition. Both cutoff choices depend only on G.

The selected minimizers are continuous wherever their defining fibers are unique. Here is the compactness argument including arbitrary selections away from that set. Let $M \subseteq X \times Y$ be the closed minimizer relation, with Y compact. For an open neighborhood V of the unique minimizing value at x_0, project $M \cap (X \times V^c)$ to X. This projection is closed because Y is compact, and it omits x_0 by uniqueness. Its complement is therefore a neighborhood of x_0 on which every selected minimizer lies in V. This is continuity at x_0. For frames the relation is closed because it is the intersection of the closed center-normalization condition and the score inequalities against all normalized competitors; scores are continuous finite path expressions. The predictor minimizer relation is closed by the corresponding continuous cost inequalities. Applying the compactness argument proves continuity on the respective minimum branches. Their strict, finite cutoff tests make those branches open. These complete the uniqueness and continuity inputs used in the analytic-frame construction.

6.15 Minimizer self-consistency and uniqueness

For a finite pair family define $\mathrm{term}(h; x, y, k) = \mathrm{price}(h_x W_{xyk} h_y^{-1})$, and cost as its average over all three indices. Reversal of the pair exchanges x,y and replaces k by its mirror permutation, preserving the summed term. The diagonal terms vanish because $W_{xxk} = 1$. Consequently, replacing only h_x by a competitor b changes the total cost by twice the change in the off-diagonal x-row. Global minimality implies that this off-diagonal row is minimal. Adding the frozen self term $\mathrm{price}(b h_x^{-1})$, which is nonnegative and is zero at $b = h_x$, proves the stronger comparison with the whole row in which every h_y, including h_x, is held fixed.

The normalized whole-row price is precisely the weighted squared-distance problem to the matrices $h_y \ W_{xyk}^{-1}$, up to a positive constant and a term independent of b. Expanding the squared

Frobenius norms, or completing the square about their average, identifies its minimizers with nearest group points to feedback$(h)_x$. Thus h_x is nearest to its feedback. This proves the self-consistency assertion used in both alignment arguments.

A uniform small-transport condition gives uniqueness modulo a common left multiple, independently of the numbers of sites or pair labels. Suppose $\|\operatorname{mat}(W_{xyk}) - 1\| \le d$. Comparison with the constant array bounds the average squared pair separation of any minimizing h by $4d^2$. If 1 is nearest to its centroid, averaging this bound at a suitable index gives $\|\operatorname{centroid}(h) - 1\| \le 2d$. The feedback differs from the centroid by at most d, so every feedback is within 3d of 1.

Choose δ from the focus two-point estimate with coefficient 1/4, and $d = \min(\delta/4, 1/4)$. Given two minimizers, multiply each on the left by the inverse of a nearest point to its own centroid. This preserves minimization and makes 1 nearest to both centroids. Let $u = \max_x \|\operatorname{mat}(h_x) - \operatorname{mat}(k_x)\|$ and take an index attaining it. Write $B = \operatorname{feedback}(h)_x - \operatorname{feedback}(k)_x$. Unitary invariance gives $\|B\| \le u$. Subtracting the centroid difference before estimating gives the sharper bound

$$\|B - (\operatorname{centroid}(h) - \operatorname{centroid}(k))\| \le ud,$$

because each summand factors as $(\operatorname{mat}(h_y) - \operatorname{mat}(k_y))(\operatorname{mat}(W_{xyk})^{-1} - I)$. Both centroid-minus-1 matrices are normal to the Lie algebra by the nearest-point normal equation. Their difference has tangent projection zero. Thus $\|\operatorname{projection}(B)\| \le ud$. Self-consistency and the focus estimate, applicable since $3d < \delta$, now imply

$$u \le \|\operatorname{projection}(B)\| + \tfrac{1}{4}\|B\| \le ud + u/4 \le u/2.$$

Therefore $u = 0$. Undoing the two centroid normalizations shows that the original minimizers differ by one common left group element. A fixed center value removes this last freedom.

6.16 Finite-pair alignment and the route census

The synchronization estimates compare overlapping local predictor charts. They cannot identify all gauges by declaration: each query has its own route, and every route contributes to the finite multiplicity count. The common macro gauge is obtained only after these comparison errors have been estimated. This supplies the coordinate compatibility used by the affine rows and protected mean in the following steps.

Let W_{xyk} be a finite family of group-valued pair transports with its reversal symmetry, and let h minimize the averaged pair price of $h_x\ W_{xyk}\ h_y^{-1}$. Put $\operatorname{income}(x) = \operatorname{average}_{y,k} \operatorname{chord}(W_{xyk})$ and $m = \operatorname{average}_x \operatorname{income}(x)$. Comparison with the identity array and $\operatorname{chord}^2 \le 2$ chord show that its minimum price is at most m. Factoring $h_x\ h_y^{-1}$ through the pair discrepancy gives

$$\operatorname{average}_{x,y} \|\operatorname{mat}(h_x) - \operatorname{mat}(h_y)\|^2 \le 8 \operatorname{colorScale}^2 m.$$

Indeed the squared triangle inequality bounds each $\operatorname{chord}(h_x h_y^{-1})^2$ by 2 times the squared discrepancy plus 4 $\operatorname{chord}(W_{xyk})$. Averaging and using the minimizing comparison proves the factor 8.

Let $\operatorname{centroid}(h) = \operatorname{average}_x \operatorname{mat}(h_x)$ and $\operatorname{feedback}(h)_x = \operatorname{average}_{y,k} \operatorname{mat}(h_y W_{xyk}^{-1})$. Direct subtraction gives

$$\|feedback(h)_x - \operatorname{centroid}(h)\| \le \operatorname{colorScale} \cdot \operatorname{income}(x).$$

A minimizer is nearest to its feedback at each coordinate; this coordinatewise variational assertion is the finite-pair self-consistency input. Choose a nearest group point b to the centroid and

replace h by $b^{-1}h$. The centroid now has nearest point 1. By choosing one index whose averaged squared pair distance is at most the global average and applying the squared norm inequality for averages, its distance from 1 is at most $\sqrt{8\,\mathrm{colorScale}^2\,m}$.

Take $\delta > 0$ so that the nearest-point comparison from the focus derivative holds with error coefficient 1. Put $\sigma = (\delta/2)^2/(8\,\mathrm{colorScale}^2)$ and $a = \delta/(2\,\mathrm{colorScale})$. If $m < \sigma$, the centroid is within $\delta/2$ of 1. At any index with $\mathrm{income}(x) < a$, its feedback is within δ of 1 as well. Apply that comparison to the nearest points h_x and 1 of $\mathrm{feedback}_x$ and centroid. The projection has norm at most one; therefore

$$\|\,\mathrm{mat}(h_x) - 1\| \le 2\|feedback_x - \mathrm{centroid}\,\|,$$
$$\mathrm{chord}(h_x) \le 2\,\mathrm{income}(x).$$

At the remaining indices $\mathrm{chord}(h_x) \le 2 \le (2/a)\,\mathrm{income}(x)$. Averaging proves the alignment estimate with $\max(2, 2/a)$. If $m \ge \sigma$, simply choose $b = 1$ and use average $\mathrm{chord}(h_x) \le 2 \le (2/\sigma)m$. Thus the single positive constant

$$C_{\mathrm{align}} = \max(2/\sigma, 2, 2/a)$$

gives $\mathrm{average}_x\ \mathrm{chord}(h_x b^{-1}) \le C_{\mathrm{align}}$ m for every finite pair family, with no smallness hypothesis on its individual transports.

Here is the route census used to bound m. For a radius-r box, replace selected coordinates of an endpoint x by the corresponding coordinates of y; call the resulting site $d_p(x, y)$. Swapping those coordinates between x and y is an involution of endpoint pairs. Therefore every site function f satisfies

$$\sum_{x,y} f(d_p(x, y)) = N \sum_{x} f(x).$$

Likewise swapping a site's j coordinate with an independent height in $-$r,. . . ,r is an involution of (site,height), giving

$$\sum_{x,\mathrm{height}} f(update_j(x, height)) = (2r + 1) \sum_{x} f(x).$$

A straight segment in direction j has chord at most the sum of the internal positive j-link chords on its whole coordinate line. For a positive segment this follows by the product triangle inequality and then enlargement to the full line; for a negative segment reverse the walk and use inversion invariance. Append one new axis j to an ordered route. Its initial part ends at $d_p(x, y)$, and its last segment is bounded by that site's full j-line sum. The two involution identities bound the sum of this added cost over endpoint pairs by $N(2r + 1)\,\mathrm{edgeSum}(U)$. Induction over the distinct axes proves the route bound $\mathrm{length}(p)N(2r + 1)\,\mathrm{edgeSum}(U)$; four axes give the constant 4. This establishes the route estimate and the finite-pair alignment used in synchronization. The preceding minimizer self-consistency argument gives the coordinatewise nearest-feedback condition used here.

6.17 Predictor query and synchronization estimates

Let Y_p be connector values with nonnegative patch weights summing to one. For arbitrary group elements x,y put $q = xy^{-1}$ and

$$\tau = \sum_{p} weight_p\,\mathrm{chord}(x^{-1}Y_p y).$$

Cyclic invariance of chord identifies this with $\sum$ weight$_p$ chord$(Y_p q^{-1})$. The weighted matrix mean $m = \sum$ weight$_p$ mat(Y_p) satisfies $\|m - \mathrm{mat}(q)\| \le$ colorScale $\cdot\tau$. In the best-predictor branch its value b is a nearest group point to m. Hence

$$\begin{aligned}\|\,\mathrm{mat}(b) - \mathrm{mat}(q)\| &\le \|\,\mathrm{mat}(b) - m\| + \|m - \mathrm{mat}(q)\| \\ &\le 2\,\mathrm{colorScale}\cdot\tau.\end{aligned}$$

In the other branch the minimum connector price is at least predictorCut. Since price$(g) =$ chord$(g)^2/2 \le$ chord(g) for $0 \le$ chord$(g) \le 2$, comparison with q gives $\tau \ge$ predictorCut. The universal chord bound 2 then gives the same common estimate in both branches:

$$\mathrm{chord}(x^{-1}\mathit{predict}(Y)y) \le 2\max(1, \mathrm{predictorCut}^{-1})\tau.$$

Thus the query estimate uses the actual fallback outcome as well as the nearest-point outcome.

We apply the preceding finite-pair alignment estimate to synchronization. The ordered route sum on a radius-r box obeys

$$\begin{aligned}&\sum_{p,q} \mathrm{chord}(\mathrm{travel}(p \to \mathit{qinordero})) \\ &\le 4N(2r+1)\,\mathrm{edgeSum}(U), N = \mathit{numberofpatchpoints}.\end{aligned}$$

It follows by charging each coordinate segment to its internal positive edges; the precise route-count argument is developed with the finite-pair input below. Averaging this inequality over orders and endpoints gives $N \cdot \mathit{meanPairChord} \le 4(2r+1)\,\mathrm{edgeSum}(U)$.

In the minimizing frame branch, gauge by H and set $b_p = \mathrm{frame}_p\, H_p^{-1}$. This array minimizes the corresponding finite pair energy: gauge covariance of route products preserves that energy, and any competing array can be normalized at the center by a common left multiplication without changing it. The finite-pair cluster inequality gives c with average$_p$ chord$(b_p c^{-1}) \le C$ meanPairChord. Therefore

$$\begin{aligned}&\sum_p \mathrm{chord}(c^{-1}\mathit{frame}_p H_p^{-1}) \\ &\le 4C(2r+1)\,\mathrm{edgeSum}(\mathrm{gauge}\,HU).\end{aligned}$$

The fallback branch is elementary and does not use that alignment inequality. Its positive frame mass makes its axial tag eligible, so its normalized route gauge d satisfies edgeSum(gauge$(d, U)) \le$ axialFactor$(2r+1)\,$faceSum(U). Put $g = dH^{-1}$. Edgewise factorization and the chord triangle inequality give

$$\mathrm{edgeSum}(\mathrm{gauge}\,g1) \le \mathrm{edgeSum}(\mathrm{gauge}\,dU) + \mathrm{edgeSum}(\mathrm{gauge}\,HU).$$

For this pure gauge, route transport from p to q is $g_p\, g_q^{-1}$. Choose q whose sum over p is at most its average over q, and apply the ordered route sum above. Taking $c = g_q$ yields

$$\begin{aligned}&\sum_p \mathrm{chord}(c^{-1}\mathit{frame}_p H_p^{-1}) \\ &\le 4(2r+1)[\mathrm{edgeSum}(\mathrm{gauge}\,HU) \\ &+\mathit{axialFactor}(2r+1)\mathit{faceSum}(U)].\end{aligned}$$

With $A = \max(4C, 4)$, both branches prove the synchronization estimate below. The finite-pair cluster inequality and route census proved in this chapter supply its two inputs.

The patch weight bound used in the macro predictor estimate also has a direct counting proof. Its weight is the distribution of a sum of J independent uniform small-box rolls. Separate the first roll from the other J−1 rolls. For every fixed choice of those remaining rolls, at most one first roll can give a specified sum. Since the small box has L^4 points (L is odd), that sum has probability at most L^{-4}. Thus L^4 $\text{weight}_p \le 1$, with no assumption that different roll histories give different sums.

6.18 Common macro gauge and predictor price

The planting argument uses the axial estimate already proved and a frame synchronization inequality. The latter states that for a usable frame on a patch of radius r and any gauge H there is a constant group element p with

$$\begin{aligned}\sum_{\text{patch}} \text{chord}(p^{-1} frame(q) H(q)^{-1}) \\ \le A(2r+1)[\text{edgeSum}(\text{gauge}\, HU) \\ + axialFactor(2r+1) faceSum(U)].\end{aligned}$$

The preceding synchronization proof establishes this inequality. Its constant A depends only on the group; the block-length factors are displayed explicitly.

Let j be the smoothing index, set $\ell_0 = 2(j+7)+1$, $D = 16\ell_0$ and $C = A\ell_0(D + axialFactor \cdot \ell_0)$. The macro radius is the patch radius plus 4L; consequently 2 macroR $+1 \le \ell_0$ L. Apply the four-axis axial construction on this whole macro box to obtain a single H with edge sum at most D L·port. Every patch translated by $centerRoot(y) = Ly$, $y \in [-2,2]^4$, lies in this macro box. Its internal edge and face sums are bounded by the full macro sums by nonnegativity and translation. Apply synchronization separately on each such translated patch using the same H. The displayed bounds give an element p_y with frame error sum at most CL^2 port. This proves the planting estimate used below and preserves one common macro gauge for all neighboring roots.

To price a predictor connector between adjacent roots y and $y + e_i$, factor it as the left frame error, the straight length-L transport in gauge H, and the inverse right frame error. Its chord is bounded by the sum of these three chords. Let the two frame-error sums be at most KL^2 port and the macro edge sum at most TL port. The sum of the straight-path chord over all patch starting points is at most L times the macro edge sum: each step's starting-point translation is injective, its links remain in the macro box, and the nonnegative edge census bounds that partial sum. This bound counts all L steps even if they repeat a torus link.

The patch weight bound is L^4 $\text{weight}_p \le 1$. Therefore the weighted connector chord sum τ satisfies

$$L^4 \tau \le L^2(2K + T)\,\text{port}\,.$$

The predictor query inequality bounds its chord, in this left and right frame, by 2 $\max(1, \text{predictorCut}^{-1})\tau$. Multiplying by L^2 and using $L > 0$ proves

$$\begin{aligned}L^2\,\text{chord}(p_y^{-1} Q_{(y,i)} p_{(y+e_i)}) \\ \le 2\max(1, \text{predictorCut}^{-1})(2K+T)\,\text{port}\,.\end{aligned}$$

This proves the indicated-predictor estimate from synchronization and the predictor query inequality.

It also supplies the coarse curvature comparison. Four such gauged predictors around an elementary coarse plaquette have chord at most α port $/L^2$ each, where $\alpha = 2\max(1, \text{predictorCut}^{-1})(2C+$

D). Their endpoint p factors cancel in the plaquette product. Triangle and conjugation invariance yield $\mathrm{chord}(\mathrm{face}\, Q) \le 4\alpha\,\mathrm{port}\,/L^2$. For any V, compare its four links to Q. The two-factor inequality

$$\mathrm{chord}(ab(cd)^{-1}) \le \mathrm{chord}(ac^{-1}) + \mathrm{chord}(bd^{-1})$$

follows by factoring the product into one discrepancy and a conjugate of the other. Iteration, with inversion invariance for the reversed links, gives $\mathrm{chord}(\mathrm{face}\, V) \le \mathrm{chord}(\mathrm{face}\, Q) + \mathrm{fuel}(Q, V)$, where fuel is the sum of its four wire discrepancies. This is the local curvature comparison used by the clean-patch and shell arguments.

6.19 The nonlinear shell link comparison

Fix the forest and its actual slice. To compare a fine link (x,i), close it with the reverse tree hike from its first coarse root to x and the tree hike from $x + e_i$ to its second coarse root. On the slice both hike transports are 1, so the resulting walk has transport $U(x, i)$, endpoints the two embedded coarse parents, and length at most 8L+1. Lift the whole walk to the cover starting at 0 over its first root. Every point lies in the macro box.

The planting estimate supplies on this same macro box a gauge H and comparison elements p_y with

$$\begin{aligned} &\mathit{totalEdgeChord}(\mathrm{gauge}\, HU) \le DL \cdot \mathrm{port}, \\ &\sum_{\mathrm{patch}} \mathrm{chord}(p_y^{-1} \mathit{frame}_y(q) H(q + \mathit{centerRoot}(y))^{-1}) \\ &\qquad \le CL^2 \cdot \mathrm{port}\,. \end{aligned}$$

The second inequality is read at the distinguished patch point, whose frame is 1, to give $\mathrm{chord}(H(\mathit{centerRoot}(y))p_y) \le CL^2$ port. The indicated-predictor estimate on adjacent roots gives

$$\begin{aligned} \mathrm{chord}(p_0^{-1} Q p_{\mathrm{ei}}) &\le \alpha L^{-2}\,\mathrm{port}, \\ \alpha &= 2\max(1, \mathrm{predictorCut}^{-1})(2C + D). \end{aligned}$$

These planting and indicated-predictor estimates are inputs to this step.

The triangle inequality along the lifted forest walk bounds its gauged transport by

$$u \cdot \mathrm{port}, u = (8L + 1)DL.$$

The period assumption 100 macroR $\le \mathit{fineSide}$ makes the cover wrapping injective on the macro box. If the two coarse parents are equal, the lifted endpoint is therefore 0; conjugation invariance bounds $\mathrm{chord}(U(x, i))$ by u port and its replica comparison link is 1. If they differ, the centered parent-hop identity makes the second parent the positive coarse neighbor, and injectivity identifies the lifted endpoint with $\mathit{centerRoot}(e_i)$, even across a periodic seam. Factoring the gauged predictor through p_0 and $p_e i$ gives its chord bound

$$v \cdot \mathrm{port}, v = 2CL^2 + \alpha/L^2.$$

Compare the fine link and predictor in that same macro gauge and use the triangle inequality to get $\mathrm{chord}(U(x, i)Q^{-1}) \le (u + v)\,\mathrm{port}$. Comparing Q with the desired coarse link V then adds precisely its wire discrepancy, with coefficient one. This proves the shell comparison with constant u+v uniformly in the period and in the given forest. The room hypotheses supply its size assumptions, so the same constant is available in the local driving inequality.

6.20 The local driving inequality

Let lens(u) be the active part of the fine tangent field. At a coarse site o define current(o,u) as the square root of the sum of $\|\operatorname{lens}(u)_e\|^2$ over the radius-3B fine-link neighborhood, and response(o,u) as the square root of the sum of squared selected motion rows over radius macroR+2. A summand's norm is bounded by the corresponding square root because all squared norms are nonnegative.

Fix a hot fine link e and put $o = \operatorname{parent}(e)$. Suppose a gauge z makes the background links in radius 2B about o have chords at most h. Every row root within macroR+2 of o is selected in the identity orientation: hotness puts e far enough from the seed set, and the room radius leaves this entire neighborhood clear. Its oriented motion norm is bounded by response(o,u). Apply the transported-to-flat row comparison to the gauged background and the gauged lens. The required tangent values lie in the radius-3B neighborhood of o, so their norms are bounded by current(o,u). Gauge covariance and norm preservation give, at each such row,

$$\|barField_R(Ad(z)\operatorname{lens}(u))\| \le E_o,$$
$$E_o = response(o,u) + K_{\mathrm{st}} \cdot 2\,\mathrm{colorScale}\cdot(B+1)h \cdot current(o,u).$$

The local link estimate for a tree-slice tangent field w with zero trunk values states

$$\|w_e\| \le K \sum_{qinmacro\,\mathrm{box}} \sum_{plaquettesp} \|\operatorname{curl}(w)_{(q,p)}\|$$
$$+\|predictorLinearBill(w)_{(o,axis(e))}\|.$$

Here K is the common local comparison constant for the actual predictor tags; the gauge-flattening and local-driving arguments in this chapter supply it. Apply this estimate to $w = Ad(z)\operatorname{lens}(u)$. It vanishes on trunks, since these are not hot. The direct-curvature rows in the preceding bound imply $\|\operatorname{curl}(w)\| \le E_o/\sqrt{\mathrm{direct}}$, and the noise row implies $\|predictorLinearBill(w)\| \le E_o/\sqrt{\mathrm{penalty}}$. All macro-box plaquette roots belong to the radius macroR+2 neighborhood. Thus

$$\|\operatorname{lens}(u)_e\| \le ZE_o,$$
$$Z = K|\operatorname{box}(\mathrm{macroR})||PlaqAxes|/\sqrt{\mathrm{direct}} + 1/\sqrt{\mathrm{penalty}}.$$

This proves the local driving estimate used in the plate floor. A clean-patch gauge provides h of rate $-1/2$ with its radius factor already included in the flattening bound; hence the coefficient of current tends to zero fast enough for the squared-neighborhood absorption previously described. The tree-slice link estimate in this chapter gives the constant uniformly in the tags.

6.21 Transported affine rows

For coarse sites x,y choose the coordinate-order route from 0 to the centered lift of y−x and start it at x. Its endpoint is y and its length is at most 4 separation(x,y), since each of its four coordinate lengths is bounded by that separation. Let $c_x(V;y)$ be its actual parallel transport in V. For a fine tangent field u, spin each link value at its coarse parent y by $\mathrm{Ad}(c_x(V;y))$, then apply the cutoff flat linear row rooted at x. This defines the straightened row. The signed oriented row first changes the input coordinates by the actual tangent transformation associated with the replica of V, then applies this straightened row in the signed background.

These constructions are linear in u. The tangent change at a reversed link includes its background adjoint factor and a minus sign; it is not a mere permutation of color coordinates. Both this

tangent change and the carrier adjoint preserve the norm of each color vector. Thus the straightened and oriented rows have the same $K_{\mathrm{st}}\|u\|$ bound as the cutoff flat row. At $V = 1$, every carrier is 1 and the straightened row is exactly the cutoff flat row.

Gauge covariance is exact. The carrier transforms as $z(x)c_x(V; y)z(y)^{-1}$, because its endpoint is y. Combining its adjoint with the adjoint on u cancels $(z(y))^{-1}z(y)$, leaving the common factor $Ad(z(x))$ at the row root. The flat row commutes with this common color action. Signed tangent transformation commutes with the gauge change after pulling the gauge to the signed sites; hence the oriented row transforms by the adjoint at its signed root as well.

The locality and comparison bounds can be proved without controlling the background outside the queried hull. A carrier to y depends only on the radius-d hull about x if 4 separation$(x, y) \le d$. If link chords there are at most $h \ge 0$, replace all links outside the hull by 1. The carrier is unchanged and its chord is at most length $\cdot h \le dh$ by the walk estimate. Similarly, if two tangent inputs differ by at most M on the radius R+predictionReach+4 input stencil, truncate their difference to zero outside this stencil. The resulting global field has norm at most M, and linearity and locality give a row difference at most K_{st} M.

It follows that, when $4(R + \text{predictionReach} + 4) \le d$ and the input is bounded by M on the needed stencil,

$$\begin{aligned}&\|straightened_R(V, u) - barField_R(u)\|\\ &\le K_{\mathrm{st}} \cdot 2\,\text{colorScale} \cdot dhM.\end{aligned}$$

Indeed each carrier changes the input by at most 2 colorScale·dh M; apply the preceding row comparison. The signed case has one additional error: at a reversed stored link the tangent transformation at V differs from that at 1 by at most 2 colorScale h M. A signed link root moves by at most one site, so radius d+1 bounds in the original orientation supply the required radius-d bounds after turning. Adding the carrier and reversal errors gives

$$\begin{aligned}&\|orientedRow_R(V, u) - orientedRow_R(1, u)\|\\ &\le K_{\mathrm{st}} \cdot 2\,\text{colorScale} \cdot (d + 1)hM.\end{aligned}$$

All fields in this inequality refer to the same signed row and input; no independent branch outcome is selected.

For a single-link input of color a, the signed tangent transformation and the carrier spin still produce a single-link input with color norm $\|a\|$. Its parent distance from the row root differs from the original signed distance by at most one. The flat impulse estimate above therefore yields

$$\begin{aligned}&\exp(width_{\mathrm{top}} \cdot distAt(link, signedRoot))\\ &\|orientedRow_R(V, single_{\mathrm{link}}a)\|\\ &\le 2K_{\mathrm{st}} \exp(width_{\mathrm{top}}(\text{predictionReach} + 7))\|a\|.\end{aligned}$$

This bound requires no smallness of V. Locality at $R = \text{till}(x)$ follows from the same carrier reach and the room inequality, using a half-sized hull before the signed change to absorb its one-site root shift.

6.22 Affine accuracy on the protected stencil

The local hypothesis here records both actual field agreement and a tangent bound. Say that U is grafted to V by u on a radius-d hull if it agrees there with $\exp(u)$ replica(V) and $\|u_i\| \le H$ on the same hull. This condition is preserved by a coarse gauge transformation with u changed by its adjoint action. Under a signed coordinate change it is preserved on radius d-1, because the stored

positive branch of a reversed link may shift the root by one site. Restriction to a smaller hull and enlargement of H also preserve it.

Fix a row p, its root i, $d = B/4$ and $R = \mathrm{till}(x)$. Suppose V agrees with exp b on that hull, $\|b\| \le H$, and U is grafted by u there. Replace u outside this hull by zero to obtain v with $\|v\| \le H$, and set $a = linearCopy(b)$, so $\|a\| \le H$. The logarithm-product estimate supplies w with

$$\exp w = \exp v \exp a,$$
$$\|w - v - a\| \le JH^2, \|w\| \le (1 + J)H$$

for eventually small H. The equality is linkwise and uses the exact replica exponential identity. Locality of the entry row now identifies its actual value C with entry(exp w,exp b), and its background value Y with entry(exp a,exp b).

Let E_0 bound the error between these two logarithmic entries and their corresponding full linear costs; let M_0 bound the full linear cost for w. These are the uniform entry-near estimates at the enlarged envelope (1+J)H. Then $\|C\| \le M_0 + E_0$. To estimate the affine remainder write $z = \mathrm{linearCost}(w, b)$, $z_0 = \mathrm{linearCost}(a, b)$. Linearity gives the exact decomposition

$$\begin{aligned} &C - Y - straightened(V, u) \\ &= (C - z) + uncut(w - v - a) + (z_0 - Y) \\ &\quad +(uncut(v) - truncated(v)) \\ &\quad +(truncated(v) - straightened(V, u)). \end{aligned}$$

The five terms have bounds E_0, q_1, E_0, q_2, q_3 respectively, where, for the common stencil constant K and exponential-chart constant A,

$$q_1 = KJH^2,$$
$$q_2 = KH \exp(-width \cdot R),$$
$$q_3 = K \cdot 2\,\mathrm{colorScale} \cdot B \cdot AH^2.$$

The first uses the full stencil bound on w−v−a. The second uses its exponential tail bound. For the third, V's links in the hull have chords $\le AH$; the adjoint transport comparison bounds the background-to-flat row difference by the displayed radius-times-quadratic error. Locality first replaces u by v, so the bound never uses u outside the hull. The room inequality $4(R + \mathrm{predictionReach} + 4) \le d$ ensures that every needed input lies inside it.

Consequently the common affine error can be chosen as

$$E_{\mathrm{aff}} = 2E_0 + q_1 + q_2 + q_3.$$

It also bounds $\|Y - \mathrm{linearCost}(a, b)\|$. If H has rate $-1/2$, the entry-near estimate gives E_0 rate -1 and M_0 rate $-1/2$; the quadratic, cutoff, and radius-weighted quadratic terms all have rate -1. Thus $M_0 + E_0$ has rate $-1/2$ and E_aff has rate -1. The predictor entry-near estimate and the uniform full and tail stencil bounds of this chapter supply the constants used here.

6.23 The protected mean and its stationary comparison

The protected mean is the stationary comparison variable retained when vertical fluctuations are integrated. Its affine accuracy is tested on the prescribed stencil. The proof first estimates the discrepancy of the actual nonlinear rows from their flat counterparts, then inserts those bounds in the stationary identity. The cancellation concerns these actual rows; it does not assert that every nonlinear minimizer is equal to the flat one.

Here the scale, geometry, selected rows, and gauge forest are held fixed throughout. Write A for the counting-coordinate plate, D for the selected linear form, b for its linear bill, and S for the diagonal dormant-coordinate operator. For every fine tangent field v the pricing identity gives

$$Av + b = D^{\mathsf{T}} priced(v) + Sv.$$

Coordinates are expressed in the fixed orthonormal color basis. All norms of fields in this argument are supremum norms; a color-vector norm is bounded by the sum of the absolute values of its color coordinates. Let Y bound both the row sums of D and the row sums of D^{T}. Its matrix entries vanish beyond distance B/4.

We first produce a comparison field in the same gauge as the actual background. A clean background on a ball of radius d+1, with $d \le 32B$, admits an axial gauge in which each link in the radius-d ball has chord at most

$$u(x) = 16(2(32B(x)+1)+1)^5 |PlaqAxes| \cdot 2\,\mathrm{fiery}(G, x, L).$$

The axial-patch estimate supplies this bound using the clean plaquette toll and the room temperature comparison. The local logarithm extension then gives a globally defined bounded logarithmic field a agreeing with that gauged background on the smaller ball, with $\|a\| \le C\ u(x)$. Thus the bound $w = C$ u has rate $-1/2$. The axial-gauge and bounded logarithmic-extension argument of this chapter proves this assertion.

For clarity, the logarithm-product estimate used by the subsequent comparisons is a genuine group-chart estimate. Set $h(p,q) = \log(\exp p \exp q)$. Differentiating at (0,0), using the derivative of the matrix exponential and the inverse-chart identity, gives $Dh(0)(p,q) = p+q$. Analytic Taylor expansion on a sufficiently small ball therefore gives, with fixed positive K,r,

$$\begin{aligned} \exp h(p,q) &= \exp p \exp q, \\ \|h(p,q) - p - q\| &\le K \max(\|p\|, \|q\|)^2, \\ \|h(p,q)\| &\le K \max(\|p\|, \|q\|). \end{aligned}$$

Applying this link by link proves the same uniform bounds for the logarithm of two successive increments. No identification of a group-valued density with a tangent vector is made here.

Let spring(a) be the stationary response with zero values on the chosen forest trunks. It is linear in a. The stationary response estimate bounds each color coordinate for $\|a\| \le 1$ by

$$C_0 = Response.total(j, \eta, C) + \exp \eta.$$

Consequently its field norm is at most $C_s\|a\|$, where

$$C_s = \max(1, \max(0, C_0)\dim(Color)).$$

Indeed this is the coordinate bound summed over colors for unit input; rescaling a by $\|a\|^{-1}$ proves the general case, and linearity handles $a = 0$. Signed coordinate changes preserve the field bound and commute with the linear copy. The latter assertion follows either from the copy formula or by differentiating the exact identity for its exponentiated replicas.

If z is the axial gauge transformation, put $v = Ad(z^{-1})\,\mathrm{spring}(a)$. Then v vanishes on trunks and $\|v\| \le C_s$ w. Gauge covariance of the selected path, away from the seed region, identifies its value at v with

$$request(\mathrm{gauge}(z, V)) + orientedRow(\mathrm{gauge}(z, V), \mathrm{spring}(a)).$$

The source uses one common clean hull for all rows and columns subsequently read. On this hull the uncut stationary identity is exactly

$$\begin{aligned}\mathrm{linearCost}(\mathit{linearCopy}(a), a) - \mathrm{linearGuide}(a)\\ +\mathit{uncutField}(\mathrm{spring}(a)) = 0.\end{aligned}$$

To see this, the stationary shifted field is the sum of the linear copy and spring(a). Linearity of the uncut field and the defining stationary equation identify its image with linearGuide(a). This proves the displayed cancellation before any estimate is applied.

Subtract this identity from the gauged real path. There are four errors: the base-stroke linear approximation, the guide linear approximation, replacement of the background-oriented row by the flat truncated row, and replacement of the truncated row by the uncut row. Denote the first two envelopes by A_1 and A_2; the guide and local affine comparison estimates give each rate -1. If C_{ch} bounds chord(exp a) by $C_{\mathrm{ch}}\|a\|$ and K_{st} bounds the stationary stencils, the last two errors are bounded respectively by

$$\begin{aligned}U_1 x = K_{\mathrm{st}} \cdot 2\,\mathrm{colorScale} \cdot (B(x)+1) \cdot C_{\mathrm{ch}} w(x) \cdot C_s w(x),\\ U_2 x = K_{\mathrm{st}} C_s w(x) \exp(-\mathit{width} \cdot \mathrm{till}(x)).\end{aligned}$$

The first has rate -1 because it is quadratic in w with only the radius factor; the second has rate -1 by the cutoff decay. Hence

$$A_0 = A_1 + A_2 + U_1 + U_2$$

has rate -1 and bounds every path of v whose root lies outside the 20B seed neighborhood and whose radius-2B background hull agrees with exp a. These bounds use the previously specified guide, affine-row, stencil, and stationary identities on that same hull.

Now let P be the nonsymmetric inverse-column pad, $E = AP - I$, and $\mu = -P^{\mathsf{T}} b$. Symmetry of A gives $E^{\mathsf{T}} = P^{\mathsf{T}} A - I$, whence the exact identity

$$\mu - v = E^{\mathsf{T}} v - P^{\mathsf{T}}(Av + b).$$

For a coordinate i, only the column window of P at i is read by the second term. If $|(Av+b)_j| \le M$ on that window and $\|v\| \le W$, the inverse-window estimates already proved imply

$$|\mu_i - v_i| \le \mathit{payE} \cdot M + \mathrm{scraps}(B) \cdot W.$$

This uses only the elementary inequality $|\sum_j M_i j\ t_j| \le K \sum_j |M_i j|$ when $|t_j| \le K$ on the nonzero entries of the row. It requires no bound on t outside that row support.

For coordinates outside the 11B seed neighborhood, S $v = 0$: on a trunk $v = 0$, and any other such coordinate is active. The pricing identity and the transpose row bound therefore show that path bounds A_0 within distance B imply

$$|(Av + b)_i| \le Y \cdot \mathit{scale} \cdot A_0.$$

For a selected path p, locality and the row bound of D also imply

$$\begin{aligned}\|\mathit{realPath}_p(u) - \mathit{realPath}_p(v)\|\\ \le \dim(\mathit{Color})\mathit{scale}^{-1} Y U\end{aligned}$$

whenever $|u_i - v_i| \le U$ on the radius-B coordinate neighborhood of its root.

Take a path root outside the 30B seed neighborhood. Its clean hull has radius 8B+1. Flatten on radius 8B and use the resulting single comparison field v for every coordinate needed below. Its paths within distance 6B satisfy the A_0 bound. Each column window read at distance B from the root lies within distance 3B, so the residual bound above applies throughout that window. All these rows remain outside the required 11B seed neighborhood. Thus, with

$$W = C_s w, M = YscaleA_0,$$
$$U = payE \cdot M + \mathrm{scraps}(B) \cdot W,$$
$$E_{\mathrm{mean}} = A_0 + \dim(Color)scale^{-1}YU,$$

we obtain $\| \mathrm{realPath}_p(\mu)\| \leq E_mean$. All envelopes are eventually nonnegative. The scraps bound at any prescribed negative power gives rate -1 for U and hence for E_mean. The root is selected by the same seed-distance condition. This proves the protected-mean estimate used in the watch argument, with one common choice of gauge and comparison field for its entire finite stencil.

The axial patch estimate follows by lifting the torus patch to its radius-r cover box. The axial gauge construction bounds the total edge chord by 16(2r+1) times the face-chord sum. If every face chord is at most q, the sum is at most $(2r+1)^4$ |PlaqAxes| q. Descend the gauge to the torus using injectivity on the box when $2r < n$. Both endpoints of every radius-d bond lie in that box when $d + 1 \leq r$, and each individual edge chord is bounded by the total nonnegative sum. Thus the bound is $16(2r+1)^5$ |PlaqAxes| q, as used above.

6.24 Local normalization differences

The condensation letter at a fine link e is

$$condensation(e) = [joint(A_{\mathrm{blank}}, window; e) - joint(A, window; e)]/2.$$

The blank geometry has the same forest, scale, windows and tags, but background 1 and empty seeds. Its hot links are exactly the nontrunks, and all its rows are selected. No second forest or tag choice is made.

We first compare joint diagnostics on the same finite window U. For positive plates A,Z let $f_A = A_U - I$ and $g_A(s) = I + s(A_U - I)$, with $0 \leq s \leq 1$, and similarly for Z. The window intensity is $f_A\ g_A(s)^{-1}$, evaluated at its distinguished diagonal coordinate. Suppose all entries of A−Z on U have absolute value at most $e_0 \geq 0$. The census bounds $|U| \leq \Delta = \mathrm{ballCost}(B+2)m$, so the row sum of $f_A - f_Z$ and of $g_A - g_Z$ is at most $d = \Delta e_0$. Let $C_{\mathrm{inv}} = payE$ bound both inverse row sums uniformly along this interpolation, and let Q+1 bound the row sum of f_Z.

The resolvent identity gives

$$g_Z^{-1} - g_A^{-1} = g_Z^{-1}(g_A - g_Z)g_A^{-1},$$

so its row sum is at most $C_{\mathrm{inv}}^2 d$. Decomposing the difference of the two intensities into $(f_A - f_Z)g_A^{-1} - f_Z(g_Z^{-1} - g_A^{-1})$ yields the diagonal bound

$$pulseHead(e_0) = [C_{\mathrm{inv}} + (Q+1)C_{\mathrm{inv}}^2] \cdot \Delta e_0.$$

All inverse matrices here are the principal-window inverses at that same s. Integrating from 0 to 1 and summing over the color coordinates of e gives

$$|joint(A; e) - joint(Z; e)| \leq \dim(Color)pulseHead(e_0).$$

Continuity in s, supplied by positivity of the interpolated principal matrices, justifies subtraction of these integrals. The same calculation without taking a difference bounds each individual joint diagnostic by $\dim(Color)(Q+1)C_{\mathrm{inv}}$.

To apply the local difference bound, take e outside the 30B seed neighborhood and flatten its background in a single radius-6B patch. If the resulting link chords are at most h_0, compare each actual gauged form entry to the blank one by testing on a single coordinate impulse. Trunk columns are zero in both forms. A nontrunk column in the needed neighborhood is hot in both geometries. Rows farther than B/4 from that column vanish in both; all other rows are selected by the seed margin. The transported-row comparison already proved therefore gives

$$|D_{\mathrm{gauged}}(p,i) - D_{\mathrm{blank}}(p,i)| \le swell,$$
$$swell = scale \cdot K_{\mathrm{st}} \cdot 2\,\mathrm{colorScale} \cdot (B+1)h_0.$$

The dormant diagonal terms of their plates agree in this neighborhood. Expand $D^{\mathsf{T}}D$ entrywise using $ab - cd = (a-c)b + (b-d)c$. The two transpose row sums are at most Y, so

$$|A_{\mathrm{gauged}}(i,j) - A_{\mathrm{blank}}(i,j)| \le 2swellY =: repair$$

throughout every window read at e. This is an entrywise comparison only on the patch, with the census factor retained in pulseHead.

Gauge invariance of joint follows by conjugating the plate by the block diagonal orthogonal color action, which preserves every link window. The same conjugation preserves its principal inverse, interpolation intensity and color trace at e, hence its integrated diagnostic. Thus the gauged plate can replace the actual plate in the condensation difference.

Flattening gives $h_0 = Z_{\mathrm{ch}}\ w(x)$ with w of rate $-1/2$. Hence repair has rate $-1/2$, including its radius factor. The window census Δ has rate zero. The envelope

$$E_{\mathrm{cond}}(x) = \dim(Color)pulseHead(repair(x))$$

therefore has rate $-1/2$ and bounds the absolute condensation outside the 30B seed neighborhood; dropping the factor $1/2$ only enlarges the bound. This derives the remote normalization estimate used below from the local form comparison and the already established inverse-window bounds.

6.25 Exact locality of the cut gaps

Every walk of length at most R reads only the radius-R hull of its starting point, in any projection for which neighboring sites have projected separation at most one. This includes stored negative links: their positive-branch root is one step behind the current point, still paid by the step being read. Induction on the walk proves equality of its transport for fields agreeing on that hull.

A retained plotted slot has its transport length at most R, and its final plaquette needs four more steps. Thus plotted fields agree under radius R+4 agreement. The cutoff F,I,N kernel rows are banded at radius R and their entries read only the radius-R coarse hull about the row root. For the fine F row this follows by lifting its designated route in the replica; its endpoint projects to the coarse column, and the route-length cutoff bounds that projected separation. The lane rows use the same argument on their coarse route.

For a banded row $F(V)$, its transpose at output j reads input rows within distance R. Agreement of V,V′ on radius 2R about j makes each such row entry equal; input values need agree only at those rows. Likewise $F(V)^{\mathsf{T}}FV$w is unchanged by this background replacement on that hull, with

the entire arbitrary array w held fixed. The inner row coefficients agree individually, so no locality assumption on w is needed.

Apply this to lGap. The fine plotted inputs at relevant rows require at most R further travel plus their four plaquette steps; the noise input additionally queries its actual predictor through predictionReach. The coarse plotted term needs R+4. Therefore common fine/coarse agreement on

$$D \geq 3R + 4 + \text{predictionReach}$$

at the tested anchor makes the whole lGap equal, with the same actual predictor tags on both sides. For dGap the three banded row-square operators are unchanged on radius 2R; hence the same D suffices for every fixed input array w. These prove the local gap identities used in the chart transfer, without changing any clipping indicators in an unrelated data array.

6.26 A reliable cutoff logarithm

Choose a smooth profile χ with values in [0,1], equal to one on the inner unit interval and zero beyond its prescribed outer interval, and set

$$cutLog_d(U) = \chi.innerShare(d, \text{chord}(U)) \cdot \log(\text{mat}(U)).$$

Here innerShare evaluates χ at $(\text{chord}(U)/d)^2$, so it uses a smooth squared matrix norm. Its nonzero values imply $\text{chord}(U) < 2d$. The ordinary logarithm is only required to have its chart properties on this latter shell.

The local logarithm estimates give $C, s > 0$ with $\exp(\log U) = U$ and $\| \log U \| \leq C \,\text{chord}(U)$ when $\text{chord}(U) \leq s$. Choose also a matrix radius t where log is analytic and a color radius R where $\log(\exp a) = a$. Set

$$r = \tfrac{1}{2} \min(s, t/\,\text{colorScale}, R/C), \gamma = C.$$

For every $0 < d \leq r$, the entire chord shell $\leq 2d$ satisfies all three chart conditions. In particular the logarithms there have norms $\leq R$. Applying $\log(\exp a) = a$ to $a = Ad(v)\log$ U and $a = -\log$ U, whose norms are unchanged, and using exponential conjugation and inversion, proves $\log(vUv^{-1}) = Ad(v)\log$ U and $\log(U^{-1}) = -\log$ U on the shell.

On $\text{chord}(U) \leq d$ the profile equals one, so the cutoff logarithm is the actual inverse chart. At every U its norm is $\leq \gamma \cdot 2d$: if the profile vanishes this is immediate, and otherwise use its value ≤ 1 and the shell logarithm bound. Chord conjugation and inversion invariance extend the two equivariance identities to all U, splitting into the zero and nonzero profile cases.

For any smoothly varying matrix-valued group input, the cutoff logarithm is smooth. Inside the closed outer shell this is composition with the analytic log and the smooth squared-norm profile. Outside it, continuity of chord gives a neighborhood where the profile vanishes identically. The same argument proves continuity, hence Borel measurability, as a function on G. Thus no regularity of the chosen logarithm outside its chart is used.

6.27 The sector input and its grafted signal

Fix a forest F, coarse background V, tags k and frozen slice values Y. A fine link is hot when it is not a trunk and neither of its coarse endpoints touches the 10B seed dome. Define input(u) by putting $\exp(u_e)(\text{replica}(V))_e$ on hot links, the fixed Y values on other nontrunks, and the prescribed tree value on trunks. This is on the same tree slice for every u.

Let small(e, U) be the cut logarithm of U_e (replica$(V))_e^{-1}$ at radius width·t. Define saved to be this logarithm on frozen nontrunks and zero elsewhere; lens(u) is u on hot links and zero elsewhere. The signal is saved+lens(u). Thus signal equals u on hot links and zero on trunks.

Shell access gives two facts, uniformly over every forest and actual input: a reliable logarithm on the chord shell of radius width·t, and the shell displacement bound chord$(U_e \operatorname{replica}(V)_e^{-1}) \leq width \cdot t$ whenever the local Interior condition holds. It is available by choosing width equal to the previously proved shellPrice and then taking x large enough that width·fiery lies in the reliable logarithm radius. Its constants are chosen once after L.

On a link satisfying that Interior condition, $input(u) = \exp(\text{signal}(u)) \operatorname{replica}(V)$. For a hot link this is its definition. On a trunk both the signal and the replica trunk increment are zero or identity as appropriate. On a frozen nontrunk the shell bound makes the cutoff logarithm uncut and its exponential recovers the actual discrepancy, giving the same equality. The reliable cut logarithm also supplies the all-argument bound

$$\|saved_e\| \leq \gamma \cdot 2(width \cdot t).$$

Hence local Interior conditions for every link of a radius-B hull, tangent bounds $\|u_e\| \leq M$ at its hot links, and $\gamma \cdot 2(width \cdot t) \leq M$ imply the grafted condition on that entire hull with $signal(u)$. Frozen links use the saved bound; hot links use the given tangent bound.

There is a simpler remote version requiring no Interior premise. If the hull center is outside the (11B+1) seed neighborhood, every nontrunk in its radius-B hull is hot: otherwise one endpoint would touch the 10B dome, and its parent lies within one further site, contradicting the separation. Only hot links and trunks occur in that hull. The defining identities above and hot tangent bound therefore prove grafting directly for any $M \geq 0$. These are the required and distant grafted constructions used in the packet annulus and the optional remote letters, respectively.

6.28 A common chart on the grafted hull

Suppose the coarse background V is clean on radius d+1 and U is grafted to V by a tangent field bounded by $H(x)$ on radius d, with H of rate $-1/2$. Flatten V there to exp b in gauge g, with $\|b\| \leq v(x)$ of rate $-1/2$. If D_{ch} is the exponential chord constant, the gauged coarse links have chords $\leq D_{\text{ch}}$ v. Gauge covariance of grafting writes the gauged fine links as $\exp(Ad(g)u)$ times the replica of the gauged coarse field. Their chords are therefore at most $D_{\text{ch}}\ H + D_{\text{ch}}$ v by the product triangle inequality; replica links are either 1 or a coarse link. The single envelope

$$W = H + v + D_{\text{ch}}H + D_{\text{ch}}v$$

bounds b, both kinds of link chord, and H itself, and has rate $-1/2$. All chart radii are met eventually. This is the common coarse/fine local chart; it imposes no condition outside the hull.

Apply affine accuracy in this gauge after restricting the hull from B/2 to B/4. The enlarged envelope W supplies both its background-log and tangent bounds. Entry covariance, replica covariance and covariance of the straightened row rotate the entire affine difference by the same root adjoint. Norm preservation then transports the estimate back to the original fields. This proves the strokes bounds M of rate $-1/2$ and E of rate -1 used in the soft remainder argument.

The same chart proves that the actual local branch tags and holonomies lie in any prescribed small neighborhood eventually. Given $\delta > 0$, the finite-stencil branch estimate supplies a radius r and reach k such that fine links of chord $\leq r$ imply site mass 1, pin mass 1, and predictor chord $\leq \delta/2$. Choose x so that $k \leq B/2$ and $W \leq \min(r, \delta/4)$. Apply this estimate to the gauged fine field. Site and pin masses are gauge invariant. The actual wire discrepancy is also invariant and

is at most the sum of coarse-link chord W and predictor chord $\delta/2$, hence at most δ. For a fine plaquette based at the chosen parent, replace links outside the radius-B/2 hull by 1. Its four-step product is unchanged, and its chord is $\le 4W \le \delta$. Thus the original fine plaquette toll and the actual wire both meet the small-curvature threshold. The uniform branch-cutoff and predictor estimates in this chapter give this finite-stencil estimate.

6.29 Uniform local branch activation

At the identity cover field the minimizing frame is 1 for every frame tag, the connector array is 1, and the minimizing predictor is 1. Both minimum tests are strict there. Continuity of the selected minimizers on their open minimum branches therefore gives a neighborhood of the identity cover field on which the site-frame minimum test holds and every connector predictor lies within any prescribed positive chord tolerance δ. This neighborhood can be chosen simultaneously for all axes and both endpoint frame tags because these alphabets are finite. On it every predictor fallback tag yields the same best value.

A product-topology neighborhood restricts only finitely many cover links. Choose a finite set I of such links and a common positive chord radius r small enough to meet all their coordinate neighborhoods; the conversion between chord and matrix norm is the fixed factor colorScale. Set d to the largest length of the coordinate-order route from 0 to a root in I. For any torus input, expand it at the chosen embedded coarse root. Every required cover link projects to a fine link whose coarse-parent distance from that root is at most its route length, hence at most d. Thus chord bounds $\le r$ on this parent-distance hull place the expanded field in the chosen cover neighborhood, independently of the period.

The actual connector equals the cover test connector with its actual two endpoint tags. Consequently its predictor chord is $\le \delta$. On the minimum branches the normalized site mass and pin mass equal one: their unnormalized masses are respectively the reciprocals of the numbers of frame tags and connector tags. This proves the finite-stencil branch activation used in the local chart-mark argument without requiring any tags or defaults away from the finite hull.

6.30 Global mismatch estimates from the flat identities

Let F be a cut transported row and F_0 its full flat comparison, with both row and column sums $\le C$ and difference sums $\le Cq_0$. Suppose an input f has linear comparison t with $\|t_j\| \le Ch$ and $\|f_j - t_j\| \le \mathrm{Chq}_0$. The exact decomposition

$$F^{\mathsf{T}} f - F_0^{\mathsf{T}} t = F^{\mathsf{T}}(f - t) + (F - F_0)^{\mathsf{T}} t$$

and the column-sum estimates give norm $\le 2C^2 h\ q_0$ at every output coordinate. For arbitrary bounded input w, applying this argument to $f = Fw$ and $t = F_0 w$ gives

$$\|F^{\mathsf{T}} F w - F_0^{\mathsf{T}} F_0 w\| \le 2C^2 \|w\| \infty q_0.$$

Choose C to dominate the fine plotted, coarse plotted, predictor-noise and three kernel constants previously proved. Choose the chart radius as their common minimum, and choose

$$\tau = \min(\tau_{\mathrm{kernel}}, width_{\mathrm{top}}, width_{\mathrm{below}}/L) > 0,$$
$$q_0 = \mathrm{widen}(R, h, \tau).$$

The plotted-field and predictor Taylor comparisons then all have the common input bounds above; the noise quadratic remainder uses $h \le q_0$. The F,I,N kernel comparisons likewise have the

common row and column estimates. The full flat cross-roll identity cancels the linear terms in lGap, leaving two transpose/input errors and one coarse plotted error. Hence

$$\|lGap\| \le (4C^2 + C)hq_0.$$

The full flat diagonal-roll identity cancels w against the sum of its three flat row-square operators. Applying the second estimate to F,I,N gives

$$\|dGap(w)\| \le 6C^2\|w\|\infty q_0.$$

The common constant $6C^2 + C$ bounds both. These estimates use the uncut flat identities, so both the transport error and the cutoff error are included in q_0. They hold for every cutoff R, including zero, and allow winding paths and arbitrarily short periods. The local transfer described next uses exactly these same gaps.

6.31 Localizing the mismatch bounds

Assume, in one gauge, both fine and coarse links have chords $\le h$ on the radius-d hull of a tested row, with $3R + 4 + \text{predictionReach} \le d$. The bounded logarithm extension constructs globally bounded fields a,b agreeing there with those actual links, and $\|a\|, \|b\| \le A_{\log}$ h. Put $X = \max(1, A_{\log})$ and choose $h \le \min(s_{\log}, r_global/X)$. The global logarithmic mismatch estimate then applies at envelope Xh. Locality identifies its tested gaps with the actual gauged gaps because all their reads are in the hull.

Gauge covariance rotates lGap by the root adjoint. For dGap it rotates both the row and its arbitrary input data w; this preserves the bound $\|w_j\| \le q$. Finally

$$\text{widen}(R, Xh, \tau) \le X\,\text{widen}(R, h, \tau)$$

because $X \ge 1$ and both summands are nonnegative. Thus increasing the global constant C to CX^2 gives simultaneously

$$\|lGap\| \le CX^2 h\,\text{widen}(R, h, \tau),$$
$$\|dGap(w)\| \le CX^2 q\,\text{widen}(R, h, \tau).$$

Only the independent supremum bound q on w is used. This matters when w is the actual thinned data: changing the links outside the hull need not preserve that entire data array. The global mismatch estimate and exact local gap identities proved here supply these inputs.

6.32 The quartic curvature deficit

The energy estimates separate useful positive income from the remainders that must be paid. “Income”, “price”, “loan” and “tariff” are bookkeeping names for nonnegative numerical bounds in this construction. They do not denote additional physical degrees of freedom. In particular a packet tariff is a summable upper bound on a packet’s contribution after its annular and normalization expenses have been assigned. The inequalities below show exactly which positive curvature terms absorb which remainders.

Let P_{lie} denote orthogonal projection onto the real matrix Lie algebra. For every group element U, not only those near 1,

$$P_{\text{lie}}(\text{mat}(U) - 1) = \text{colorMat}(\text{curvature}(U)).$$

Indeed subtract the skew part $(\mathrm{mat}(U) - \mathrm{mat}(U)^*)/2$ from $\mathrm{mat}(U) - I$. The remainder is self-adjoint. Its real Frobenius inner product with every skew-adjoint Lie algebra matrix is zero, by cyclicity of trace. Projecting therefore retains just the Lie projection of the skew part, which is the defining curvature.

Pythagoras for tangent and normal projections gives the exact identity

$$\begin{aligned}&\mathrm{colorScale}^2[\mathrm{chord}(U)^2 - \|\,\mathrm{curvature}(U)\|^2]\\ &\quad = \|\,\mathrm{mat}(U) - 1 - P_{\mathrm{lie}}(\mathrm{mat}(U) - 1)\|^2.\end{aligned}$$

The analytic focus fixes each group matrix near 1, and its derivative there at 1 is P_{lie}. Its Taylor remainder consequently gives constants $C, r > 0$ such that

$$\|\,\mathrm{mat}(U) - 1 - P_{\mathrm{lie}}(\mathrm{mat}(U) - 1)\| \le C\|\,\mathrm{mat}(U) - 1\|^2$$

when $\|\,\mathrm{mat}(U) - 1\| \le r$. This is an ambient analytic estimate restricted to group points; it requires no logarithm branch on an entire configuration. Substituting the chord normalization in the squared bound and setting $K = C^2\ \mathrm{colorScale}^2$ proves

$$\mathrm{chord}(U)^2 - K\,\mathrm{chord}(U)^4 \le \|\,\mathrm{curvature}(U)\|^2.$$

Choose $\delta = \min(r/\,\mathrm{colorScale}, \sqrt{1/(4K)})$. If $\mathrm{chord}(U) \le \delta$, then $K\,\mathrm{chord}(U)^2 \le 1/4$, so the last bound is at least $\frac{3}{4}\,\mathrm{chord}(U)^2$. Both constants depend only on G and are chosen before the block length. This proves the small-curvature input to the soft and coarse-C estimates.

6.33 Soft curvature and noise remainders

Define $\mathrm{flaw}(g) = [\mathrm{chord}(g)^2 - \|\,\mathrm{curvature}(g)\|^2]/2$. It is nonnegative by the curvature norm bound. On a retained fine plaquette whose toll is at most its clipping threshold, expanding the two direct/path moment coordinates gives the exact identity

$$crown = direct \cdot flaw(faceU) + \sum_{i \in Bool} \|entry_{(face,i)}U\|^2/2.$$

Adding the guide cross term and its square completes each square. Thus workF equals direct·flaw plus the two half-squares of bill $= entry - guide$. The noise identity holds without a retention hypothesis: workI equals penalty·flaw(VQ^{-1}) plus the single bill half-square. These are algebraic identities, using the actual $Q = \mathrm{prediction}(U, k)$ and the same guide.

The affine comparison already proved bounds entry by $M(x)$ of rate $-1/2$ and bill−reply by $E(x)$ of rate -1; subtracting the common guide does not change that affine error. The coarse maximum condition bounds guide by A·fiery, so bill has norm at most $m = M + A \cdot$ fiery. For vectors v,w with $\|v\| \le m$ and $\|v - w\| \le E$, factor the difference of squared norms to obtain

$$|\|v\|^2/2 - \|w\|^2/2| \le E(2m + E)/2 =: Z.$$

This has rate $-3/2$.

The small-curvature estimate supplies $J > 0$ and a chord neighborhood with

$$\begin{aligned}\mathrm{chord}(g)^2 - J\,\mathrm{chord}(g)^4 &\le \|\,\mathrm{curvature}(g)\|^2,\\ \tfrac{3}{4}\,\mathrm{chord}(g)^2 &\le \|\,\mathrm{curvature}(g)\|^2.\end{aligned}$$

The local chart-mark estimate places the plaquette holonomy and V Q^{-1} in that neighborhood. A scaled curvature bound $\|\sqrt{k}\ \mathrm{curvature}(g)\| \le M$ therefore gives $\mathrm{chord}(g) \le 2M/\sqrt{k}$ and

$$|flaw(g)| \le (J/2)(2M/\sqrt{k})^4.$$

For the fine and noise terms put respectively

$$X = direct \cdot (J/2)(2M/\sqrt{\text{direct}})^4,$$
$$Y = penalty \cdot (J/2)(2M/\sqrt{\text{penalty}})^4.$$

Both have rate -2. The common envelope $F = X + Y + 2Z$ thus has rate $-3/2$ and bounds the difference between each soft energy and its corresponding affine reply half-square sum. The two in the last term pays the two fine coordinates; the noise row needs only one. The small-curvature and local branch-activation estimates above supply this reduction.

6.34 Outer N, C and mismatch corrections

On a clean radius-B/4+1 hull, flatten the coarse field and use its bounded logarithmic extension. The guide comparison for the N row has zero linear term and envelope e of rate -1. Gauge covariance preserves its norm. Since workN is its half-square, $|workN| \le e^2/2$ has rate -2.

For the C row, the hypothesis toll$(V) \le$ temp makes its clipping multiplier exactly one, also at zero toll. The exact C-row formula and curvature norm bound then show $workC \le 0$. The quartic curvature inequality above gives

$$|workC| \le direct_{\text{below}} J/2(2\,\text{fiery})^4,$$

because temp ≤ 2 fiery. This envelope has rate -2.

For the L and D mismatch rows, use the common local gauge and logarithmic extensions of the actual fine and coarse inputs on their clean grafted hull. The layered extension gives a common envelope W of rate $-1/2$. The interior kernel comparison supplies constants $C, \tau > 0$ and, writing $e = \text{widen}(R, W, \tau)$,

$$\|lGap\| \le CWe,$$
$$\|dGap(vision)\| \le Cde, d = A \cdot \text{fiery}\,.$$

The maximum condition gives $\|vision\| \le d$. Hence Cauchy–Schwarz bounds each L inner product by $M = d$ C W e and each D half-inner-product by $V/2 = d$ C d e/2. There are two Bool coordinates, giving

$$|workL| + |workD| \le 2M + V.$$

Broadening preserves the rate $-1/2$ of W; both products have rate $-3/2$. This proves the asserted outer-envelope assembly from the layered-extension and interior-gap estimates, with actual U,V,Q shared by the two mismatch terms.

6.35 Light estimates on the annulus

The light estimate is required in two forms: a lower bound at every row and an absolute bound when that row is selected. The source decomposes it into four already specified energy types: the soft curvature/noise remainder, the N correction, the coarse C correction, and the L/D corrections. Let their envelopes be E_1, E_2, E_3, E_4 respectively. The first and fourth have rate $-3/2$; the other two are weakened to that rate. Their sum E_light is nonnegative and has rate $-3/2$.

For a selected curvature row, its extracted energy is exactly the sum of its two real-path half-squares. The soft curvature estimate consequently gives the required absolute remainder bound. For a selected noise row its extracted energy is its single real-path half-square, and the exact noise split reduces to the soft noise estimate. At an unselected curvature or noise row the extracted term is zero; the remaining energy is nonnegative, being the sum of a nonnegative flaw cost and half-squares. It therefore satisfies the lower bound without an absolute-value claim. The C row has no extracted parent contribution and is bounded by its C, L and D envelopes; the N row similarly uses its N envelope.

These statements apply to a clean radius-B coarse hull and a grafted tangent bound $w(x)$ of rate $-1/2$. Under a signed coordinate change first restrict the grafted hull to radius B/2+1, then turn it to obtain radius B/2; the extra site pays the reversed-link shift. The coarse maximum hypothesis is read in the same turned background.

For a packet annulus, take a row in the 200B seed dome but outside the 2B dome, together with the Wide bounds through 300B and tangent control through 304B. Local Interior bounds and shell access give the required grafted condition, provided $\gamma \cdot 2(width \cdot t) \le w(x)$. The same Wide bounds give the clean hull and make the original debit equal light: retained curvature selectors are one by the actual local toll bounds. Thus $debit \ge -E_light$ on the entire annulus. This uses no derivative of a hard selector and no globally small gauge. The soft-remainder, outer-correction and grafted-signal estimates of this chapter provide these four envelopes and the required shell access.

6.36 Physical operator and path-series bounds

We record why the allowances and maximum responses in the energy argument are uniform in the torus and are chosen before L. Suppose a cover kernel between types α, β satisfies $\sum_z e^{\eta|z|}\|K_{\alpha\beta}(z)\| \le C$, with $\eta \ge 0$. Each physical summand is a scalar kernel coefficient times adjoint transport along its prescribed route, subject to the endpoint congruence and route cutoff. Adjoint transport is an isometry, including on winding routes; its operator norm is at most one. Thus the summand is bounded by the absolute cover coefficient with its endpoint indicator.

For a fixed input site and cover displacement, the endpoint congruence selects exactly one output site. For a fixed output site and displacement it selects exactly one input site. Interchanging the finite site sum and absolutely convergent cover sum therefore bounds a physical row by $|\beta|C$ and a column by $|\alpha|C$. This argument takes absolute values before periodization, retaining every displacement when endpoints coincide.

For the fine-to-coarse kernel, the input type includes its fine residue and its two-layer plaquette type. The per-type cover price is C/L^2; there are L^4 residues. The row bound is consequently a/L^2 and the column bound aL^2. Noise and tight kernels have row and column bounds at most a common a. One may take a to be the maximum of one and the cover price multiplied by each of the three finite type cardinalities. This choice is independent of L, W.

The support assertion concerns kernel incidence: a nonzero cut term has endpoint distance at most its route length, at most $B/100000$, hence at most $B/50$. For fine routes use the parent map, which takes adjacent fine sites to equal or adjacent coarse sites. The route from the lifted fine residue to Lz ends at the embedded coarse site $\mathrm{parent}(x) + \mathrm{wrap}(z)$; this remains true for wrapping routes by additivity of wrapping and the cover route endpoint identity. No larger support claim about all data read by an operation is needed.

For path shortening, retain the entire label index i, writing its coefficient and slot as (c_i, s_i). Define the nonnegative matrix

$$M_{pu} = \sum_i |c_i| \, \mathbf{1}_{\{d\,\mathrm{length}(s_i) \le R,\ \mathrm{landing}(s_i,p)=u\}}.$$

The weighted path hypothesis is $\sum_i |c_i| e^{\sigma \mathrm{length}(s_i)} \le P$ for each source plaquette type. It gives row sums at most P. For a fixed label and source type, the landing equation fixes its source site uniquely; summing over source types gives column sums at most qP. Absolute convergence permits all these finite-sum exchanges. Curvature norm is bounded by plaquette chord and adjoint transport preserves norm, so

$$\| \operatorname{shorten}(U,p) \| \le \sum_u M_{pu} \operatorname{toll}(U,u).$$

If the selector keeps p, it bounds the toll at distance k by $te^{k/\rho}$. Landing distance is at most slot length, and the stage hypothesis $\sigma\rho \ge 1$ makes the weighted path sum applicable. Thus the shortening norm is at most Pt.

The common stage capacity

$$A_{\mathrm{stage}} = (q+1)P + \sqrt{\mathrm{direct}}$$

dominates row sums, column sums, and the direct clamp scale. Both branches of the thinned field therefore have norm at most $A_{\mathrm{stage}} t$: its shortening branch uses the kept estimate, and its direct branch uses the multiplier $\min(1, t/r)$ and curvature norm at most r. At $r = 0$ the curvature vanishes.

The required reach margins hold eventually. If $t = \mathrm{temp}(x+a)$ and $te^{k/\rho} < 2$, then $t \ge 1/\sqrt{x+a}$ implies $k/\rho < \log x$ once $2\sqrt{x+a} \le x$. For large $\log x$, this gives $k + 1 \le B/50$, since $B \ge (\log x)^2$. The cut path length $\mathrm{till}(x)+1$ also lies below $B/50$. Parent distances cost at most the fine distance plus one, by the parent foreshortening inequality. These observations supply both the shortening stencil reach and the discarded-witness reach used above.

Finally $a = G.\mathrm{slopeConstant} \log L \ge 0$: the minus Killing form is a sum of squares, so the Casimir and slope constant are nonnegative. At fixed shift a, eventually $\mathrm{temp}(x) \le 2\mathrm{temp}(x+a)$; use $\sqrt{x+a} \le 2\sqrt{x}$ and monotonicity of the nonnegative logarithmic factor in the temperature. Put $A = A_{\mathrm{top}} + A_{\mathrm{below}}$, and let C be the common operator allowance. The coarse vision has norm at most $2At$. Applying the fine and noise operator row bounds to it gives

$$\|\mathrm{vision}\| \le dt, \qquad \|Mv\| \le dL^{-2}t, \qquad \|Mw\| \le dt, \qquad d = 2A(1+C) > 0.$$

Both stage lane counts may use this same d. It is fixed before L, W, while the common late cutoff is taken afterwards. These prove the operator and maximum hypotheses used throughout the preceding energy estimates.

6.37 Fine and noise income, including discarded rows

Here is the summation argument producing the fine and noise inequalities used below. Write $a > 0$ for the direct quadratic coefficient, $b > 0$ for its linear coefficient, $p > 0$ for the noise penalty, and $t > 0$ for the fine threshold. Let r_i be the plaquette chord, K the listed plaquettes with $r_i > t$, and $Z = \sum_{i \in K} r_i$. The actual plot conditions imply that every listed plaquette outside K has $t/4 \le r_i \le t$. Any plaquette with $r_i > t$ whose parent is in the required five-radius dome is listed. These are local statements about the one configuration under consideration.

The two components of a fine row, with their corresponding coarse responses v_i^0, v_i^1, are

$$D_i = \mathrm{clip}_t(r_i) + \tfrac{1}{2}\|v_i^0\|^2 - \langle v_i^0, \mathrm{shoot}_i \rangle, \qquad E_i = \mathbf{1}_{\mathrm{kept}(i)} \tfrac{1}{2}\|u_i\|^2 + \tfrac{1}{2}\|v_i^1\|^2 - \langle v_i^1, u_i \rangle.$$

Here $\| \mathrm{shoot}_i \| \le \sqrt{a}\, r_i$, and clipping is $ar_i^2/2$ for $r_i \le t$, and btr_i otherwise. Suppose the previously constructed response bounds give $\|v_i^\nu\| \le \epsilon t$. Completing the scalar norm square gives

$$D_i \ge \tfrac{1}{2}(\sqrt{a}\, r_i - \|v_i^0\|)^2 \quad (r_i \le t).$$

Thus, if $\epsilon \le \sqrt{a}/8$, each listed small plaquette contributes at least $at^2/128$. If $\epsilon\sqrt{a} \le b/4$, each large listed plaquette contributes at least $3btr_i/4$. Unlisted rows in the summation region have $r_i \le t$ and hence nonnegative direct part. Consequently

$$\sum_i D_i \ge \frac{3b}{4} tZ + \frac{a}{128} t^2 |\mathrm{costs} \setminus K|.$$

It remains essential to bound the discarded second components after summation. Let $\rho > 0$ be the fixed relaxation length and

$$H_\rho = (8\rho + 1)^4 \, |\mathrm{PlaqAxes}|.$$

A row fails the kept condition only if some plaquette j satisfies $r_j > t \exp(\mathrm{dist}(i, j)/\rho)$. Such a row belongs to the ball about j of radius $\lfloor \rho \log(r_j/t) \rfloor$, including all plaquette types. For $y \ge 1$, this ball has at most $H_\rho y$ elements: use

$$2\lfloor \rho \log y \rfloor + 1 \le (8\rho + 1) e^{(\log y)/4}$$

and raise to the fourth power. The local witness-reach estimate places each witness j in the five-radius dome, so it belongs to K. The discarded set X therefore satisfies $t|X| \le H_\rho Z$. This count takes place on the original fine lattice and introduces no factor L^4.

Let $M_{ij} \ge 0$ be the previously bounded shortening majorant, with both row and column sums at most C. Its support-reach estimate shows that, for $i \in X$, every contributing $j \notin K$ has $r_j \le t$. Hence

$$\|u_i\| \le Ct + \sum_{j \in K} M_{ij} r_j, \qquad \sum_{i \in X} \|u_i\| \le C(H_\rho + 1) Z.$$

Kept rows have $E_i = \|u_i - v_i^1\|^2/2 \ge 0$. Discarded rows satisfy $E_i \ge -\epsilon t \|u_i\|$. Choosing $\epsilon C(H_\rho + 1) \le b/4$ yields

$$\sum_i (D_i + E_i) \ge \frac{b}{2} tZ + \frac{a}{128} t^2 |\mathrm{costs} \setminus K|.$$

In particular this argument never presupposes that a discarded row has a good halo.

For noise rows put $n_i = \mathrm{wire}_i$. Their exact energy is bounded below by $\frac{1}{2}(\sqrt{p}\, n_i - \|v_i\|)^2$. On listed aisles the plot gives $n_i \ge wt$. The uniform noise-response bound $\|v_i\| \le Ct$, together with $C \le \sqrt{p}\, w/2$, therefore gives $p n_i^2/8$ per listed aisle; all other noise rows remain nonnegative.

The order of these choices matters. The response theorem supplies the common constant C before L is chosen and gives $\epsilon = C/L^2$. First take L large enough for the three displayed restrictions on ϵ; next take $w \ge 2C/\sqrt{p}$; finally take the late scale parameter at those fixed values. The fine coefficient delivered by this construction is exactly $a/128 > 0$.

6.38 Macro multiplicity and the spending census

The macro census used here has a period qualification, already included in the step hypotheses. Put $R = j + 8$. The macro box radius is $(j+7)L$. If twice this radius is smaller than the fine period, two displacements in that box with the same periodic image coincide, coordinate by coordinate. Thus $M(i, z) \le 1$. A nonzero multiplicity also implies $\mathrm{dist}(i, \mathrm{parent}\, z) \le R$: multiply the coarse distance by L, use the distance from the embedded base to z, at most $(j + 7)L$, and the fine site's distance to its embedded parent, at most $L/2$. The triangle inequality gives the stated integer bound. Consequently

$$\sum_i M(i, z) \le (2R + 1)^4.$$

The port identity follows by interchanging its finite cover-displacement and plaquette-type sums and grouping identical periodic plaquettes, with their multiplicities. The period hypothesis is used for the column bound, not for that identity.

The store's base set has at most four times its number of shares: a fine plaquette has at most four corner parents, a noise bond at most two pins, and each flag contributes one site. Hence any collection of coarse faces based in its radius-r dome has cardinal at most

$$V(r)\,\text{shares}, \qquad V(r) = 4(2r+1)^4 q.$$

For $B \geq 1$, $V(2B) \leq V(2)B^4$. This proves the bounds on the outer and affected face sets, using respectively the original store and the spending store. The exact flag condition says each flag is touched by a coarse face with toll above the coarse temperature. Since a plaquette has at most four corners, the number of flags is at most four times the number of those faces.

For completeness, every coarse face whose toll exceeds a fixed $\delta > 0$ has an actual spending witness within distance R, once the scale parameter is late enough. Write T for the large spending threshold. If no spending base lies in that neighborhood, then every fine input in its macro port has toll at most T: listed inputs exceeding T would themselves be spending witnesses, while unlisted inputs have toll at most $t/2 \leq T$. Likewise every one of its four boundary wires is at most T, using the unlisted bound $2Wt \leq T$. Its toll is then at most

$$\left(4 + DL^{-2}q|\operatorname{box}((j+7)L)|\right) T.$$

At fixed L, W, this is eventually at most δ, a contradiction. All relevant inputs lie in the plot dome by the distance bounds above, so the local predictor comparison applies.

The limits used here are explicit. With fine inverse coupling $\beta_F \to \infty$, $T = \beta_F^{-49/100} \to 0$, whereas $t = \beta_F^{-1/2} P(\beta_F)$ with P smaller eventually than every positive power. For each fixed m, use $P(\beta_F) \leq \beta_F^{1/200}$ and $m \leq \beta_F^{1/200}$ to obtain $mt \leq T$. Also $B = \lceil (\log x)^2 \rceil \to \infty$, allowing $B \geq 5R$. One common pre-L census constant is

$$N = V(R) + V(2) + R + (2R+1)^4.$$

It simultaneously bounds the macro columns and reach, the number of δ-large coarse faces by N times spending shares, and the outer and affected face counts by NB^4 times their respective shares. This supplies the census used in the departure and point-charge bounds.

6.39 Transfer of affected departure costs

Fix the same plot, and let s now denote its affected coarse faces, with $m = |s|$. Use constants $a, d, D \geq 0$ for the operator allowance, maximum response, and predictor-curvature comparison respectively. These letters are local to this argument. Put

$$\mathcal{B} = L^{-2} \sum_{\text{swell}(t)} \operatorname{toll} U + \sum_{\text{aisles}} \operatorname{wire}(Q, V), \quad u = \sqrt{\text{high.direct}} + d, \quad v = \sqrt{\text{low.direct}} + d, \quad g = \sqrt{p}.$$

Retain the two nonnegative excess arrays exactly on the fine swell at threshold t and on all listed noise aisles. The macro transfer estimate constructs a nonnegative coarse excess H such that, in the three-radius dome,

$$\operatorname{toll} V \leq rt + H, \qquad \sum H \leq b, \qquad r = \text{macroSmall}(j, D, L, W), \quad b = \text{macroLarge}(N, D)\mathcal{B}.$$

This uses the actual usable-plaquette inequality $\operatorname{toll} V \le \operatorname{fuel}(Q,V) + DL^{-2}\operatorname{port}(U)$. The macro transfer estimate has the following explicit form. Let $q = |\mathrm{PlaqAxes}|$, and let $M(i,z)$ be the integer macro-path multiplicity, counting cover displacements separately even when their periodic images coincide. Its row sum is at most

$$S_{j,L} = q\,|\operatorname{box}((j+7)L)|,$$

while the macro-store census gives $\sum_i M(i,z) \le N$ on the physical torus. The port identity is $\operatorname{port}(i,U) = \sum_z M(i,z)\operatorname{toll}(U,z)$. For the retained arrays E, F, set

$$H(p) = \sum_{e\in\partial p} F(e) + DL^{-2}\sum_z M(p.1,z)E(z).$$

The boundary sum has four terms, retaining repetitions. For each fixed plaquette type and each one of its four boundary positions, the map from its base site to that bond is injective, including on the torus. Hence its total is at most $4q\sum F$. Summing the macro term and using its column census gives

$$\sum_p H(p) \le 4q\sum F + DqNL^{-2}\sum E \le (4q + DqN)\mathcal{B}.$$

Thus $\operatorname{macroLarge}(N,D) = 4q + DqN$.

In the five-radius dome, the plot conditions give $\operatorname{toll} U \le t + E$ and $\operatorname{wire}(Q,V) \le 2Wt + F$. For unlisted inputs these follow from the plot's upper thresholds; for listed inputs they follow by retaining the full input size in the excess. Every face in the three-radius dome is usable: its required coordinate box has radius one and is inside the plot's listing dome. Each of its four boundary bonds is in the five-radius dome. The four wire terms therefore contribute at most $8Wt + \sum_{\partial p} F$. The macro stencil also stays inside the five-radius dome, by its reach bound and $B \ge 1 + N$; its row sum bounds the small part of the port by $S_{j,L}t$. The predictor-curvature comparison now proves the asserted split with

$$\operatorname{macroSmall}(j,D,L,W) = 8W + DL^{-2}S_{j,L}.$$

The row count follows simply by summing the defining indicator over all cover displacements and plaquette types. It allows coincident periodic images, so no injectivity of the cover box is assumed.

The operator row and column bounds transfer the local input split to the $2m$ tested Boolean slots as follows:

$$\sum \|\Phi^* X_f\| \le 2m\,aL^2ut + 2au\mathcal{B},$$
$$\sum \|(\Phi')^* Z\| \le 2m\,ag(2Wt) + ag\mathcal{B},$$
$$\sum \|X_n\| \le 2mvrt + 2vb.$$

For the first inequality, separate the uniform small input ut from the plotted nonnegative excess. The wide column bound costs aL^2 on the small input, while the row bound a/L^2 multiplies the total excess, at most $2u\sum E$. Thus the large-input term retains its L^{-2} saving. The noise argument uses narrow row and column bounds a, and the coarse argument uses the plotted excess bound directly. In each case the stencil radius takes the tested two-radius dome into the input three- or five-radius dome, so only the given local split is used.

The departure L is the slotwise inner product of the vision, of norm at most dt, with $\Phi^* X_f + (\Phi')^* Z - X_n$. Cauchy–Schwarz gives its summed lower bound by minus dt times the sum of the three right sides above. The remaining departure D requires no large-input saving. Indeed its residual vector has norm at most

$$(d + ad + ad + a^2 d)t,$$

since the wide factor aL^2 cancels the response bound dt/L^2, while the narrow and tight terms cost adt and a^2dt. Pairing with the vision and summing the two half-weighted slots gives

$$\mathrm{workD}(p) \geq -d_0 t^2, \qquad d_0 = d(d + ad + ad + a^2 d).$$

All third-range coarse faces charged by the payments relation are contained in s. Their toll sum is at most $mrt + b$. As the coarse temperature is at most $2t$, their linear charge can consequently be included in the same bound. Writing $\ell = \mathrm{low.linear}$, the result is

$$\ell \,\mathrm{temp} \sum_{\mathrm{tides}} \mathrm{toll}\, V - \mathrm{talll}(s) \leq C_{\mathrm{pre}}\, t\mathcal{B} + C_{\mathrm{seat}}\, t^2 m,$$

where the exact available choices are

$$C_{\mathrm{pre}} = d(2au + ag + 2v\ \mathrm{macroLarge}(N, D)) + 2\ell\ \mathrm{macroLarge}(N, D),$$
$$C_{\mathrm{seat}} = 2d(aL^2 u + 2agW + vr) + d_0 + 2\ell r.$$

The first constant is fixed before L, W; the second may depend on both. The store census gives $m \leq NB^4\,\mathrm{shares(spend)}$. Thus the seat term is bounded by $(NC_{\mathrm{seat}} + 1)t^2 B^4\,\mathrm{shares(spend)}$, with a positive constant. This is precisely the affected-departure charge needed in the point-charge assembly. In particular the fine excess here is cut at t, while only the census uses the larger spending threshold.

6.40 The three coarse clipping ranges

The curvature estimate above yields the coarse row bounds without differentiating its selector branches. Let $r = \mathrm{toll}(V, i)$, $s = \mathrm{temp}(x) \geq 0$, $a = \mathrm{low.direct} > 0$, and $b = \mathrm{low.linear}$. Put $\alpha = 1$ if $r = 0$, and otherwise $\alpha = \min(1, s/r)$. The shortening components cancel exactly: on a kept row their contribution is its shortening half-square, and on a discarded row both tested shortening terms vanish. Expanding the remaining direct component gives

$$\mathrm{workC}(i) = a(\alpha - \alpha^2/2)\| \,\mathrm{curvature}(V, i)\|^2 - \mathrm{clip}_s(r).$$

Choose $K, \delta > 0$, depending on the group alone, so that for $r \leq \delta$,

$$r^2 - Kr^4 \leq \| \,\mathrm{curvature}\,\|^2, \qquad \tfrac{3}{4} r^2 \leq \| \,\mathrm{curvature}\,\|^2.$$

If $r \leq s \leq \delta$, then $\alpha = 1$ and $\mathrm{workC}(i) \geq -aKs^4/2$. If $s < r$, then $0 \leq \alpha = s/r \leq 1$, so $\alpha - \alpha^2/2 \geq \alpha/2$. Therefore

$$\mathrm{workC}(i) \geq \frac{as}{2r} \| \,\mathrm{curvature}\,\|^2 - bsr.$$

In the middle range $s < r \leq \delta$, the curvature lower bound and the clip hypothesis $b \leq a/8$ give $\mathrm{workC}(i) \geq asr/4$. In the outer range $r > \delta$, discard the nonnegative curvature term to obtain $\mathrm{workC}(i) \geq -bsr$.

Finally $s \leq \delta$ eventually. The supplied temperature estimate $x\,\mathrm{temp}(x)^3 \leq 1$ implies this whenever $x > 1/\delta^3$. Thus the same late-scale choice gives all three bounds, with small-range debit coefficient $aK/2$ and a group-dependent fixed middle-range cutoff δ.

6.41 Threshold ratios and absorption of stitch errors

All limits in this paragraph are taken after fixing L. Set $y = \beta_F = x + k$, where k is the fixed coupling shift, $B = \lceil (\log x)^2 \rceil$, $R = \lfloor B/100000 \rfloor$,

$$t = \mathrm{temp}(y), \qquad T = y^{-49/100}.$$

The temperature is $v^{-1/2}\exp((\log\log v)^2)$ for positive v in these late-scale calculations. For every fixed real shift k, $\mathrm{temp}(x+k)/\mathrm{temp}(x) \to 1$. Indeed its logarithm is

$$(\log\log(x+k))^2 - (\log\log x)^2 - \tfrac{1}{2}(\log(x+k) - \log x).$$

On $[1,\infty)$ the derivative of $(\log v)^2$ has absolute value at most two. The mean value estimate bounds the first difference by twice $|\log(x+k) - \log x|$, which tends to zero. Exponentiation proves the ratio limit.

For every $r > 0$ and integer $n \geq 0$, $B^n = o(x^r) = o(y^r)$, where the latter notation means $B^n/y^r \to 0$. To see this directly, eventually $B \leq 2(\log x)^2$, and every fixed power of $\log x$ is smaller than x^r by a factor tending to zero. The fixed shift does not affect that comparison. The monitor bound gives, for every $p > 0$,

$$v^{-1/2} \leq \mathrm{temp}(v) \leq v^{-1/2+p}$$

eventually. It follows that $B^n\mathrm{temp}(x) \to 0$ and

$$B^n t/T \to 0, \qquad B^n T \to 0, \qquad B^n T^2/t \to 0.$$

For the first limit choose $p = 1/200$, leaving the decaying power $y^{-1/200}$; for the last use the lower bound on t, leaving $y^{-48/100}$. In particular the square of T, rather than T itself, is the quantity small relative to t.

For each $\tau > 0$, $e^{-\tau R} \leq y^{-1}$ eventually. In fact $B \leq 100000R + 100000$, while $\log y \leq 2\log x$ eventually. Since $(\log x)^2 \leq B$, the quadratic growth of this lower bound for R dominates $2\log x/\tau$. Thus

$$B^n e^{-\tau R} \to 0, \qquad B^n e^{-\tau R}/t \to 0.$$

Define the stitch size $h = CB^{10}T$ for fixed $C > 0$, and its widening error $e = (1+R)h + e^{-\tau R}$. The preceding limits, with $n+10$, $n+20$, or $n+1$ in place of n, give

$$B^n h \to 0, \quad B^n h^2/t \to 0, \quad B^n e \to 0, \quad B^n he/t \to 0.$$

For the last two statements use $R \leq B$, expand e, and bound each nonnegative summand by one of the already established limits; in the mixed tail term also use $h \to 0$. These are the precise limits that pay the unaffected-row stitch losses after multiplying by the dome census. They involve no torus period or number of store shares.

6.42 Finite-store budget inequalities

Write $a = \mathrm{linear}_{\mathrm{top}}$, $p = penalty$, $t =$ fiery and $u = largeT \geq t$. Let F_t be the listed fine plaquettes with toll $> t$, F_u those with toll $> u$, and $R = \sum_{F_t}$ toll. Then $t|F_t| \leq R$ and $u|F_u| \leq R$. The supplied fine tally bound is (a/2)tR+d $t^2|costs$ minus $F_t|$. If $m \leq a/8$, $m \leq d$ and $c_F \leq a/4$, use $2m + c_F \leq a/2$ and split costs into F_t and its complement to obtain

$$tallyF \geq mt^2|costs| + mtu|F_u| + c_F tR.$$

This separately pays all original fine shares, the larger spend subset, and the fine mismatch fee.

Every listed noise bond satisfies wire $\geq Wt \geq t$. With $R_I = \sum_l isted$ wire^2, this implies $t^2|aisles| \leq R_I$. For the spend subset wire $> u$, also $tu|\operatorname{spend}_a isles| \leq R_I$. If $C \leq pW/16$, then C $t \cdot$ wire $\leq (p/16)$ wire^2 bondwise, since $Wt \leq$ wire. Using $m \leq p/32$, the supplied noise bound $tallyI \geq (p/8)R_I$ therefore gives

$$\begin{aligned} tallyI &\geq mt^2|aisles| + mtu|\operatorname{spend}_{\text{aisles}}| \\ &\quad + Ct \sum_{\text{listed}} \text{wire}\,. \end{aligned}$$

The listed-wire lower bound comes directly from Plot read at that bond's endpoint, which belongs to the store dome.

For the coarse row split outs into three ranges: toll $\leq s =$ temp, $s <$ toll $\leq \delta$, and toll $> \delta$, where $s \leq \delta$. The coarse row estimate gives respectively a lower bound $-B_q uartic\ s^4$, a positive bound at least 4m t^2 when $m \leq$ direct$_b\, elow/64$ and $t \leq 2s$, and a lower bound $-$ linear$_b\, elow$ s·toll. Sum these disjoint ranges, enlarging the small-range count to all outs because its error is nonnegative. This proves

$$\begin{aligned} tallC + linear_{\text{below}} s \sum_{toll>\delta} \text{toll} \\ + B_{\text{quartic}} s^4 |outs| \\ \geq 4mt^2 \cdot number(s < \text{toll} \leq \delta). \end{aligned}$$

Finally Payments bounds flags by four times the number of s-large tides and bounds the δ-large tide count by K shares(spend). Since the latter tide set is a subset of the former, subtract it and obtain

$$|flags| \leq 4[number(s < \text{toll} \leq \delta) + K \operatorname{shares}(\text{spend})].$$

These are exactly the four finite-store inequalities used in the assembly. They are algebraic consequences of the fine/noise/coarse producers and Payments, with no compatibility assumption on the periods. In particular the fine and noise spend gains are paid at tu, while the original seed shares are still paid at t^2.

6.43 Assembly of the point-charge gain

Use the same store, actual input, predictor and background in all estimates below. Put $t =$ fiery(x), $s =$ temp(x), $u = largeT(x)$, and $B = radiusB(x)$. The fine energy estimate bounds its tally below by

$$\begin{aligned} linear_{\text{top}}/2t \sum_{largefinetolls} \text{toll} \\ + dt^2 \cdot number(otherlistedfinetolls), \end{aligned}$$

with $d > 0$ after choosing L,W. The noise estimate gives $tallyI \geq (penalty/8) \sum_l isted$ wire^2. The coarse estimate, store census and hit-row mismatch estimate supply the remaining terms described below.

Choose a positive m no larger than linear$_{\text{top}}$ /8, d, penalty/32 and direct$_b\, elow/64$. The hit mismatch estimate has fine-loss coefficient $c_F = C/L^2$ and noise-loss coefficient C:

$$\begin{aligned} tallL_D(hit) + c_F t \sum_{\text{large}} finetolls + Ct \sum_{\text{listed}} \text{wire} \\ +Zt^2B^4 \cdot \text{shares(spend)} \\ \geq linear_{\text{below}} \cdot s \sum_{\delta-largecoarsetolls} \text{toll}\,. \end{aligned}$$

Choose L first so that $c_F \leq \text{linear}_{\text{top}}/4$. Choose W afterwards with $W \geq 1$ and $C \leq penalty \cdot W/16$. Since every listed noise wire is at least Wt, its quadratic positive energy absorbs the linear C t wire loss. The fine positive energy similarly absorbs the c_F loss. The coarse inequality pays the large coarse-toll term, leaving only the small quartic coarse error.

The flag census is

$$\begin{aligned} number(flags) \leq 4[number(tides_s minus tides_\delta) \\ +K \cdot \text{shares(spend)}]. \end{aligned}$$

The output-row census is at most K B^4 shares. Choose x large so that $s \leq \delta$, $t \leq u$, $t \leq 2s$, and

$$\begin{aligned} (4mK + ZB^4)t^2 \leq mtu, \\ C_{\text{quartic}} s^4 K B^4 \leq (m/4)t^2. \end{aligned}$$

The first inequality lets the large toll/noise spend pay both its own flags and the mismatch remainder; the second pays the small coarse error by one quarter of the desired gain. The unaffected mismatch estimate costs at most another $(m/4)t^2$ shares. The N tally is nonnegative.

To obtain these last two scale inequalities, use $t/s \to 1$, $t/u \to 0$ with $B^4 t/u \to 0$, and $B^4 s^2 \to 0$. Dividing the first desired inequality by tu reduces it to a quantity tending to zero. For the second use $s \leq 2t$ and the vanishing $B^4 s^2$. These choices are uniform in the live period.

Split the L/D tally into hit and unaffected rows. Add the fine, noise and coarse budget inequalities, the flag census, and the two quarter-error bounds. Their positive remainders pay m t^2 per listed fine/noise/flag share before the two quarter losses. The result is

$$assessed \geq (m/2)t^2 \,\text{shares}\,.$$

The exact fine/noise/coarse budget rearrangements supplying these positive remainders are developed next. This assembly fixes the quantifiers: the fine coefficient is paid by L, the noise coefficient by W, then $m > 0$ is fixed, and finally the common delay pays Z and the quartic and unaffected errors. No later change of L or W is used.

6.44 Unaffected mismatch costs in the point-charge proof

The point-charge producer combines the E01–E09 energy estimates with one additional unaffected-row estimate. The latter is derived from the mismatch bounds proved above. Fix an unaffected coarse plaquette row, and let $t =$ fiery. The coarse maximum gives $\|vision\| \leq A$ t. The stitching estimate supplies a gauge in which the actual fine and coarse links are bounded on its radius-2B/5 hull by $h(x) = stitchCost(x)$, with no change of forest or tags. Choose the delay so that this hull contains 3 till(x)+4+predictionReach and h lies in the interior chart radius. Put $e =$ widen$(\text{till}(x), h, \tau)$.

The local gap bounds and Cauchy–Schwarz give

$$workL + workD \geq -loss,$$
$$loss = 2(At)(Che) + (At)(CAte).$$

The factor two pays the two L coordinates; the two D half-inner-products together give the last term. This estimate applies to each row outside the hit subset of the store's output rows.

The store census bounds their number by N B^4 shares. The stitching-scale limits give

$$B^4he/t \to 0, B^4e \to 0.$$

Indeed after dividing N B^4 loss by t^2 its two terms are exactly 2NAC B^4he/t and NA^2C B^4e. Hence, for every prescribed $\varepsilon > 0$, the total unaffected mismatch cost is at least $-\varepsilon$ t^2 shares eventually. Sum the row bounds with the unchanged actual U,V and predictor. This proves the additional unaffected-row estimate rather than assuming it in the point-charge chain. The stitching bound, its two displayed scale limits, and the census and maximum estimates of this chapter supply its inputs.

The final point-charge statement has order of choices: choose a block lower bound n, then $L \geq n$; choose W sufficiently large, then a positive μ; finally choose x sufficiently large, uniformly over all periods with coarse side at least 100(1+B), all funded step data, forests and actual Plot inputs. It asserts $assessed \geq \mu\ t^2$ shares pointwise. Passing from the supplied grant to this statement just intersects its delay with positivity of x and t and specializes its universal query. The fine, noise, coarse-clipping and flag budgets in this chapter prove this combination of energy estimates.

6.45 Exact assignment of the energy to packet cores

A row location has four types: a fine plaquette, a coarse noise bond, a coarse plaquette carrying C together with L and D, or an N row. Their principal values are respectively workF, workI, workC+workL+workD and workN, all at the same actual input, coarse background and predictor. Their coarse anchor is the parent of the fine plaquette root or the coarse root of the other types.

A fine plaquette has two parent moment rows; a noise bond has one; the remaining row locations have none. Each moment row occurs as a parent exactly once. Select a signed row when its signed coarse anchor lies outside the 6B ambient seed dome. Its surcharge is the corresponding real affine path half-square if selected and zero otherwise. Extracted is the sum of its parent surcharges. Put $debit = principal - extracted$ and liftSq equal to the orientation average of all extracted rows.

The actual kernel localization identity, including the invariant noise sum, now gives

$$wholeGap + \sum_{\text{noise}} monetary$$
$$= \text{average}_s \sum_{\text{all}} rowlocationsdebit_s + \text{liftSq}\,.$$

This is exact recombination: finite subtraction cancels the extracted rows, while C,L,D remain bundled at their common coarse plaquette. Noise covariance under a signed change is accompanied by the actual transformed tags, and its finite bond sum is invariant under that bond permutation.

For a store Z contained in the ambient seed place, every posted 2B core row lies inside the ambient 6B dome, so it is unselected. Its extracted value is zero, and its debit equals its principal energy. Summing the four principal types over this dome is exactly $assessed(Z)$: expand the finite sums and group C,L,D at their anchor. This proves the core identity used in the packet gain without discarding any extracted energy.

Finally, a Wide condition through $d \geq 5B$ implies Plot on its 5B restriction and is preserved by a signed rotation of fields, tags, forest and store together. Apply the point-charge inequality in each rotated tree slice to get $assessed \geq \mu t^2$ shares. Rotation preserves shares. Averaging gives the mean assessed bound used above. This completes the assignment and averaging step, using the preceding point-charge estimate.

6.46 Paying a packet's annulus and obtaining its tariff

For a store Z contained in the ambient seed place, the core identity identifies its assessed energy with the sum of debits on the posted 2B rows. Its assigned energy sums the same debits through 200B. Subtract the two finite sums. Every annulus row has debit at least $-Y(x)$, by the annulus estimate already proved. The number of posted rows is at most

$$\text{shares}(Z)\,\text{ballCost}(200B+1)spots(L)$$

in each orientation, where $spots(L)$ counts the row types per coarse site. Consequently assessed$-E(x)\,\text{shares}(Z) \leq assigned$, with $E = Y$ ballCost$(200B+1)spots(L)$. The radius factor has rate zero, so E retains rate $-3/2$.

The point-charge estimate on the actual wide sector gives mean assessed energy at least $\mu\ t^2$ shares(Z). Average the preceding inequality over the same signed orientations. This proves

$$quoteNear(Z,u) \geq [\mu t^2 - E(x)]\,\text{shares}(Z).$$

Its hypotheses include Wide through 300B, tangent control through 304B, and the same shell-access bound. No connectedness of Z was required. The exact core assignment and point-charge inequality proved in this chapter supply this step.

The ancillary packet expenses can be bounded uniformly before using this gain. The all-field joint diagnostic bounds the condensation by a constant. The Haar logarithm is eventually bounded by one. Each swept neighborhood and the ticket fee is at most a fixed constant times bill$(x)^4$, and the negative logarithm of the frozen measure is at most a fixed constant times bill(x). Thus for one fixed $K_up > 0$,

$$\begin{aligned}&\text{steamLog} + tollFee \cdot \text{shares}\\ &\leq -\beta quoteNear + K_{\text{up}}\,\text{bill}(x)^5\,\text{shares}\,.\end{aligned}$$

Now consider an admitted nonzero packet. Its seed-consistency indicator supplies containment in the ambient seed place and the correct actual flags. Nonzero ticket and sector profiles imply the Wide hypotheses on its 300B dome. Its nonzero enforcement profile supplies hot tangent control 2a t through 304B. The width condition $\gamma \cdot width \leq a$ supplies the matching frozen tangent bound. Hence the lower bound for quoteNear applies with that same field, forest and tags.

Since E has rate $-3/2$ and $\beta \leq 2$ betaF, eventually $\beta E \leq 1 \leq \text{bill}^5$. Combining the last two bounds and exponentiating yields

$$|Pack(Z,u)| \leq \exp\left[1 + K_{\text{upbill}}(x)^5 - \beta\mu t^2\right]\text{shares}(Z).$$

This is the packet income with cap 1, per-share expense $1 + K_upbill^5$ and per-share gain $\beta\mu t^2$. A zero packet satisfies it immediately. The estimate is pointwise before the frozen-Haar and Gaussian integrations and applies to the entire packet factor once, including its own tag and guard costs.

6.47 Compulsory packet factors and normalization expenses

For a seed store A, let $\mathrm{visit}_d(A)$ be the fine bonds with at least one coarse endpoint in its radius-d dome, and $pens(A)$ the nontrunk bonds in $\mathrm{visit}_{(10B)}\,A$. Let

$$m_g(\beta) = \mathrm{chartScale}(\sqrt{2\pi}\,\beta^{-1/2})^{\dim(\mathrm{Color})}.$$

The packet exponent is defined exactly by

$$\begin{aligned}\mathrm{steamLog} = &-\beta[quoteNear(A,u) + farBills(A)]\\ &+\sum_{e\in \mathrm{visit}_{200B}(A)} condensation(e)\\ &+\sum_{e\in \mathrm{plaza}(\mathrm{dome}_{304B}(A))} gentry(e,u)\\ &-|pens(A)|\log m_g(\beta).\end{aligned}$$

Here farBills is the signed-orientation average of the extracted mean half-squares at the posted 200B rows; it is nonnegative. The entire packet factor is exp(steamLog) times the contraction-order profile, the actual tag tickets on the 300B dome, the sector profiles there, and tangent enforcement on the 304B plaza. It is nonnegative. A nonzero factor forces every individual enforcement cone in that plaza to be nonzero.

The order, sector and enforcement profiles lie in [0,1]. The tickets have upper bound $\exp(\mathrm{tollFee}(patch, 300B)\cdot \mathrm{shares}(A))$. Hence

$$|steam| \le \exp[\mathrm{steamLog} + tollFee\cdot \mathrm{shares}(A)].$$

This bound keeps the ticket cost of the whole packet; it is not repeated separately in neighboring log components.

Suppose $|condensation| \le Q_0$, $|gentry| \le C_0$, and $-\log m_g \le K_0$, with these bounds nonnegative. Endpoint census gives

$$|\,\mathrm{visit}_d(A)| \le \mathrm{shares}(A)\,\mathrm{ballCost}(d+2)L^4|Axis|.$$

The gentry plaza is contained in $\mathrm{visit}_{304B}(A)$, and pens in $\mathrm{visit}_{10B}(A)$. Dropping the nonpositive term $-\beta$ farBills yields

$$\begin{aligned}\mathrm{steamLog} \le &-\beta quoteNear\\ &+\mathrm{shares}(A)[Q_0 swept(200B+2)\\ &+C_0 swept(304B+2) + K_0 swept(10B+2)],\end{aligned}$$

where $swept(d) = \mathrm{ballCost}(d)L^4|Axis|$. This is the exact ancillary expense to be compared with the packet's core gain.

The measure expense has a uniform logarithmic bound over $0 < \beta \le 2$ betaF. Expand

$$\begin{aligned}-\log m_g &= -\log chartScale\\ &-\dim(\mathrm{Color})\log\sqrt{2\pi} + (\dim(\mathrm{Color})/2)\log\beta.\end{aligned}$$

Its upper bound depends only on log(2 betaF), so eventually it is at most $C_measure\cdot Delay.\mathrm{bill}(x)$ for one positive constant. No lower beta-band bound is needed for this upper estimate. The point-charge core bound and the packet annulus estimate show that the gain in quoteNear dominates these counted expenses.

6.48 The short correction family

The short corrections and Gaussian chains are organized so that locality can be priced before the full connected expansion is summed. A chain letter remembers its genuine path support and itinerary; changing to a smaller nominal support would invalidate the later hit counts. The inverse and logarithm series are absolutely convergent under the displayed smallness conditions, so their pointwise exponent identities may subsequently be integrated.

Budget language as inequalities The correction families are organized by the upper bounds they spend and the positive estimates that absorb those bounds. A price is an explicit coefficient in such an inequality; income is a proved lower estimate; and a payment is the displayed domination after the parameters have been chosen in the permitted order. These terms do not describe a new probability law. Retaining their weights and exponents is essential because the same finite store must absorb the fine, noise and coarse contributions simultaneously.

The local short letters have three kinds. At a signed row root o outside the 200B seed neighborhood, the energy letter is

$$-\beta orientation \cdot [enforce_{(2a)}(u, booth(o)) \cdot light(u) + extracted(\mu)].$$

It is zero inside that neighborhood. The second letter is the local condensation term, set to zero when either endpoint of its fine bond touches the 200B seed dome. The third is the cut-off Haar logarithm gentry, retained only at hot links outside the plaza of the 304B seed dome. Their geometric sleeves are fixed neighborhoods of their row or bond captions.

On the original seeded sector the light term equals the original energy debit at a remote row. The noise row equality is definitional. For the curvature row, turn the actual fields, tags and seed store together; the seeded sector and positive likelihood give the local Interior bounds. If the soft retention threshold at a possible curvature read is below 2, the room inequality places that read in this interior and its plaquette toll is at most $t/2 \le t$. If the threshold is at least 2, the universal chord bound supplies retention. Thus the retained high-stage path field is unchanged there, and the exact split-curvature identity identifies light with debit. This reasoning uses actual sector membership only for this equality, not for the all-field bounds below.

Let $A(x)$ be the protected-mean path envelope already proved, of rate -1. The extracted term is a sum of at most two nonnegative half-squares of parent path values. At a remote row its absolute value is therefore at most $2(A^2/2) = A^2$.

Take $J = E_1 + E_2 + E_3 + E_4$ from "Light estimates on the annulus," with grafted tangent bound $4a\,\mathrm{fiery}(x)$. The soft-square and mismatch estimates there have rate $-3/2$, and the two quartic or squared-guide corrections have rate -2; weakening the latter gives J rate $-3/2$. Take $E_{\mathrm{cond}} = \dim \mathfrak{g}\ \mathrm{pulseHead(repair)}$ from the preceding local plate comparison. Its repair envelope has rate $-1/2$, its window census has rate zero, and its fixed inverse-window constants do not depend on x, so E_{cond} has rate $-1/2$. These are the actual two bounds needed here. The Haar estimate gives $|\,\mathrm{gentry}\,| \le C\,\mathrm{fiery}$. To use the light estimate for arbitrary input u, split according to whether its enforcement profile is zero. If it is nonzero, its protection property bounds every tangent value read by the booth by $4a\,\mathrm{fiery}$. The distant-field identity then supplies precisely the grafted condition on the common clean radius-B hull. The profile is at most one, so its product with light is bounded by J. Outside the $200B$ exclusion all these hull and parent-path conditions follow from the same seed separation.

For $0 < \beta \le 2$ betaF and $0 \le orientation \le 1$, the energy letter is bounded by

$$W_{\mathrm{short}}(x) = 2betaF(x)[J(x) + A(x)^2].$$

Its rate is $-1/2$: betaF has rate 1, J has rate $-3/2$, and A^2 has rate -2. The condensation and Haar letters have bounds E_cond and C·fiery of rate $-1/2$ as well. Thus their sum

$$E_{\mathrm{short}} = E_{\mathrm{cond}} + C \cdot \mathrm{fiery} + W_{\mathrm{short}}$$

is a common eventually nonnegative envelope of rate $-1/2$ for all three kinds, all fields and all admissible β. A bond whose endpoints touch a dome has its parent in the one-site enlargement; conversely a bond not touching it has its parent outside it. These elementary endpoint facts justify the condensation exclusion used here. The light-remainder, remote-condensation and grafted-signal estimates of this chapter establish the three inputs to this assembly, together with the mean and Haar bounds.

6.49 A fixed multiple of the small-field scale for the base stroke

Interpreted typography. The bound below is read as $4\,\mathrm{pcap}(4)(\sqrt{\mathrm{direct}_{\mathrm{top}}} + \mathrm{paths}_{\mathrm{top}})t$. Both the direct row and the summed path row carry the common small-field factor t. This repairs the extraction's missing parentheses and mistyped terminal subscript; it is the convention used in the subsequent base-stroke constant C_1.

For an integer n define $pcap(n) = 32(2(n+1)+1)^5|PlaqAxes|$. Axial flattening on radius n+1 and temp $\le 2t$ make the background links on radius n have chords $\le pcap(n)$t. The crucial point is that n is fixed separately for each finite query below, rather than increasing to B.

First consider any fine plaquette of replica(V) with a clean radius-five hull about its coarse parent. Use $n = 4$. Replica links in that hull are either identity or a coarse link in the chosen gauge, and hence have chords $\le$ pcap$(4)t$. Replace links outside the hull by 1. Its plaquette product is unchanged, so its chord is at most 4 pcap(4)t. The same bound applies to its curvature norm.

For the fine base moment row, every retained path slot ends at a plaquette whose parent has a clean radius-five hull inside the original radius-B clean region. Its actual path length is bounded by till(x), and the room inequalities ensure this inclusion. Its adjoint transport preserves curvature norm. Summing against the absolute path coefficient budget, and including the direct curvature coordinate, therefore gives

$$\begin{aligned} &\|Xf_R(\mathrm{replica}(V))\| \\ &\le 4\,\mathrm{pcap}(4)(\sqrt{\mathrm{direct}_{\mathrm{top}}} + \mathrm{paths}_{\mathrm{top}})t =: C_1 t. \end{aligned}$$

No single gauge is needed for every slot: only its curvature norm is estimated.

For the noise base stroke, instead fix $n = \mathrm{predictionReach}(D)$. Flatten on that fixed hull and extend its logarithm to a global bounded field b with $\|b\| \le A$ t, $A = A_{\log}\, pcap(n)$. Its linear copy a has norm at most $\|b\|$ and exponentiates exactly to replica(exp b). Predictor locality identifies the actual base noise stroke in this gauge with zTest(a, b), with its original tags. Its uniform Taylor estimate gives

$$\|zTest(a,b)\| \le J\|(a,b)\|^2 + J\|(a,b)\| \le 2JAt$$

eventually, since $At \le 1$ and the chart radius. Gauge covariance transfers the bound back. Thus $C_1 + 2JA$ is one positive constant bounding every clean-hull base stroke by that constant times t. This proves the crowns input to the global mean bound without spending an extra growing-radius factor.

6.50 Global mean and stationary residual

The global mean bound uses selected rows, the frozen cutoff field, and the inverse-window bounds; it is separate from the sharper remote path estimate. Let the clean-background base stroke be bounded by $A_0 t$, every guide by Mt, the oriented cutoff operator by N, and the saved field by $\gamma \cdot 2\mathrm{Kt}$, where K is the fixed shell width. On a selected row, seed separation supplies the clean hull needed for the base-stroke estimate. At zero hot input,

$$\|priced(0)\| \le scale[A_0 + M + 2N\gamma K]t =: b_1 t.$$

An unselected row is zero and satisfies the same bound. This uses the saved-field bound globally, without requiring a common global gauge.

The linear bill is the transpose form applied to priced(0). Its row-sum bound Y gives $|b_i| \le Y b_1 t =: b_0 t$. Let P be the inverse-column pad and $\mu_c ount = -P^{\mathsf{T}} b$. Its column bound $C_{\mathrm{inv}} = payE$ gives $|\mu_c ount, i| \le C_{\mathrm{inv}}\ b_0 t$. Summing its color-coordinate bounds yields $\|\mu_e\| \le \dim(Color) C_{\mathrm{inv}}\ b_0 t$.

The stationary residual has a useful exact form. Symmetry of A^{-1} gives

$$A\mu_{\mathrm{count}} + b = -A(P - A^{-1})^{\mathsf{T}} b.$$

The pad discrepancy and plate row bound Q therefore imply

$$|A\mu_{\mathrm{count}} + b_i| \le Q\,\mathrm{distPay}(B) b_0 t.$$

Choose $C_m ean = 1 + \max(b_0, b_1, \dim(Color) C_i nv b_0, 0)$. It simultaneously bounds priced(0), the linear bill and the mean by $C_m ean$ t. The residual envelope above has any prescribed rate k: scraps(B), and hence distPay(B), has rate k+1/2, while t has rate $-1/2$. This supplies the mean/residual hypotheses of the correction budgets. The clean base-stroke and uniform maximum estimates proved here give A_0 and M.

6.51 Exact Gaussian exponent rearrangement

Let D be the form matrix, $S_s leep$ the dormant diagonal, $A = D^{\mathsf{T}} D + S_s leep$, μ the count-coordinate mean, b the linear bill, and C_{ref} the positive reference covariance. The lifted tangent input is μ+gz, with $g = \beta^{-1}/2$. The affine-square sum has the exact polarization identity

$$\begin{aligned}\mathrm{liftSq}(\mu + gz) &= \mathrm{liftSq}(\mu)\\ &+ g\langle z, b + D^{\mathsf{T}} D\mu\rangle + g^2 \langle z, D^{\mathsf{T}} Dz\rangle/2.\end{aligned}$$

This follows by expanding each squared affine row and summing the finite color coordinates. The dormant mean vanishes, so $S_s leep\ \mu = 0$. Define $ghost(z) = \langle z, S_s leep\ z\rangle/2$. Using $\beta g^2 = 1$ and $\beta g = g^{-1}$ gives

$$\begin{aligned}&-\beta\,\mathrm{liftSq}(\mu + gz) - ghost(z)\\ &= -\beta\,\mathrm{liftSq}(\mu) + \langle z, hatted_\beta\rangle\\ &+ \langle z, C_{\mathrm{ref}}^{-1} - Az\rangle/2 - \langle z, C_{\mathrm{ref}}^{-1} z\rangle/2.\end{aligned}$$

The last term is the reference Gaussian exponent. The preceding quadratic term is the correction expanded by the itinerary chain; the linear term is the first correction. The inserted ghost supplies exactly the dormant diagonal needed to complete A, and its normalized dormant Gaussian integral is one as shown earlier. This equality is pointwise algebra and remains valid where a later profile vanishes; it assumes no sector event.

6.52 Pointwise Gaussian chain corrections

Write A for the completed plate, C_{ref} for the symmetric reference covariance, and $S_{\mathrm{err}} = A\,C_{\mathrm{ref}} - I$ for its small discrepancy. The linear correction vector is

$$hatted_{\beta} = -gStep(\beta)^{-1}(Adrift + b).$$

For each fine link e its scalar correction is $\sum_c olor\ z_{(e,j)}$ $\mathrm{hatted}_{(e,j)}$. Multiply this by aperture(e), the cone profile at twice the chosen enforcement radius on hot links and the value 1 on other links, to obtain qfirst.

For an itinerary of n+2 link steps, the quadratic correction is half the bilinear pairing between its starting and ending z vectors with the ordered matrix-block product

$$-S_{\mathrm{err}} \cdots -S_{\mathrm{err}} A,$$

containing n+1 discrepancy factors and one terminal A. Multiply it by the two endpoint apertures to obtain chain. Summing itineraries at fixed starting link and n restores exactly the matrix product $(-S_{\mathrm{err}})^{n+1}A$ in that quadratic form: the intermediate link and color sums are precisely matrix multiplication. The terminal A is retained in both the formula and its price.

Dormant plate rows are identity rows and the reference covariance has the same identity completion. Therefore S_{err} has zero entries whenever either endpoint is dormant. The linear bill and mean vanish there too, so hatted is zero at dormant coordinates. A nonzero quadratic chain consequently has hot starting and ending links: every nonzero discrepancy factor forces hot endpoints, and the final A cannot join a hot link to a distinct dormant identity row. Induction over its discrepancy factors proves this for every n.

Let $\|\mu_e\| \le K$ t and $g = \mathrm{gStep}(\beta) > 0$. At a hot endpoint where its aperture is nonzero, the cone bounds $\|\mu_e + g\ (atmos(z))_e\| \le 4a$ t. The triangle inequality gives

$$\|atmos(z)_e\| \le g^{-1}t(4a + K) = heat(a, K, t, \beta).$$

Thus an envelope $U \ge heat$ bounds every endpoint color coordinate. The aperture is at most one. If $|\,\mathrm{hatted}_{(e,j)}\,| \le M$, the finite color sum yields

$$|qfirst_e| \le \dim(Color)UM.$$

This holds at dormant links as well because their linear correction vanishes.

For the chain, put $R_c = \dim(Color)$, $d_{\mathrm{pay}} = \mathrm{distPay}(\eta, r, Q, m, B)$, and $M_{\mathrm{err}} = Q\ d_{\mathrm{pay}}$. The discrepancy has entry bound M_{err}, whereas the terminal plate has entry bound Q. Summing its n+1 intermediate color indices bounds each entry of the prescribed block product by

$$R_c^{n+2} Q M_{\mathrm{err}}^{n+1}.$$

The convenient extra color factor is included in the source's product estimate. The two endpoint color sums therefore bound the half-bilinear form by

$$\begin{aligned} &\tfrac{1}{2}U^2 R_c^{n+4} Q M_{\mathrm{err}}^{n+1} \\ &= centreQ(U) \cdot etaS^n, \\ centreQ(U) &= \tfrac{1}{2}U^2 R_c^4 Q(Qd_{\mathrm{pay}}), \\ etaS &= R_c Q d_{\mathrm{pay}}. \end{aligned}$$

A nonzero term must also be armed: every itinerary hop has parent distance at most 3B. Otherwise its corresponding discrepancy or plate block vanishes by the finite-range estimates. Define capQ to equal the displayed center-times-ratio power on armed itineraries and zero on the others. Splitting into zero aperture, zero bilinear form, and nonzero cases proves $|chain| \le capQ$ pointwise for every z. This explains why these optional terms use endpoint cutoff bounds rather than a separate Gaussian tail event.

6.53 Convergent inverse and logarithm expansions

Weighted absolute row and column prices multiply under matrix multiplication. If a matrix E has price $q < 1$ at a nonnegative width, E^n has price q^n, so its entries form absolutely convergent geometric series. The finite identity $\sum_{n<N} E^n(I-E) = I-E^N$ and $E^N \to 0$ give $\sum E^n = (I-E)^{-1}$, proving invertibility at the same time.

There are two different orders of multiplication needed here. If $AP - I = E$, then

$$A^{-1} = \sum_{n\ge 0} P(-E)^n.$$

If $AC - I = E$, then

$$C^{-1} = \sum_{n\ge 0} (-E)^n A,$$
$$C^{-1} - A = \sum_{n\ge 0} (-E)^{n+1} A.$$

These follow by multiplying the geometric inverse of I+E on the appropriate side. The order is essential because the matrices need not commute. Absolute entry convergence permits the fixed matrix multiplications and the removal of the zeroth term.

For a real matrix A whose interpolation $G_s = I + s(A - I)$ remains invertible on [0,1], the elementary determinant derivative gives

$$d/ds \log|\det G_s| = trace((A - I)G_s^{-1}).$$

To derive it at $s = r$, factor $G_s = [I+(s-r)(A-I)G_r^{-1}]G_r$. The coefficient of t in $\det(I+tB)$ is trace B, as follows directly from the determinant polynomial. Differentiate the factored determinant and then $\log|\cdot|$. The inverse is continuous along this compact interpolation by the adjugate formula and nonvanishing determinant. Integrating therefore gives

$$\log|\det A| = \int_0^1 trace((A - I)G_s^{-1})ds.$$

For the positive plates used here $\det A > 0$, so this is $\log\det A$.

Let P_s be its inverse-column pad and $E_s = G_s\ P_s - I$. The uniform discrepancy price $\rho = \text{scraps}(B) < 1$ gives

$$(A - I)G_s^{-1} = \sum_{n\ge 0} (A - I)P_s(-E_s)^n.$$

The pad P_s occurs once in each term; only the residual $-E_s$ is raised to the power n. This follows from $G_sP_s = I + E_s$, hence $G_s^{-1} = P_s(I + E_s)^{-1}$.

Each diagonal summand is continuous in s and bounded in absolute value by $(Q + 1)payE \cdot \rho^n$. The geometric bound is summable independently of s. Hence its interval integrals may be summed:

the sum of their integral norms is bounded by the same geometric series on an interval of length one. Write $lact_i(n)$ for these diagonal integrals and $\mathrm{localBase}_i = \mathrm{lact}_i(0)$. We obtain

$$\log\det A = \sum_i localBase_i + \sum_i \sum_{n\geq 0} lact_i(n+1).$$

Thus the local principal-window diagnostic is only the zeroth term; the remaining determinant arms restore the full inverse exactly.

For any discrepancy E of price $q < 1$, apply the same argument to I+sE and expand its inverse as $\sum_n (-sE)^n$. The integrands $E(-sE)^n$ have uniform bound q^{n+1}, giving

$$\log|\det(I+E)| = \sum_{n\geq 0}\sum_i \int_0^1 [E(-sE)^n]_{(i,i)}\,ds.$$

In particular $E = A\,C_{\mathrm{ref}} - I$ yields $I + E = A\,C_{\mathrm{ref}}$. Both A and C_{ref} have positive determinants, so the right side equals $\log\det A + \log\det C_{\mathrm{ref}}$. These source-arm terms therefore represent the actual covariance normalization defect. All exchanges here use uniform summable entry bounds, followed by finite coordinate sums; no volume factor is treated as a local price.

6.54 Itinerary receipts and scalar arm identities

An itinerary of k hops is a list of k successor fine links, with the initial link retained separately. A hop is armed at radius d when some coarse endpoint of the first link is within d of some coarse endpoint of the next. Each parent is within one site of either endpoint, so an armed hop moves parents by at most d+2.

The town of a link is the mB halo of its two coarse endpoints. It contains its parent and lies in the radius-mB+1 ball about that parent, hence in the radius-(m+1)B ball when $B > 0$. It can therefore be covered at scale B with at most $\mathrm{steady}(m) = \mathrm{ballCost}(m+1)$ tiles. The visited receipt is the union of the towns of every link in the itinerary, including both its initial and final links.

If $d+2 \leq mB$, each successor parent belongs to the previous town and to the remaining visited receipt. These two sets meet. The covering inequality for meeting supports and induction give $\mathrm{coverNumber}_B(receipt) \leq (k+1)\,\mathrm{steady}(m)$. For its tenure bound, any point in the first town lies within mB+1; a point in the rest is within d+2 of its next parent plus the induction bound there. Thus, when $d+2 \leq mB+1$, every receipt point lies within (k+1)(mB+1) of the initial parent. These establish the two geometric inputs used in the address count.

For a sequence of matrices $P_s(0), P_s(1), \ldots$ and a fixed itinerary, multiply their color blocks in its link order. Keep only itineraries returning to their initial link, take the color trace of that block product, and integrate s from 0 to 1. Call this scalar its stamp. Expanding matrix multiplication shows that summing stamps over all k-hop itineraries restores the integrated diagonal trace of $P_s(0)\cdots P_s(k-1)$. There are finitely many itineraries at fixed k; continuity of the matrix factors gives integrability and permits interchanging these finite sums and the integral.

For the determinant arm choose $P_s(0) = A - I$, $P_s(1) = \mathrm{pad}(I + s(A-I))$, and $P_s(j) = -\,\mathrm{error}(I + s(A-I))$ for $j \geq 2$. Its order-n label has n+3 hops, so its sum restores the local integrated intensity at grade n+1, retaining n+1 discrepancy factors. For the source arm choose $P_s(0) = flawed(A)$ and $P_s(j) = -s\ flawed(A)$ for $j \geq 1$. Its order-n label has n+1 hops and restores the corresponding integrated grade-n note. These are exact scalar identities; the finite color trace belongs inside each arm label.

If each factor has entry bound l_j, the absolute stamp is at most $R_c^{k+1}\prod_{j<k} l_j$. This follows by the finite block-product entry bound, the final color sum, and the interval length one. If the

itinerary is not armed at the common factor range, one factor block is zero, so its stamp is zero. This proves the arm representations and the pointwise pricing mechanism used above.

6.55 Counting all itinerary addresses

The order-n itinerary price is $HC(QR_{mprice})^n$. Both the chain coefficient bound Q^n and the route/support price R^n_{mprice} arise at every step. The extraction's unparenthesized $HCQR^n_{mprice}$ is interpreted accordingly, consistently with the repeated chain estimate later in the same construction.

An order-n ledger with u extra steps consists of a starting fine link and an itinerary of n+u hops. Its receipt is the prescribed union of enlarged neighborhoods visited along that itinerary. Put density$(L) = L^4|Axis|$ and abundance$(L, d) =$ ballCost$(d + 2)$ density(L). From any starting link an armed hop of radius d has at most abundance(L, d) possible next links: every such link has parent in the radius-d+2 neighborhood, and the fine-link caption census has density(L). Induction over the hops therefore bounds the number of armed k-hop itineraries by abundance$(L, d)^k$.

The receipt geometry gives two bounds when $d + 2 \le mB$ and $B > 0$. Its B-covering number is at most (k+1)steady(m), and if a fixed site x belongs to its receipt, its starting link has parent within distance (k+1)(mB+1) of x. The underlying definitions of visited receipts give these covering and tenure bounds; they remain the immediate geometric inputs here. Their effect on the count is as follows.

For coefficients $0 \le w_{(n,e,p)} \le C\ Q^n$ on armed addresses and zero elsewhere, the tariff at order n includes $\exp(2\varepsilon n)$ and $\exp(A_s upport\,$coverNumber$)$. Bound the latter by $\exp(A_s upport(n + u + 1)\,\text{steady}(m))$. The number of eligible starting links is at most

$$\text{ballCost}((n + u + 1)(mB + 1))\,\text{density}(L).$$

The elementary inequality $n + u + 1 \le (u + 1)2^n$ and the four-dimensional ball formula give

$$\text{ballCost}((n + u + 1)g) \le \text{ballCost}((u + 1)g)16^n.$$

Multiplying this starting-link count, the armed-hop count, and the coefficient/tariff bounds yields, at each x, an order-n total at most

$$\begin{aligned} &HC(QR_{\text{price}})^n, \\ H &= \text{ballCost}((u + 1)(mB + 1))\,\text{density}(L) \\ &\cdot \text{abundance}(L, d)^u \exp(A_{\text{support}}(u + 1)\,\text{steady}(m)), \\ R_{\text{price}} &= 16\,\text{abundance}(L, d) \exp(A_{\text{support}}\,\text{steady}(m) + 2\varepsilon). \end{aligned}$$

When Q $R_{\text{price}} \le 1/2$, its sum over n is at most 2HC. All terms are nonnegative. To pass from the finite order-n sums to the countable ledger sum, first bound every finite subset within its finitely many order fibers, then apply the convergent geometric bound. This proves the asserted sales inequality and its summability, retaining every address even when different itineraries revisit the same link.

6.56 Pointwise determinant and source arm prices

For the determinant arm the ordered factors along its n+3 hops have respective entry prices Q+1, payE, and n+1 copies of scraps(B). Each factor has range at most 3B: the plate term uses its B/2 stencil, the inverse-column pad its window, and its discrepancy the enlarged window stencil. The

common interpolating parameter $s \in [0,1]$ preserves these bounds and their positive floors. The product entry estimate, followed by the color trace and s integral, gives

$$\begin{aligned} |ArmA_{(n,e,p)}| &\le R_c^{n+4}(Q+1)payE \cdot \operatorname{scraps}(B)^{n+1} \\ &= centreA \cdot etaA^n \end{aligned}$$

on armed itineraries, and zero otherwise. Here $centreA = R_c^4(Q+1)payE \cdot \operatorname{scraps}(B)$, $etaA = R_c \cdot \operatorname{scraps}(B)$, and R_c is the color dimension.

For the source arm all n+1 factors have entry price Q·distPay(B); the factors beyond the first also contain the scalar interpolation multiplier, whose absolute value is at most one. Their range remains 3B. Thus

$$\begin{aligned} |ArmS_{(n,e,p)}| &\le R_c^{n+2}(Q \operatorname{distPay}(B))^{n+1} \\ &= centreS \cdot etaS^n \end{aligned}$$

with $centreS = R_c^2 Q$ distPay(B) and $etaS = R_c$ Q distPay(B). These bounds are the actual scalar arm prices before exponential expansion. They use the ordered-product and trace-integral representation of the arms given in the itinerary section; the pointwise estimate sums their finite internal color indices and integrates over an interval of length one.

6.57 From itinerary envelopes to support budgets

The general itinerary bound takes an armed-path coefficient of the form $Y(x)\kappa(x)^n$, supported on its predetermined receipt of halo $i \ge 4$ and fixed number u of extra steps. Let abundance($L, 3B$) bound the number of possible next links in an armed hop. Both this abundance and the fixed-step prefactor $expoHead(L, i, B, 3B, u, A_support)$ have rate zero. The geometric receipt estimate gives the sales bound

$$2expoHead \cdot Y$$

as soon as

$$16 \operatorname{abundance}(L, 3B) \exp(A_{\text{support}} \cdot \operatorname{steady}(i) + 2\varepsilon)\kappa \le 1/2.$$

If κ has any negative rate, that condition holds eventually; multiplying it by the rate-zero abundance does not change its decay. Thus an envelope Y of rate b yields a sales bound of rate b. This numerical deduction is uniform in geometry; the itinerary address count and halo-cover estimate supply the quantities expoHead and steady used here.

Apply it to the determinant and source arms separately. The determinant arm uses $u = 3$, center $R_c^4(Q+1)payE \cdot \operatorname{scraps}(B)$ and ratio $R_c \cdot \operatorname{scraps}(B)$. The source arm uses $u = 1$, center $R_c^2 Q \cdot \operatorname{distPay}(B)$ and ratio R_c Q·distPay(B). All centers and ratios have every prescribed negative rate, so their sales budgets can have any prescribed rate b. Their pointwise bounds retain one scalar per actual arm, rather than splitting a trace into independently labeled colors. Applying the same bounds to the blank geometry supplies its determinant arm independently. The pointwise arm identities and the geometric itinerary estimate are still the inputs to this budget calculation.

6.58 Numerical budgets for the Gaussian chain letters

The mean/residual hypotheses are retained separately. They assert $\|\mu_e\| \le C \cdot t$ and $|(A \operatorname{drift} + b)_i| \le E(x)$, together with the shell access bound. For the physical beta band $0 < \beta \le 2$ betaF put

$$\mathrm{burst}(x) = \sqrt{2\,\mathrm{betaF}(x)}.$$

Then $\mathrm{gStep}(\beta)^{-1} = \sqrt{\beta} \le \mathrm{burst}$, and burst has rate 1/2. The bounded cutoff field used by the chain estimates therefore has common envelope

$$U(x) = burst(x) \cdot \mathrm{fiery}(x) \cdot muRate(C, H),$$

of rate zero. This envelope is independent of β and the Gaussian variable. Its nonnegativity and the mean bound supply the hypotheses of the pointwise chain estimate.

If E has rate k$-1/2$, then $b(x) = burst(x)|E(x)|$ has rate k. The first mean-correction envelope is

$$J(x) = \dim(Color)U(x)b(x),$$

of rate k. Its caption is the parent of its fine bond. With support town contained in radius-(m+1)B and fine-link caption census $L^4|Axis|$, its sales price is at most

$$\begin{aligned}&\mathrm{ballCost}((m+1)B)L^4|Axis| \cdot J(x)\\ &\cdot \exp(A_{\mathrm{support}}\,\mathrm{ballCost}(m+1)),\end{aligned}$$

also of rate k.

The higher quadratic-chain envelopes are capQ evaluated at U. Their series sales bound follows from the itinerary estimate with center coefficient proportional to

$$\tfrac{1}{2}U^2 flagP^4Q^2payE \cdot \mathrm{scraps}(B)$$

and ratio coefficient flagP·Q·payE·scraps(B). The latter has every prescribed negative rate; the former has rate k for any prescribed k because U has rate zero. The itinerary summation consequently gives the chain sales bound of rate k. The support charge is its series index, and halo $m \ge 4$ is retained. The pointwise Gaussian chain and short-correction estimates supply capQ and J; this calculation retains their coupling and residual factors in the common budget.

For a finite caption family of density n with supports contained in radius-dB balls, a constant coefficient envelope C_0 has sales price at most ballCost(dB)n C_0 $\exp(A_s upport\,\mathrm{ballCost}(d))$. Indeed a fixed site belongs to at most ballCost(dB)n supports, and each support has B-covering number at most ballCost(d). This is the fixed-envelope counting used for the first correction and local Short letters.

6.59 True path supports and bundle geometry

The native support of a walk is the finite set of its visited sites, including its start and endpoint. Every visited site lies within its actual walk length of the start, by induction and the one-step separation bound. A walk of length at most M is therefore covered by one radius-M tile. For a longer walk, split off its first M steps. The two native supports meet at the split endpoint. The meeting-support covering inequality and induction give

$$M\,\mathrm{coverNumber}_M(native(w)) \le M + \mathrm{length}(w), M > 0.$$

This uses the walk length with all repetitions and winding, even when the finite native set is much smaller.

A curvature slot's actual support is the native support of its transport walk followed by the four-step plaquette. This equality follows by concatenating visit lists and wrapping their cover coordinates. Its slot length is exactly that concatenated length. A bundle of d slots rooted at x is the union of their supports with x inserted. All its supports meet at x, so the same covering inequality gives

$$M\,\mathrm{coverNumber}_M(\mathit{bundle}) \le (1+d)M + \mathit{totalLength}.$$

Every bundle site also lies within totalLength of x, since each individual slot length is at most their nonnegative sum.

For a new support Y contained in $\mathrm{halo}_l\, B(\mathrm{parent}(\mathit{bundle}))$, the parent-distance inequality gives its extent from parent(x) at most (lB+1)(1+totalLength). Apply the coarse footprint bound with $M = B$ and the last bundle covering inequality. Define

$$\begin{aligned} \mathit{baseFee} &= \mathrm{preCost}(l) + \mathit{conversion}(L,l)(1+d),\\ \mathit{legFee} &= \mathit{conversion}(L,l)/B,\\ \mathit{originFee} &= lB+1. \end{aligned}$$

Then $\mathrm{coverNumber}_B(Y) \le \mathit{baseFee} + \mathit{legFee}\cdot\mathit{totalLength}$. The anchored-to-hit principle applies with degree equal to true totalLength, extent originFee, and those two covering coefficients. At each fixed fine root, its input budget is exactly the original slot-family length budget; the root delta and this budget form a product budget, so no volume factor is inserted beyond the parent population L^4. This proves $\mathrm{bundle}_h\, it$, including the tail $\exp(-\varepsilon m)$ for $\mathit{totalLength} \ge m$ and its explicit payment. Atom footprints are exactly bundles of their slot legs; an ordered path pair is the two-slot bundle and has the sum of its two actual lengths.

6.60 Connected center traversal and halo tiling

The covering currency uses the periodic maximum-coordinate separation. A center set is connected at range b if any two centers can be joined through centers with successive separation at most b. A cover at radius M covers the support by the M halos of its centers and requires those centers to be connected at range 2M+1. Thus disconnected pieces of a support still incur the cost of connecting their cover; it is not an unrestricted ball-counting number.

There is a unit-speed line from x to y of duration separation(x,y): in each centered-lift coordinate of y−x, truncate the displacement to the interval [−j,j] at time j, then wrap and add x. Each coordinate changes by at most one per time step, and it stops at y once j reaches their maximum absolute displacement.

For a finite center set D connected at range b, build a traversal by adjoining one unvisited center at a time. Connectivity supplies an edge from the visited subset to its complement: take the first exit along any center path to an unvisited point. At the occurrence of its visited endpoint in the current traversal, insert the line to the new center and its reverse. The detour has length at most 2b and preserves all old visits. Starting with one center and repeating gives a unit-speed traversal visiting every center with length at most 2b|D|. The empty case has the zero-length trivial traversal. This proves the ring-path bound used in extent and collapse.

For tiling a radius-kB halo of one site, take the image of the integer grid $[-k,k]^4$ scaled by B and translated to that site. Divide each centered coordinate of a halo point by B. Its quotient lies in that grid and its remainder has absolute value at most B, so one grid center covers it. The grid centers are connected through grid steps; wrapping may identify centers and can only reduce their number. Hence at most $(2k+1)^4$ centers suffice.

If X has a radius-mB cover with center set D, enlarge its balls to radius (m+1)B. Their original connecting steps are at most $2mB + 1 \le 2(m + 1)B$, so these enlarged halos meet along the center connections. Tile each enlarged ball by the preceding grid and join their connected covers at the meetings. This gives

$$\mathrm{coverNumber}_B(X) \le |D|[2(m + 1) + 1]^4.$$

The same line construction proves that halos of centers connected at range 2M+1 are connected at unit range: the line between two centers lies in their two M halos, with the possible middle adjacent step included. These provide the connected-halo and tiling facts used by the collapse proof without requiring any injectivity of a periodic halo.

6.61 Coarse collapse of a connected covering

Write L for the block side. A fine site is within L/2 of its embedded parent. Embedded coarse-site separation is at least L times coarse separation, coordinatewise from centered lifts. The triangle inequality therefore gives

$$L\mathit{separation}(\mathrm{parent}(x), \mathrm{parent}(y)) \le \mathit{separation}(x, y) + L.$$

For a unit-step fine-site path $R_0, \dots, R_N$, sample its parents every LB steps, $B > 0$. Consecutive sampled parents are within 2B by this inequality, and every parent along the path is within 2B of one sample. The set D of samples is connected at that range and satisfies

$$LB|D| \le N + LB.$$

This includes the final partial segment and allows repeated samples.

Realize the scale-M cover of a support X by a center set A of size $\mathrm{coverNumber}_M(X)$ whose union of M-balls is connected. Join two centers when their distance is at most 2M+1. This center graph is connected: a lattice path through the union passes from one ball to another only through overlapping or adjacent balls. A finite connected graph has a spanning tree, obtained by adding an edge from the already reached vertices to a new vertex until all vertices are reached. Traverse this tree out and back. Each tree edge is replaced by a lattice path of length at most 2M+1, and each is traversed twice, so the resulting unit-step path visits every center and has length at most $2(|A| - 1)(2M + 1) \le |A|(4M + 2)$. The empty cover has the constant path and length zero. Join any two sites of X to covering centers by paths of length at most M. Their distance is at most |A|(4M+2)+2M. This proves the extent estimate, on the periodic lattice as well as on the integer cover.

Sample the center traversal at scale LB as just described. If $M \le 2B$, the parent of any support point is within 3B of its center's parent, and that center's parent is within 2B of D. Hence $\mathrm{halo}_{(sB)}(\mathrm{parent}(X))$ is covered by the radius-(s+5)B halo of D, with its inherited connectedness. Tiling that halo at scale B costs at most |D|ballCost(s+6). Combining the sample count, $N \le |A|(4M + 2)$, $M \le 2B$ and $B \ge 1$ gives

$$\begin{aligned} &L\,\mathrm{coverNumber}_B(\mathrm{halo}_{(sB)}(\mathrm{parent}(X))) \\ &\le [L + 10\,\mathrm{coverNumber}_M(X)]\,\mathrm{ballCost}(s + 6). \end{aligned}$$

For X empty both sides are immediate; otherwise its realized cover has at least one center, which pays the fixed endpoint term in this inequality. Dividing by L yields exactly the preCost and conversion constants used in the coarse hit budget. Thus all spatial block factors in that budget follow from parent separation, a connected cover traversal, and halo tiling.

6.62 Anchored budgets converted to hits

The geometric counting principle is the following. Suppose each fine anchor has budget $\sum_{o_i=y} b_i e^{Ad_i} \le \gamma$, the parent map has at most P fine sites per coarse site, and

$$x \in X_i \Rightarrow \mathit{separation}(x, \mathrm{parent}(o_i)) \le W(1+d_i),$$
$$\mathrm{coverNumber}_B(X_i) \le v + v' d_i.$$

Let $0 \le g_i \le C b_i (1+d_i)^k$, and assume $a \ge 0$, $\varepsilon > 0$ and $av' + 2\varepsilon \le A$. At fixed degree $d_i = j$, the coefficient with its new support weight is bounded by

$$b_i \exp(Aj) \cdot C(1+j)^k \exp(av - 2\varepsilon j).$$

All relevant anchors lie above the coarse radius-W(1+j) ball about x. Their number is at most P ballCost$(W)(1+j)^4$. Group a finite set of labels by anchor and apply the original anchored budget at each anchor. The resulting degree-j total is at most

$$\exp(av)\gamma CP\,\mathrm{ballCost}(W)(1+j)^{4+k}\exp(-2\varepsilon j).$$

The exponential series gives $(1+j)^{4+k} \le \varepsilon^{-(4+k)}(4+k)! e^{\varepsilon} e^{\varepsilon j}$. Summing $j \ge m$ proves a hit budget with tail factor $\exp(-\varepsilon m)$ and payment

$$\exp(av)\gamma CP\,\mathrm{ballCost}(W)$$
$$\cdot \varepsilon^{-(4+k)}(4+k)!\exp(\varepsilon)/(1-\exp(-\varepsilon)).$$

The argument bounds every finite label subset, so it proves the nonnegative countable summability as well.

For a block parent of side L, the site decomposition into parent and residue gives $P = L^4$. For an old support Y at scale $M \le 2B$, the coarse collapse estimate gives

$$\mathrm{coverNumber}_B(\mathrm{halo}_{\mathrm{lB}}(\mathrm{parent}(Y)))$$
$$\le \mathrm{preCost}(l) + 10\,\mathrm{preCost}(l)\,\mathrm{coverNumber}_M(Y)/L,$$
$$\mathrm{preCost}(l) = \mathrm{ballCost}(l+6).$$

Its extent estimate gives separation from the parent of any chosen anchor at most (l+20)$B(1+$ coverNumber$_M(Y)$). These two bounds supply v,v′ and W in the general principle. The connected traversal and coarse-collapse arguments above prove these covering and extent bounds; the displayed summation then supplies the hit budget. Inspecting a support adds at most four fine-site steps, hence at most five parent steps, so its enlarged support is included after increasing l by five.

The reserve inequalities have the required quantifier order. For old exponent $A > 2\varepsilon$, choose L strictly above 10 preCost(l)a/(A-2ε). Then $10a\,\mathrm{preCost}(l)/L + 2\varepsilon \le A$, before choosing x. For path supports the coefficient is instead 10preCost(l)/(LB); at fixed L it tends to zero as $B \to \infty$, so any $\sigma > 2\varepsilon$ suffices eventually. The same delay ensures the old fine radius radiusB(betaF) $\le$ 2 $\mathit{radiusB}(x)$. Fixed powers of these radii and the monitor are subpower in betaF, whereas $\exp(-c\,\mathrm{till}(x))$ and $\exp(-\mathit{cmonitor}^z)$, $c, z > 0$, beat every inverse power. These numerical facts justify the strict-exponent absorptions used above, without changing L after the delay.

6.63 Support tariffs and path-tail conversion

For nonnegative coefficient envelopes w_i, supports X_i and integer extra charges e_i, the sales condition means that at every coarse site x,

$$\sum_{i:x\in X_i} w_i \exp(2\varepsilon e_i)$$
$$\exp(a\,\mathrm{coverNumber}_B(X_i)) \le \gamma,$$

with the corresponding nonnegative series summable. Disjoint alphabets add their γ bounds. A finite additional orientation index costs the sum of the orientation bounds; rotating or translating supports preserves a bound since the site map is bijective and does not increase their covering number. If X_i is nonempty and $a, \varepsilon \ge 0$, each $w_i \le \gamma$ follows by reading the inequality at one site in X_i.

To transfer an old anchored bound at scale M, reserve its extra tariff first: set $v_i = w_i \exp(2\varepsilon\,\mathrm{coverNumber}_M(Y_i))$. Then

$$v_i \exp((\delta - 2\varepsilon)\,\mathrm{coverNumber}_M(Y_i))$$
$$= w_i \exp(\delta\,\mathrm{coverNumber}_M(Y_i)).$$

The old anchored bound is therefore unchanged after reducing its support exponent by 2ε. If the new X_i lies in the l-halo enlargement of the parent image of Y_i, $M \le 2B$, and the geometric reserve inequality holds, the coarse support estimate applies to v. It allows a new coefficient $z_i \le C w_i (1 + \mathrm{cov}_M(Y_i))^k$ and pays its new 2ε tariff. Its resulting budget is the explicit footprint payment with fine-child factor L^4, extentCost(l, B), degree k, original budget γ, support exponent a and geometric preCost(l). The anchor-to-hit summation and the geometric coarsening proof are given in (NF9)–(NF10) below, with these parameters.

For rooted path bundles the same transfer starts from a true length budget $\sum_i w_i \exp(\delta\,\mathrm{length}_i) \le \gamma$. Set $v_i = w_i \exp(2\varepsilon\,\mathrm{length}_i)$, leaving width $\delta - 2\varepsilon$ for the geometric conversion. The bundle support estimate then supplies sales for any $z_{(\mathrm{root},i)} \le C w_i (1 + \mathrm{length}_i)^k$ on an allowed enlarged bundle footprint. Unlike an anchored activity family, its root varies over every fine site; the payment's originFee and L^4 factors account for that census.

An inspected order-r atom footprint lies inside the (l+5)-halo enlargement of its actual path bundle. Applying the path transfer to its magnitude and exponential length budget gives the atom payment. Its degree-k polynomial mark is retained in the payment. For paired tails the mask halo adds only two further coarse halos: enlargement$_l(pairFoot)$ is contained in the $(l+2)$-halo enlargement of the two-path bundle, by the parent-distance inequality. This supplies the tail payment with $r = 2$.

The tail's small factor can be computed before support conversion. Write x,y for the two actual slot lengths. A nonzero omitted pair has x $d > R$ or y $d > R$, so $(x + y)d \ge R$. The pair-price bound and half of the exponential length reserve give

$$pairPrice \cdot \exp((\sigma/2)(x + y))$$
$$\le 2\exp(-\sigma R/(2d))$$
$$|u_x| \exp(\sigma x) |u_y| \exp(\sigma y).$$

Sum using the product of the two original path budgets, each at most $C_p aths$. Thus the half-width budget is at most $2\exp(-\sigma R/(2d)) C_p aths^2$. At $R = \mathrm{till}(x)$ this has every prescribed negative rate. The remaining width $\sigma/2 - 2\varepsilon$ is positive under $8\varepsilon < \sigma$, leaving the geometric reserve needed for support conversion and the polynomial mark.

All footprint payments multiply their scalar envelope and original budget by fixed constants and a four-dimensional extent factor. If those envelopes have rates u,v and the extent has rate zero, payment has rate u+v. Consequently any strictly larger exponent absorbs the remaining fixed and subpower factors. This proves the atom and tail rate deductions used by the inherited six-family assembly, conditional only on the geometric footprint conversion.

6.64 Inherited coefficients and their enlarged supports

Fix halo l and positive ε with 4ε below the incoming support exponent, and 8ε below both the fine and coarse path widths. A sufficiently large block length L converts the incoming anchored support budget to the enlarged coarse footprints, leaving the prescribed support price A_{support} and extra-charge tariff. The same L works for the regular and forced activity classes. This conversion uses the original incoming envelopes, before any new marking reserve is constructed.

For a regular occurrence, take its supplied nonnegative envelope b_i. On its healthy support its absolute value is bounded by b_i using the zeroth-order incoming bound; outside that condition the masked regular occurrence is zero. Its anchor remains the old regular anchor. For a forced coefficient use its incoming envelope d_i without a health restriction. Choose an inventory point in each forced support as anchor; the incoming hit bound implies the weaker anchored bound at that point. Apply the same coarse support conversion separately to each family. Its geometric payment factor has rate zero, because its extent is a fixed multiple of B and L is fixed. The regular incoming envelope has rate $-11/10$, so any price $betaF^u$ with $u > -11/10$ is available eventually. The forced incoming bound can be chosen with rate v-1, yielding a sales price $betaF^v$ for any fixed v.

For an order-k path atom multiplied by a scalar η with $|\eta| \le D(x)$, let D have rate u and its total atom budget have rate v. The healthy atom bound gives the pointwise envelope

$$z_{(site,j)} = D(x)\,\mathrm{fiery}(x)^k magnitude(atom_j).$$

Its rate before support conversion is u$-$k/2+v. If this is strictly below a desired exponent r, choose an intermediate exponent s between them, convert the inspected supports, sum over the finite signed orientations, and absorb that fixed number into $betaF^r$. This retains the atom's actual path length in its extra charge.

The omitted path-pair tails have a stronger property: their exponential length budget pays any prescribed power $betaF^r$ after a fixed scalar multiplier D of any fixed rate. The support conversion uses $8\varepsilon < width$ and the actual cutoff till(x), not endpoint separation. Sum the finitely many plaquette types and signed orientations at exponent r-1 and absorb their fixed count into exponent r. The detailed atom and tail support-conversion estimates are still the inputs to these two deductions.

The inherited logarithmic exponent has six disjoint families: regular occurrences, cubic atoms, quartic atoms, loop atoms, fine path tails, and coarse path tails. Its scalar phase is $\theta = -orientation \cdot \beta$. The beta band and $|orientation| \le 1$ imply $|\theta| \le 2$ betaF; the coarse tail uses $-\theta$ and obeys the same bound. The cubic and quartic multipliers therefore have rate 1, giving exponents $1 - 3/2 = -1/2$ and $1 - 4/2 = -1$. The loop multiplier has rate zero and uses its own incoming loop budget. The two tail families allow arbitrary negative powers.

Given $r > -1/2$, choose $-1/2 < s < r$. All six families have sales bounds $betaF^s$ eventually, with the incoming regular bound weakened to s. Their disjoint-union price is at most 6 $betaF^s \le betaF^r$ eventually. The forced family retains its independent $betaF^{r'}$ price for any r′. This proves the inherited dictionary bound with the actual fine and coarse beta factors already included. Every footprint remains nonempty: regular and forced supports retain their anchor or inventory, atom

footprints retain the rooted path origin, and tail-pair footprints retain their designated root. These facts support the individual-letter smallness argument used in the optional dictionary assembly.

6.65 The non-warning primitive dictionary

Before taking connected components, keep three disjoint kinds of logarithmic letter: inherited fine-scale log letters, local short-correction letters, and Gaussian chain letters. The last kind has five families: the negative actual log-determinant series, the positive blank-reference log-determinant series, the mean-source series, the first linear mean correction, and the remaining quadratic chain correction. In the integrand their signs and factors are respectively $-\mathrm{ArmA}/2$, $+\mathrm{ArmA}_{\mathrm{blank}}/2$, $+\mathrm{ArmS}/2$, qfirst and chain. These signs belong to the exponent itself.

For an inherited log letter take its actual old exponent evaluated on the lifted fine field and multiply it by the enforcement profile of its own full footprint at parameter 2a. For an inherited forced activity use its actual K coefficient with the same footprint enforcement. These profiles lie in [0,1], so they cannot increase the supplied absolute bound. The full exponent farLog is the disjoint sum of these inherited exponents, the Short correction, and the five chain families. Each optional activity is $\exp(\mathrm{farLog})-1$.

Supports are assigned before evaluating any coefficient: inherited letters use their old support enlarged by the prescribed halo; short letters use their sleeves; the chain series use their receipt supports, and the first correction uses its town support. Their extra charges are the old extra charges, zero for short letters and first corrections, and the series index for the other chain letters. Every support is nonempty, including coefficients which happen to vanish. This keeps the support price independent of a field-dependent selection.

The numerical assembly has the following exact form. Let the Gaussian-chain estimates give envelope families with a common Sales bound $N(x)$ of rate $-1/2$. Applying the same estimate to the blank geometry controls the positive reference determinant term with its own envelope, not by identifying it with the actual determinant. Let the local Short estimate have envelope $Z(x)$, whose support sales bound $X(x)$ has rate $-1/2$. For the latter count, the number of captions per coarse site is at most

$$|Signed|spots(L) + 2itineraryDensity(L),$$

and a sleeve of halo l is contained in the radius-(l+1)B neighborhood of its caption. Its common sales price is bounded by

$$\begin{aligned}&\mathrm{ballCost}((l+1)B)\cdot captionDensity\cdot Z(x)\\&\cdot\exp(v\,\mathrm{ballCost}(l+1)).\end{aligned}$$

The radius census has rate zero, so this retains rate $-1/2$.

The inherited log and forced activities have their own preassigned envelopes from the incoming representation. Combine the log envelope, the Short envelope, and the chain envelopes by disjoint union to obtain b_0. Their sales prices add. Choose exponents $-1/2 < s < b < 0$. The inherited log budget can be taken as $betaF^s$, whereas X+N has rate $-1/2$, so

$$T(x) = betaF(x)^s + X(x) + N(x)$$

has rate s. Eventually $T \le 1$ and $2T \le betaF^b$. These thresholds are chosen uniformly before quantifying over the geometry and fields.

The combined sales bound and nonempty support imply $b_0 i \le T \le 1$ for every letter. The absolute logarithm is bounded by $b_0 i$, and

$$|\exp(t) - 1| \le 2b_0 i when |t| \le b_0 i \le 1.$$

Thus $y_i = 2b_0 i$ bounds both the raw logarithm and its Mayer activity, with combined sales price at most $betaF^b$. The inherited forced envelopes retain their separately prescribed $betaF^u$ budget for any fixed u. A dummy geometry at the same x, cutoff and width is used only to read the common numerical envelopes; the bounds themselves apply to every admitted geometry, every forest, and every field. This completes the dictionary assembly conditional on its inherited, Short and chain estimates, which are unfolded next.

6.66 Local driving estimate for the nonlinear Gaussian floor

The local driving estimate supplies the quantitative bridge between the geometric predictor and the activity expansion. It must be applied to the specific retained and noise coordinates constructed above, on their declared healthy chart. The proof keeps the quadratic positive floor while assigning the cubic, quartic and path errors to the finite envelopes already selected. The remaining subsections convert those anchored envelopes into the support-hit bounds required for the marked integration.

We now supply the local estimate used in the earlier quadratic-plate absorption argument. Fix one real accepted sector and its forest. For a variation u, let u_{hot} be its restriction to the active bonds. At a coarse site o, define

$$I_o(u) = \left(\sum_{\mathrm{dist}(o,\mathrm{addr}(e)) \le 3B} \|u_{\mathrm{hot}}(e)\|^2 \right)^{1/2}, \quad R_o(u) = \left(\sum_{\mathrm{dist}(o,\mathrm{root}(p)) \le R_{\mathrm{mac}}+2} \|\,\mathrm{motion}(u)_p\|^2 \right)^{1/2}.$$

Each individual term in either sum is bounded by its square root, because all summands are nonnegative. Suppose a gauge transformation z makes all background links within distance $2B$ of o have chord at most $h \ge 0$. The accepted local stencil estimate then says that the flat row and the background-oriented row, applied to $\mathrm{Ad}_z\, u_{\mathrm{hot}}$ at a rod rooted within $R_{\mathrm{mac}} + 2$ of o, differ by at most

$$C\, 2c_G(B+1)h\, I_o(u).$$

This uses only the background links on that local box. Adjoint action preserves the color norm. The rod is selected because an active bond is separated from the forbidden seeds by the prescribed buffer; the triangle inequality places its entire stencil inside the clean region. Gauge covariance identifies the oriented row with the corresponding measured motion component. Its norm is at most $R_o(u)$. Consequently the flat row is bounded by

$$E_o(u) = R_o(u) + 2Cc_G(B+1)h\, I_o(u). \tag{NF1}$$

Here is how the local Poincaré estimate turns (NF1) into control of a bond. Write w for the logarithmic color variation of $\mathrm{Ad}_z\, u_{\mathrm{hot}}$. On every forest trunk u_{hot} is zero, hence its collected logarithmic variation is zero. The forest link estimate therefore bounds the variation at an active bond e, with $o = \mathrm{addr}(e)$, by

$$\|u_{\mathrm{hot}}(e)\| \le K \sum_{x \in \mathrm{box}(R_{\mathrm{mac}})} \sum_{\alpha \in \mathrm{PlaqAxes}} \|\,\mathrm{curl}\, w(x,\alpha)\| + \|\,\mathrm{bill}\, w(o, e_{\mathrm{axis}})\|.$$

The low face row in (NF1) is $\sqrt{q_{\mathrm{dir}}}\,\mathrm{curl}\, w$, where $q_{\mathrm{dir}} > 0$ is the direct quadratic weight. Its norm is at most $E_o(u)$, so each face norm in this sum is at most $q_{\mathrm{dir}}^{-1/2} E_o(u)$. The pin row is $-\sqrt{k}\,\mathrm{bill}\, w$,

with $k > 0$ the fixed penalty, and likewise its bill norm is at most $k^{-1/2}E_o(u)$. All the face roots lie within $R_{\mathrm{mac}} + 2$ of o, by the parent-distance bound. Thus

$$\|u_{\mathrm{hot}}(e)\| \le UE_o(u), \qquad U = K|\operatorname{box}(R_{\mathrm{mac}})||\operatorname{PlaqAxes}|q_{\mathrm{dir}}^{-1/2} + k^{-1/2}. \tag{NF2}$$

No global pure-gauge representation of the background has been assumed.

The clean-background flattening estimate supplies the gauge in (NF1) with $h \le Ys(x)$, uniformly over the actual clean boxes, where $s(x)$ has rate $x^{-1/2}$ in the source's rate calculus. The local chord chart supplies this bound by writing each transformed link as the exponential of a Lie-algebra element of norm at most $s(x)$, then using its uniform chord bound. The chart radius is fixed before taking x large. Put

$$M(x) = 2UCc_G(B(x) + 1)Ys(x).$$

Equation (NF2), and the trivial zero bound on inactive bonds, give

$$\|u_{\mathrm{hot}}(e)\| \le UR_{\mathrm{addr}(e)}(u) + M(x)I_{\mathrm{addr}(e)}(u).$$

The neighborhood census bounds the sums of the squared local quantities by their global sums times

$$\operatorname{swept}(L, d) = \operatorname{ballCost}(d)L^4|\operatorname{Axis}|.$$

Squaring the preceding inequality with $(a + b)^2 \le 2a^2 + 2b^2$, summing, and requiring $2M(x)^2 \operatorname{swept}(L, 3B(x)) \le 1/2$, yields

$$\sum_e \|u_{\mathrm{hot}}(e)\|^2 \le 4U^2 \operatorname{swept}(L, R_{\mathrm{mac}} + 2) \sum_p \|\operatorname{motion}(u)_p\|^2. \tag{NF3}$$

This smallness holds on a common tail: the radius census has rate zero, while M has rate $-1/2$, so the absorbed quantity has negative rate. All constants are chosen after the block side L, before the sector and the variation.

Finally the real plate is the motion Gram matrix plus the identity on inactive coordinates. The exact normalization is

$$\langle v, \operatorname{form}^T \operatorname{form} v\rangle = \operatorname{scale}^2 \sum_p \|\operatorname{motion}(u)_p\|^2.$$

The Euclidean square norm splits into its active and inactive squares. If $a > 0$ is a common constant in (NF3), the plate therefore obeys

$$\langle v, \operatorname{plate} v\rangle \ge \min(\operatorname{scale}^2/a, 1)\|v\|^2.$$

This derives the uniform Gaussian floor from one accepted branch's local forest inequality and clean-background flattening, rather than assuming a floor for the stacked action.

The forest inequality used in (NF2) is obtained directly from the nonlinear slice-shell inequality. For a logarithmic direction w that vanishes on the forest trunks, consider $U_t = \operatorname{represent}(tw)$. Then U_t remains on the same forest slice for every real t, and $U_0 = 1$. All predictor frame weights are usable near $t = 0$: the minimum condition holds at the identity and on a neighborhood, and there are finitely many coarse roots. The shell inequality, with comparison background the identity, gives

$$\operatorname{chord}(U_t(e)) \le K \operatorname{port}_o(U_t) + \operatorname{chord}(\operatorname{prediction}(U_t)(o, e_{\mathrm{axis}})).$$

Each local port is bounded by the full sum of plaquette chords over the macro box. This may overcount periodic plaquettes, which only enlarges a nonnegative sum.

If a group-valued curve V_t starts at the identity and has matrix derivative $\mathrm{colorMat}(v)$, the normalization of chord and the color norm imply

$$|t|^{-1}\,\mathrm{chord}(V_t) \longrightarrow \|v\|.$$

Indeed the derivative quotient of its matrix is $\mathrm{colorMat}(v)$; take its norm and divide by the fixed color scale. Apply this fact to the edge, to each plaquette holonomy, and to the actual predictor. Their derivatives are respectively the collected edge variation, its linear curl, and the predictor bill. Divide the shell inequality by $|t|$, let $t \to 0$ through nonzero real values, and pass through the finite port sums. The resulting inequality is exactly the forest inequality used in (NF2). In particular its pin term is the derivative of the nonlinear predictor chosen by the construction.

This argument also gives the fixed-block vertical Hessian gap. If $X = \mathrm{box}(R_{\mathrm{mac}})$, squaring the forest inequality and applying Cauchy–Schwarz to the two finite sums yields

$$\sum_e \|w(e)\|^2 \le \beta \sum_b \|\operatorname{curl} w(b)\|^2 + \gamma \sum_b \|\operatorname{bill} w(b)\|^2,$$

with

$$\beta = 2K^2L^4|\operatorname{Axis}||X|^2|\operatorname{PlaqAxes}|, \qquad \gamma = 2L^4.$$

For fixed offset the map from coarse root and plaquette type to the fine plaquette is injective, so summation over coarse roots costs $|X|$, even when different offsets represent the same periodic site. Parent fibers have exactly L^4 elements. These two observations account for all factors and avoid any period-dependent norm equivalence.

On a vertical forest coordinate v, the stretched field has collected value zero on trunks and v_e on the complementary bonds. The Hessian of the tested action is bounded below by

$$a \sum_b \|\operatorname{curl} w(b)\|^2 + k \sum_b \|\operatorname{bill} w(b)\|^2,$$

where $a > 0$ is the curvature-metric floor. The pin term has this sign because the vertical coarse mismatch is minus the bill. Therefore

$$\mathrm{Hess}(v,v) \ge \frac{\min(a,k)}{1+\beta+\gamma} \sum_{e \notin \mathrm{trunks}} \|v_e\|^2. \tag{NF4}$$

The Hessian restriction follows by composing the joint analytic action with the linear vertical inclusion; its second derivative is the joint second derivative evaluated on two vertical vectors.

The local stencil comparison in (NF1) has a separate elementary transport proof. A carrier from x to y has path length at most $4\,\mathrm{dist}(x,y)$. When all links on a surrounding radius-d box have chord at most h, and that path lies in the box, the carrier's chord is at most dh, by the product chord inequality. To apply it using only local hypotheses, replace all links outside the box by the identity; locality leaves the carrier unchanged. The matrix adjoint estimate

$$\|\operatorname{Ad}_g b - b\| \le 2c_G\,\mathrm{chord}(g)\|b\|$$

then bounds straightening of a local variation of size M by $2c_G dhM$. A bounded stencil of norm at most C, depending only on its declared lookup neighborhood, transfers this to a row error at most $2Cc_G dhM$. One can see the locality step explicitly by cutting the difference of the two variations to that neighborhood; its sup norm is at most the local bound, and the stencil sees the same difference.

Under a signed coordinate change, a reversed link contributes a further adjoint correction of at most $2c_G hM$; forward links contribute none. The transformed lookup is contained in the original box enlarged by one. The total comparison is thus

$$\| \operatorname{orientedRow}(V,u) - \operatorname{orientedRow}(1,u)\| \le 2Cc_G(d+1)hM.$$

Taking the box and variation bounds specified before (NF1) proves the asserted estimate. Adjoint action and the signed displacement preserve each individual variation norm.

The same transport calculation preserves the entrywise decay needed for the matrix upper price. A single-bond impulse remains a single impulse after the signed displacement and carrier conjugation, and its color norm remains unchanged. The untransported stencil's impulse decay therefore implies

$$e^{w\,\operatorname{dist}(e,s\ \operatorname{focus}(p))}\| \operatorname{orientedRow}_{s,p}(V,\mathbf{1}_e b)\| \le 2Ce^{w(\operatorname{predReach}+7)}\|b\|,$$

where $w>0$ is the common stencil width. The extra one in this exponent pays the possible signed-link endpoint shift. Thus the local comparison does not introduce a volume factor into the matrix estimate.

Finally the small logarithmic background in (NF1) is constructed, not assumed. For a clean radius-$(d+1)$ box, $d \le 32B(x)$, axial gauge flattening gives a gauge z with chord at most

$$u(x) = 16(2(32B(x)+1)+1)^5|\operatorname{PlaqAxes}|\,2\operatorname{fiery}(x)$$

on the radius-d box. The accepted temperature comparison bounds the clean plaquette threshold by $2\operatorname{fiery}(x)$. A fixed local logarithm chart extends these transformed links to a whole logarithmic array b, agreeing on this box, with $\|b\| \le C_{\log}u(x)$, once $u(x)$ lies in its fixed chart radius. The radius factors have rate zero and the fiery threshold has rate $-1/2$; hence $H(x) = C_{\log}u(x)$ has rate $-1/2$ and eventually lies in that chart. This supplies the function $s(x)$ used in (NF3).

The logarithmic product estimates used here follow from the analytic local map $B(a,b) = \log(\exp a \exp b)$. It satisfies $B(0,0)=0$ and $DB(0,0)(a,b) = a+b$, by differentiation of matrix multiplication and the local inverse relation between logarithm and exponential. A bounded second derivative on a sufficiently small closed ball therefore gives

$$\|B(a,b)-(a+b)\| \le C\|(a,b)\|^2, \qquad \|B(a,b)\| \le C\|(a,b)\|.$$

The local exponential recovers the product on that ball. Applying this at each link gives the same constants for every finite lattice in the sup norm, since the operation is coordinatewise. This justifies use of a genuine left group chart in the comparison, including its quadratic error.

The extension of local logarithms used above has a particularly simple form. The analytic matrix logarithm chart satisfies $\log 1 = 0$ and, on a fixed neighborhood,

$$\exp(\log U) = U, \qquad \|\log U\| \le C\ \operatorname{chord}(U).$$

The second inequality follows from the local Lipschitz bound of the analytic chart and $\|\operatorname{mat}(U) - 1\| = c_G \operatorname{chord}(U)$. Choose the chord radius small enough to lie both in the Lipschitz neighborhood and in the neighborhood of the right-inverse identity. On the requested local set of edges define $b_e = \log U_e$, and outside it define $b_e = 0$. This is a whole finite array with $\|b\|_\infty \le Ch$; its exponential agrees with U on the requested set. No continuous cutoff or global logarithm of U is involved.

The adjoint estimate in (NF1) is likewise a direct matrix inequality. For a unitary U and a color matrix X,

$$UXU^* - X = (U-1)XU^* + X(U^*-1).$$

Unitary invariance of the Frobenius norm and submultiplicativity bound its norm by $2\|U-1\|\|X\|$. Dividing by the color scale proves the stated physical-color inequality. The path chord inequality follows by induction on path length, using the product triangle inequality and invariance of chord under inversion; it also applies to paths with repeated edges.

A related estimate needed in the local vertex remainders is

$$\| \operatorname{curvature}(\operatorname{face}(\exp b)) - \operatorname{curl} b\| \le C\|b\|_\infty^2$$

on a fixed sup-norm ball, with the same constant for every lattice. To prove it consider the analytic finite-dimensional map

$$(a_0,a_1,a_2,a_3) \longmapsto \operatorname{curvature}(\exp a_0 \exp a_1 \exp(-a_2)\exp(-a_3)).$$

It vanishes at zero and its derivative there is $a_0 + a_1 - a_2 - a_3$. Taylor's theorem with bounded second derivative gives the quadratic error and a uniform bound on this linear part. Substituting the four actual face coordinates gives the result; their maximum norm is at most $\|b\|_\infty$, even when periodicity identifies two or more of these coordinates.

6.67 Common envelopes for the non-packet activity letters

The regular-log and forced-letter estimates must use one envelope independent of the background sector, its forest, and the Gaussian variable. The assembly is as follows. Fix $\varepsilon > 0$ below one quarter of the support exponent and below one eighth of each of the two stencil widths, a nonnegative footprint weight v, and a retained order $l \ge 4$. Choose the block side large enough for the inherited footprint estimates. After this choice fix the positive Gaussian floor and matrix upper-price parameters, the local chart constants, and the remaining real sector parameters. Suppose the covariance-window error has rate -1.

For every exponent $-1/2 < b < 0$ and any prescribed forced-letter exponent u, there are common nonnegative weights w_i for forced letters and y_p for the other logarithmic letters such that

$$|\operatorname{bag}_i(U)| \le w_i, \qquad |\operatorname{farLog}_p(z)| \le y_p, \qquad |e^{\operatorname{farLog}_p(z)} - 1| \le y_p. \tag{NF5}$$

Their weighted footprint sums have respectively budgets β_F^u and β_F^b. These are the sales inequalities used in the connected expansion: the weights include the specified footprint exponential, overlap price, radius charge and retained derivatives. All the ancillary costs enter before the budget is reduced to the displayed powers.

To see the exponent bookkeeping, the local estimates give a base envelope $b_p^{(0)}$, obtained by combining the inherited physical letters, the finite local logarithmic letters and the covariance-window letters. Their combined sales budget has the form

$$T(x) = \beta_F(x)^s + X(x) + N(x), \qquad -1/2 < s < b,$$

where X, N have rate $-1/2$. Thus T has rate s; eventually $T \le 1$ and $2T \le \beta_F^b$. The envelope is chosen by evaluating the uniform estimates on an empty-seed dummy sector at the same radius and cutoff, then retaining their background-independent weights. This is legitimate because the sector hypotheses hold there and the estimates themselves quantify over all real backgrounds. It makes no choice based on the subsequently integrated Gaussian value.

The base envelope bounds each logarithm and is individually at most one, since a nonnegative weighted sales sum bounds each of its summands. For $|q| \le a \le 1$, the elementary bound $|e^q - 1| \le 2a$ supplies the Mayer envelope. Consequently set $y_p = 2b_p^{(0)}$; multiplying its sales budget by two costs at most β_F^b, as just arranged. The forced-letter envelope keeps its separately prescribed budget. The real localization factors all lie in $[0, 1]$, so multiplying by them can only decrease these pointwise bounds. This establishes (NF5) with one set of weights usable in every sector and for both the raw logarithms and their Mayer factors.

The letters combined in (NF5) are the following actual terms. Besides the inherited activity logarithms, there are three finite local types: row-energy remainders indexed by a signed orientation and a row, a diagonal logarithmic correction indexed by a fine bond, and a Haar-chart correction indexed by an active fine bond. There are also five chain types: the principal log chains for the real plate with coefficient $-1/2$, the corresponding blank-plate chains with coefficient $+1/2$, the covariance correction chains with coefficient $+1/2$, the affine first-order correction, and the residual quadratic chain. Their chain lengths are their additional charges; the affine bond term has charge zero. Their supports are the corresponding full receipts enlarged by the halo parameter lB, including every lookup used by the term.

For a row with coarse center o, the actual short exponent is zero within distance $200B$ of the seeds, and otherwise is

$$-\beta\vartheta\left[g_o(u)\,\mathrm{light}_o(u) + \mathrm{extracted}_o(\mu)\right], \qquad 0 \le g_o(u) \le 1, \quad 0 \le \vartheta \le 1. \tag{NF6}$$

The diagonal term is zero when either endpoint of its bond meets the radius-$200B$ seed dome, and otherwise equals

$$\tfrac{1}{2}\big(\,\mathrm{joint}(\mathrm{plate}_{\mathrm{blank}}) - \mathrm{joint}(\mathrm{plate})\big)_e.$$

The Haar term is the localized chart logarithm on active bonds outside the radius-$304B$ seed dome and is zero otherwise. These signs are exponent signs: the activity is $e^{\mathrm{Short}} - 1$.

On an intact real sector the row quantity light equals the actual debit. Indeed outside the dome the local slice is interior and the predictor is alive. The kept-path test is satisfied on the local region by the clean plaquette bound; outside it the exponential path threshold is at least two and the universal chord bound supplies the test. The exact row-energy splitting can therefore be applied to the actual configuration, with its existing choices.

Here is the size estimate for (NF6). A nonzero localization factor supplies the local variation bound required by the grafted-input energy estimate. That estimate bounds $|g_o(u)\,\mathrm{light}_o(u)|$ by a common $J(x)$ of rate $-3/2$. The protected-mean estimate bounds each of the at most two parent path values of μ by $A(x)$, of rate -1. The extracted term is a sum of at most two nonnegative half-squares, so it is at most $A(x)^2$. Since $\beta \le 2\beta_F(x)$, its row contribution is bounded by

$$W(x) = 2\beta_F(x)\big(J(x) + A(x)^2\big),$$

of rate $-1/2$. The diagonal-log difference has a common bound $E_{\mathrm{diag}}(x)$ of rate $-1/2$, obtained by local gauge flattening and the principal-inverse comparison for the real and blank plates with the same window geometry. The localized Haar logarithm has bound $C_{\mathrm{Haar}}\,\mathrm{fiery}(x)$, also of rate $-1/2$. Adding these three nonnegative bounds gives a common bound for every short letter.

To turn this pointwise bound into its sales bound, the center census is

$$|\,\mathrm{Signed}\,|\,\mathrm{spots}(L) + 2\,\mathrm{density}(L).$$

Every support is contained in the radius-$(l+1)B$ neighborhood of its center. Thus a common short-letter envelope $E_{\mathrm{short}}(x)$ has weighted budget at most

$$\mathrm{ballCost}((l+1)B)\,\big(|\,\mathrm{Signed}\,|\,\mathrm{spots}(L) + 2\,\mathrm{density}(L)\big)E_{\mathrm{short}}(x)e^{v\,\mathrm{ballCost}(l+1)}.$$

Its rate remains $-1/2$. This is a local census bound, independent of the total number of sites.

The inherited letters satisfy their own support estimates before they are joined to these local types. For the incoming regular activities, use their common healthy-sector sup envelopes; the masked activity is zero off the healthy set. For forced activities choose an anchor in each nonempty inventory, which lies in its support. The incoming support budget then becomes an anchored budget at that anchor. The footprint conversion expands each old support to its inspected coarse footprint, paying the retained overlap and charge weights. Its exponential reserve is available once 4ε is less than the incoming support exponent and L exceeds the fixed conversion threshold. These choices give the regular budget β_F^q for every $q > -11/10$ and the forced budget β_F^u for every prescribed u.

A degree-k translated local atom, with coefficient bound of rate a and aggregate magnitude budget of rate b, has pointwise envelope equal to that coefficient bound times $\mathrm{fiery}(x)^k$ times its fixed atom magnitude. Its sales budget therefore has any exponent strictly greater than

$$a - k/2 + b.$$

Finite signed orientations only multiply the budget by their fixed cardinality. Applying this rule to cubic and quartic terms with coefficient of rate one, and to the loop quadratic term with coefficient of rate zero, places all three below every exponent strictly greater than $-1/2$. The strict inequalities leave room for footprint conversions and their polynomial charges. True path tails, including the lower-scale tail, have an arbitrary power saving after their positive retained cutoff order is imposed; their additional orientation and plaquette-type sums are finite. The six inherited classes can consequently be combined with an intermediate exponent $s < b$, and the fixed factor six is absorbed by $\beta_F^b/\beta_F^s \to \infty$. This supplies precisely the inherited part of the common envelope construction (NF5).

The ordered-chain support estimate can be stated and proved independently of the matrix formulas. A route of length n is an ordered list of n fine bonds after a starting bond. A hop is allowed if an endpoint of one bond is within coarse distance d of an endpoint of the next. Hence their coarse centers are within $d + 2$. The number of allowed next bonds is at most

$$P_d = \mathrm{ballCost}(d + 2)\,\mathrm{density}(L), \qquad \mathrm{density}(L) = L^4|\,\mathrm{Axis}\,|.$$

Induction on length bounds the number of allowed routes from any starting bond by P_d^n, including routes with repeated bonds.

Attach to every route node the radius-mB halo of both bond endpoints. The full receipt is the union of these halos. If $d+2 \le mB$, successive halos meet. Each halo has B-cover number at most $\mathrm{steady}(m) = \mathrm{ballCost}(m + 1)$. The union therefore has cover number at most

$$(n + 1)\,\mathrm{steady}(m).$$

Every point in it is within $(n + 1)(mB + 1)$ of the starting center, by the triangle inequality along the route.

Suppose a chain letter has $n+u$ hops, with u fixed obligatory hops, charge n, and magnitude at most Cq^n on allowed routes and zero on the others. Fix a coarse site x and sum over all such letters whose receipts contain x, including their charge factor $e^{2\varepsilon n}$ and footprint factor $e^{v\ \mathrm{coverNumber}}$. Its possible starting bonds number at most

$$\mathrm{ballCost}((n + u + 1)(mB + 1))\,\mathrm{density}(L).$$

The elementary inequality $n + u + 1 \le (u + 1)2^n$ gives, in four dimensions,

$$\mathrm{ballCost}((n + u + 1)(mB + 1)) \le 16^n\,\mathrm{ballCost}((u + 1)(mB + 1)).$$

Combining the starting-point count, route count, magnitude, charge and cover costs bounds the length-n contribution by

$$HC(q\mathcal{R})^n,$$

where

$$H = \mathrm{ballCost}((u+1)(mB+1))\,\mathrm{density}(L)P_d^u e^{v(u+1)\,\mathrm{steady}(m)}, \qquad \mathcal{R} = 16P_d e^{v\,\mathrm{steady}(m)+2\varepsilon}.$$

If $q\mathcal{R} \le 1/2$, summing the geometric series gives the sales bound $2HC$, uniformly in x and in the lattice period. This counts every endpoint and its complete lookup halo before any geometric shortening. Summing over the finite route coordinates at fixed n, then over n, is justified by nonnegativity; finite partial sums obey the same bound.

The localization and chart factors used in these estimates have exact identities worth recording. Let

$$\gamma_a(u,e) = \chi_{\mathrm{inner}}(at, \|u_e\|), \qquad \Gamma_a(u,Y) = \prod_{e\in Y} \gamma_a(u,e).$$

Both lie in $[0,1]$. If $at > 0$, then $\|u_e\| \le at$ implies $\gamma_a = 1$, and $\gamma_a \ne 0$ implies $\|u_e\| \le 2at$, hence $\gamma_{2a} = 1$. Thus a nonzero Γ_a bounds every variation on its selected edge set and gives $\Gamma_{2a} = 1$ there. For the Haar chart density J, the localized logarithm is $\gamma_{2a}\log J$. If $\gamma_a \ne 0$ and $2at$ lies within the chart radius, its exponential equals J exactly.

Its all-argument size bound follows without assuming the argument is in the chart. If $\gamma_{2a} = 0$, the localized logarithm is zero. Otherwise $\|u_e\| \le 4at$. Since $J(0) = 1$, positivity and analyticity on a fixed chart give $|\log J(u_e)| \le C\|u_e\| \le 4Cat$, once $4at$ lies in that fixed neighborhood. Multiplication by $\gamma_{2a} \le 1$ preserves the bound. This proves the uniform Haar contribution used in (NF6), with the shrinking threshold chosen on one common tail.

6.68 Exact warning factors and their vanishing region

There are five detector addresses: a fine plaquette, two distinct types of coarse bond, a coarse site, and a fine site. Their raw factors are respectively the fine cutoff, the predictor-noise cutoff, the pin-frame mass, the site-frame mass, and the contracted flag charge. Each address has its fixed district containing its focus, and every district point is within one coarse step of that focus. The warning support is its radius-lB halo. There is one further warning type, the complement of a fine-bond logarithmic cutoff, with the two-endpoint support used above.

A detector is compulsory when its district meets the radius-$300B$ seed dome. Its optional warning letter is then zero. Otherwise it is

$$\Gamma_{2a}(u, \mathrm{plaza}(\mathrm{halo}_B(\mathrm{district})))(\mathrm{rawStorm}(u) - 1). \tag{NF7}$$

The fine-bond warning is $\gamma_a(u,e) - 1$ on active bonds outside the radius-$304B$ dome, and zero on all other bonds. Set

$$C_{\mathrm{storm}} = 1 + |\,\mathrm{FrameTag}\,| + |\,\mathrm{Points}\,|.$$

Every raw detector lies between zero and C_{storm}: the two cutoffs lie in $[0,1]$, each frame mass is bounded by its finite number of summands, and the contracted charge is zero or one. Thus $|\,\mathrm{rawStorm} - 1| \le C_{\mathrm{storm}}$. Since all localization factors lie in $[0,1]$, every warning has absolute value at most C_{storm}.

Write the Gaussian-coordinate input as $u = \mu + \beta^{-1/2}z$. A quiet box of radius D about o means

$$|z_{e,\alpha}| \le \xi t\sqrt{\beta} \quad \text{whenever } \mathrm{dist}(o, \mathrm{addr}(e)) \le D. \tag{NF8}$$

If $\|\mu_e\| \le Mt$, this implies

$$\|u_e - \mu_e\| \le n_G \xi t, \qquad \|u_e\| \le (M + n_G \xi)t,$$

where n_G is the number of color coordinates. The estimate follows by bounding the color Euclidean norm by the sum of the absolute coordinate values. It uses no independence between coordinates.

Here are the choices that make every ordinary detector equal to one on a quiet box. Fix the background guide bound $A > 0$. First choose $\tau = \sqrt{q_{\mathrm{dir}}}/32 > 0$, then take L large enough that

$$\frac{2}{\sqrt{q_{\mathrm{dir}}}}(A/L^2 + \tau) < \frac{1}{4}.$$

After L, take W large enough that

$$\frac{2}{\sqrt{k}}(A + \tau) < W.$$

Choose $\xi > 0$ so small that $n_G \xi \le 1$ and the row-stencil response to the noise in (NF8) is at most $\tau t/3$. A concrete admissible choice is

$$\xi = \min\left((n_G + 1)^{-1}, \frac{\tau}{3 n_G \operatorname{scale}^{-1} |Y| + 1}\right),$$

where Y is the uniform row-price bound.

The remaining row bill has three contributions: the nonlinear-to-linear row error, the row change from μ to u, and the protected row at μ. The first and third have rate -1 and hence are eventually at most $\tau t/3$ each, since the threshold t has rate $-1/2$. The middle contribution is at most $\tau t/3$ by the choice of ξ. Thus every relevant actual row bill has norm at most τt.

For a fine plaquette add its guide norm at most At/L^2; its weighted curvature is at most $(A/L^2 + \tau)t$. The local curvature-to-chord inequality and the displayed fine quota give its chord strictly below $t/4$. For a predictor bond the same argument, with guide norm at most At and penalty k, gives mismatch chord strictly below Wt. The local small-field tag identities give site-frame and pin-frame masses exactly one. These estimates apply outside the radius-$30B$ dome on a quiet radius-B box and use the actual predictor choices.

To deduce vanishing of (NF7), suppose the radius-$2B$ box about its focus is quiet. If the detector is compulsory, the letter is already zero. Otherwise its focus is outside the radius-$300B$ dome. Every point within distance B is then outside the radius-$30B$ dome, and its radius-B box is quiet, by the triangle inequality. The preceding conclusions apply there. The fine and noise cutoffs are on their unit plateaus, and both frame masses are one. Finally a nontrivial contracted flag would supply a touched plaquette with chord greater than t. Its parent lies within the same radius-B neighborhood, where its chord is less than $t/4$, a contradiction. Hence the flag charge is also one. Each of the five raw detectors is therefore one, so its warning letter vanishes.

The fine-bond complement similarly vanishes whenever its protected mean and noise place $\|u_e\|$ on the unit plateau of γ_a. Consequently a nonzero optional warning has a coordinate violating a bound of the form (NF8) within its fixed enlarged support. This is the witness statement used by the joint Gaussian catastrophe estimate. It is valid for an individual processed warning regardless of the values of the other cutoffs, so no product of independent warning probabilities is assumed.

6.69 From anchor bounds to support-hit bounds

Here is the summation underlying the inherited support budget. Let a map p from fine to coarse sites have at most P fine preimages per coarse site. Labels i have fine anchors o_i, nonnegative

weights b_i, and integer sizes d_i, with

$$\sum_{i:o_i=y} b_i e^{Ad_i} \le \gamma \quad \text{for every fine site } y.$$

Suppose their coarse supports X_i obey, for $x \in X_i$,

$$\operatorname{dist}(x, p(o_i)) \le W(1+d_i), \qquad \operatorname{cov}_B(X_i) \le v + v' d_i.$$

Let $g_i \le C b_i (1+d_i)^k$, where $g_i, C \ge 0$. If $a \ge 0$, $\epsilon > 0$, and $av' + 2\epsilon \le A$, then

$$\sum_{\substack{x \in X_i \\ d_i \ge m}} g_i e^{a \operatorname{cov}_B(X_i)} \le \frac{e^{av} \gamma C P (2W+1)^4 \epsilon^{-(4+k)} (4+k)!\, e^{\epsilon}}{1 - e^{-\epsilon}} e^{-\epsilon m}. \tag{NF9}$$

The sum on the left is over labels i, with x fixed. To prove it, first restrict to $d_i = j$. All anchors lie over the coarse ball of radius $W(1+j)$ about x, so there are at most $P(2W+1)^4(1+j)^4$ possible fine anchors. Sum the given anchored budget over these anchors and use

$$g_i e^{a \operatorname{cov}_B(X_i)} \le b_i e^{Ad_i} C (1+d_i)^k e^{av - 2\epsilon d_i}.$$

The inequality

$$(1+j)^{4+k} \le \epsilon^{-(4+k)} (4+k)!\, e^{\epsilon} e^{\epsilon j}$$

follows by retaining one term in the exponential series at $\epsilon(1+j)$. The degree-j bound is therefore the numerator in (NF9) times $e^{-\epsilon j}$. Sum its geometric tail. This proof also covers infinitely many address variants, since every finite partial sum has the same bound and all summands are nonnegative.

For a block ratio L, $P = L^4$: a fine site is uniquely specified by its coarse parent and its four residue coordinates. For the halo enlargement of the parent image of an old support, the geometric cover estimate used earlier reads

$$\operatorname{cov}_B(X_i) \le v + \frac{10v}{L} d_i, \qquad v = (2(l+6)+1)^4,$$

where $d_i = \operatorname{cov}_M(Y_i)$, $M \le 2B$, and the enlargement has halo radius lB. Each point of it is within $(l+20)B(1+d_i)$ of the parent of any chosen anchor in Y_i. Thus (NF9) applies with $W = (l+20)B$ and $v' = 10v/L$. In the inherited sales bound the extra tariff is $e^{2\epsilon d_i}$; absorb it first into the old weight, replacing the anchored exponent δ by $\delta - 2\epsilon$. The reserve condition is then exactly

$$a(10v/L) + 2\epsilon \le \delta - 2\epsilon.$$

This gives the stated payment, with its L^4 child factor and polynomial-degree cost, rather than an unspecified enlargement constant.

The geometric cover estimate just used has the following direct proof. Write L for the block ratio and use the periodic maximum-coordinate distance throughout. A site and the embedded parent of that site are at distance at most $L/2$. Embedded coarse distances are multiplied by L. The triangle inequality therefore gives

$$L \operatorname{dist}(p(x), p(y)) \le \operatorname{dist}(x, y) + L. \tag{NF10}$$

These statements use centered residue representatives and remain valid on winding paths.

Choose a realizing M-cover A of Y, with $|A| = d = \mathrm{cov}_M(Y)$. Its centers are connected at distance $2M + 1$. A unit-step walk visits them all in at most

$$N \le d(4M + 2)$$

steps. One way to construct it is to start at one center and repeatedly attach an unvisited center adjacent to the visited connected set, inserting an out-and-back detour of length at most $2(2M+1)$. For the empty cover use the empty-support case separately.

Sample this walk at times LBj, including time zero. Its parent samples form a coarse set D with

$$LB|D| \le N + LB.$$

Successive samples, and the parent of any intervening walk point and its preceding sample, are within $2B$ by (NF10). A fine point of Y is within M of a visited cover center, so its parent is within $3B$ of the parent of that center, since $M \le 2B$ and $B \ge 1$. It follows that the halo of radius lB around $p(Y)$ lies within $(l + 5)B$ of D, with the required connected covering centers.

To replace this larger-radius cover by a B-cover, tile each such ball with points spaced B apart in the four coordinate directions. For a ball of radius mB, the integer coordinate box with $|z_i| \le m+1$ suffices: coordinatewise integer division puts every point within B of one of its grid points. The grids are connected, and the connections between their centers join their union. Their combined cardinality is at most $|D|(2(m+1)+1)^4$; taking images on a torus only decreases cardinality. With $m = l + 5$, this gives

$$\mathrm{cov}_B(\mathrm{halo}_{lB}\, p(Y)) \le |D|(2(l + 6) + 1)^4.$$

Finally, $4M + 2 \le 10B$. Multiply the last inequality by L and use the sampled-walk bound to obtain

$$L\,\mathrm{cov}_B(\mathrm{halo}_{lB}\, p(Y)) \le (L + 10d)(2(l + 6) + 1)^4,$$

as claimed. The same walk bound gives $\mathrm{dist}(x, y) \le d(4M + 2) + 2M$ for $x, y \in Y$. Combining it with (NF10) and the halo radius proves the anchor-distance estimate with $W = (l + 20)B$. This completes the geometric input to (NF9).

7 Marked integration and the coupling reset

A marker is a finite algebraic label used to keep track of an insertion in a product. Distinct markers multiply only when their modes are allowed and disjoint. This nilpotent algebra makes the insertion bookkeeping finite, while the underlying scalar activity gas still requires absolute summability. The construction below first restores the actual insertions, then performs a coupling reset as an exact exponent identity, and finally integrates the primitive density with all its marked and forced terms.

7.1 The insertion algebra and its local prices

Modes and repeated occurrences The marker algebra keeps track of which insertions occur in the integrand. Its nilpotence is a finite algebraic convention, not a statement about the physical fields. Repeated occurrences of one insertion are assigned distinct occurrence marks when the corresponding product is formed. Their lattice variables can still coincide and remain the same integration variables. This is why disjointness of marker modes does not justify treating the associated scalar factors as probabilistically independent.

Fix a finite marker set V and the finite census of allowed marker modes. A mode is an allowed subset $A \subseteq V$; multiplication joins disjoint subsets and discards products outside the census. Write

$[P]_A$ for the coefficient of a polynomial in this algebra, and e_A for the corresponding monomial. The empty mode is the scalar mode and $e_{\{v\}}$ is a singleton mode. For positive marker weights w_v, put

$$t_w(A) = \prod_{v \in A} w_v, \qquad \|P\|_w = \sum_A t_w(A) |[P]_A|.$$

This norm is subadditive and submultiplicative: expand each product coefficient into disjoint pairs, apply the triangle inequality, and use $t_w(A \sqcup B) = t_w(A) t_w(B)$. The truncated exponential $\operatorname{Exp} P$, for P with zero scalar coefficient, is the finite exponential series in the nilpotent ideal. Its norm is bounded by $\exp \|P\|_w$. It obeys the exponential addition law for commuting zero-scalar arguments.

An insertion representation consists of a density exponent $E(U)$, primary functions $\Phi_v(U)$, forced profiles $\omega_v(q, U)$, and higher-degree corrections $H(S, U), K(q, U)$. Here S is a support and $q = (S_q, I_q)$ is a support together with its flagged inventory. Let $b_q(U)$ be the scalar activity profile, regarded as a polynomial in the empty mode, and let $F(U)$ be the configuration's flagged set. Restoration means the following literal finite expression:

$$X(U) = e^{E(U)} \operatorname{Exp}\left(\sum_v e_{\{v\}} \Phi_v(U) + \sum_S H(S, U) \right) \times \sum_{\mathcal{P}:\, \mathrm{Pattern}(F(U), \mathcal{P})} \prod_{q \in \mathcal{P}} \left(b_q(U) + \sum_v e_{\{v\}} \omega_v(q, U) + K(q, U) \right).$$

The pattern condition says the supports S_q are pairwise disjoint and the inventories I_q have union $F(U)$. All sums here are finite at fixed torus and census. The corrections have only modes of cardinality at least two. They are gauge invariant, measurable, supported on their stated regions and scoped to markers whose pins lie in those regions. Exponential corrections vanish off the healthy region and have zero scalar coefficient. A flagged correction can be nonzero only when $S_q \cap F(U) = I_q \neq \varnothing$, with $I_q \subseteq S_q$. These support and permission conditions are preserved by scalar multiplication.

The price of a correction family $Y_i(U)$, with supports S_i, is controlled by nonnegative envelopes a_i satisfying $|Y_i(U)|_w \leq a_i$ and the sitewise bound

$$\sup_x \sum_{i : x \in S_i} a_i e^{e\, c_B(S_i)} \leq f.$$

Here $c_B(S)$ is the minimizing B-cover number and e is the support exponent. The definition also includes summability of the envelopes. If every nonzero $Y_i(U)$ has some marker pin $o_v \in S_i$, and $e \geq 0$, then

$$\sum_i \|Y_i(U)\|_w \leq |V| f.$$

Indeed charge a nonzero term to all marker pins in its support. At least one is present. Since $e^{e c_B(S_i)} \geq 1$, summing the sitewise bound over the $|V|$ pins dominates the desired sum. The scope and positive degree conditions ensure the required pin is present. For the unmarked background profiles the same argument charges to all sites, giving $\sum_q \|b_q(U)\|_w \leq |\Lambda| u$, where u is their sitewise activity budget.

Suppose also

$$w_v|\Phi_v(U)| \le x, \qquad \sum_q w_v|\omega_v(q,U)| \le y.$$

The exponential argument then has norm at most $|V|(x+f)$, and the sum of norms of all field factors is at most $|\Lambda|u + |V|(y+f)$. Dropping pattern restrictions increases the nonnegative norm bound for the field polynomial to

$$\sum_{\mathcal{P}} \prod_{q\in\mathcal{P}} \|T_q\|_w = \prod_q (1+\|T_q\|_w) \le \exp \sum_q \|T_q\|_w.$$

Multiplication with the exponential bound proves

$$\|X(U)\|_w \le \exp\{E(U) + |V|(x+y+2f) + |\Lambda|u\}.$$

The lawful-density bound on $E(U)$ gives the uniform majorant used above. To bound a normalized mode integral, use $t_w(A)|[X(U)]_A| \le \|X(U)\|_w$, integrate against the probability prior, and divide by the positive density integral. The denominator lower bound yields the final insertion bill; the history treasury bounds that bill by $|\Lambda_k|\tau_k + |V|(x+y+2f)$. This separates the elementary polynomial estimate from the vacuum estimates required to control its density and denominator.

7.2 Restoring the prescribed singleton insertions

The correction representation must preserve the actual singleton insertions after every step. The following algebra explains that requirement. Given a represented density R, a primary function Φ, and a summable forced family g_i, its singleton density is

$$\mathrm{coin}_R(\Phi, g)(U) = R(U)\Phi(U) + e^{E_R(U)} \sum_i \mathrm{nickel}_{R,i}(U; g).$$

The second term is the first variation of the activity gas in the forced direction. Grouping the summable family by its support-and-inventory barcode gives its profile; this grouping preserves the gas variation. In the restored polynomial, the scalar mode is exactly $R(U)$. The coefficient of $e_{\{v\}}$ is exactly this singleton density: in a product contributing to a singleton mode, at most one factor can supply a singleton and every other factor supplies its scalar mode. The exponential contributes its linear coefficient and scalar coefficient one. Higher corrections contribute neither. These observations prove the scalar and singleton relay identities by finite polynomial expansion and the absolutely convergent grouping of activity coefficients.

There is a useful converse. Suppose a polynomial has the form

$$e^{E_R(U)} \operatorname{Exp}(F(U)) \operatorname{Field}(K(U), F_R(U)),$$

where F has zero scalar coefficient, the scalar coefficients of K are the original background profiles, and its singleton coefficients are the prescribed forced profiles g_v. If the singleton coefficients of the whole expression are $\mathrm{coin}_R(\Phi_v, g_v)$, then the singleton coefficients of F are precisely Φ_v. Subtract the two singleton formulas: their gas-variation terms cancel, leaving $R(U)([F(U)]_{\{v\}} - \Phi_v(U)) = 0$. The nonvanishing of $R(U)$ permits cancellation. Projection onto modes of cardinality at least two now splits F into its prescribed singleton part plus the higher exponential correction, and splits every K_q into background, prescribed singleton profile and higher flagged correction. This proves the desired restored representation, preserving the original density pointwise.

To adjust an incorrectly placed forced singleton, use attached local responses $\psi_v(S,U)$ supported on regions containing o_v, vanishing whenever that region is healthy. Set

$$J(U) = \sum_{v,S} e_{\{v\}} \psi_v(S, U).$$

The closure operation on field factors joins overlapping background factors and attached responses into their connected components, with the resulting support and flagged inventory recorded. Expansion into connected components gives the exact identity

$$\mathrm{Field}(\mathrm{Closure}(K, \psi)) = \mathrm{Field}(K) \operatorname{Exp} J.$$

Consequently replacing K by its closure and F by $F - J$ preserves the full polynomial. In the closure expansion, terms containing no response reproduce K. Terms containing at least one response have zero scalar coefficient. Terms containing at least two responses have zero singleton coefficient. The contribution containing just one response equals that response multiplied by the background portion attached to its support. Thus closure leaves the background scalar profile fixed and adds exactly the hook of ψ_v to its singleton profile. It introduces no unspecified change of singleton field.

After this identity, the preceding converse identifies the required primary functions from the known singleton densities. Keeping only higher modes produces the two higher correction families with their locality and permission conditions. Projection to higher modes decreases their weighted prices. The price of closure is controlled by the connected-component estimate: from a field budget W, response budget c, and pin occupancy M, its higher part has budget at most

$$\exp\{W + M \,\mathrm{siteFactor}(B)c\} - 1,$$

provided the support-exponent loss pays

$$(2B + 1)^4\{W + M \,\mathrm{siteFactor}(B)c\}.$$

The connected-component forest-price argument below supplies this estimate. A reweighting bound on marker monomials then transfers both higher families to the requested weights. These are the algebraic and numerical operations used in the one-step retained-insertion estimate.

The connected-component price estimate follows from a recursion on distinct labels. Here are its details. Let X_i be supports and $f_i \geq 0$, with

$$\sup_x \sum_{i : x \in X_i} f_i e^{(a+w)c_B(X_i)} \leq \gamma, \qquad a \geq 0, \quad (2B + 1)^4 \gamma \leq w.$$

For a finite label set D and root region Y, let $F_D(Y)$ be the sum of $\prod_{i \in A} p_i$, $p_i = f_i e^{a c_B(X_i)}$, over subsets $A \subseteq D$ every overlap-connected component of which meets Y. Include the empty subset with weight one. For each component choose one label whose support meets Y, and remove that label from the component. Every component of the remainder meets the removed label's support. Different original components choose different labels. The chosen labels and the corresponding remainders reconstruct the original subset, so this assignment is injective. Allowing all choices of remainders, including incompatible ones, only increases the sum of nonnegative weights. Therefore

$$F_D(Y) \leq \prod_{i \in D : X_i \cap Y \neq \varnothing} \{1 + p_i F_{D \setminus \{i\}}(X_i)\}.$$

Strong induction on $(|\mathrm{D}|)$ proves $F_D(Y) \leq e^{w c_B(Y)}$. Indeed the induction hypothesis makes each factor at most $\exp\{f_i e^{(a+w)c_B(X_i)}\}$. The sum of these exponents over supports meeting Y is

at most $|Y|\gamma$, by summing the sitewise hypothesis over Y. Finally the cover-volume bound gives $|Y| \le (2B+1)^4 c_B(Y)$, hence $|Y|\gamma \le w c_B(Y)$. The same product estimate, with this inductive bound for the smaller label sets, also gives the sharper root-region bound $F_D(Y) \le e^{|Y|\gamma}$.

If a nonempty overlap-connected subset A has union containing a fixed site x, it is one of the nonempty terms counted by $F_D(\{x\})$. For such a subset,

$$c_B\left(\bigcup_{i\in A} X_i\right) \le \sum_{i\in A} c_B(X_i).$$

To prove this cover inequality, start with one support and add supports in an order with each new support meeting the preceding union; connectedness supplies such an order. The cover number of two meeting supports is at most the sum of their cover numbers. Induction gives the claim. Since $a \ge 0$, its consequence for products and the preceding singleton-root bound imply

$$\sum_{\substack{A\neq\varnothing \text{ connected} \\ x\in\bigcup_{i\in A} X_i}} e^{a c_B(\bigcup_{i\in A} X_i)} \prod_{i\in A} f_i \le e^{\gamma} - 1.$$

For countably many labels, first bound every finite subsum by taking D to contain all its labels. Nonnegativity then proves summability and the same bound. No repeated-label expansion or logarithmic series is involved.

Apply this estimate to the disjoint union of background-factor labels and response labels, with their respective envelopes. Their sitewise budgets add, so $\gamma = W + M\,\mathrm{siteFactor}(B)c$. The envelope of a connected component is the product of its label envelopes, by submultiplicativity of the weighted polynomial norm. Selecting only the permitted components can only decrease the nonnegative sum. Grouping them by their resulting support and inventory preserves the sitewise bound. This proves the closure price and exponent loss stated above.

The exact closure identity also has a direct inverse to its component decomposition. Expand the original field polynomial and $\prod_j (1+u_j)$, choosing background factors and a subset of responses. Join labels whose supports intersect. A nonzero chosen response lies off its healthy region, so its support meets a flagged site. That site belongs to a chosen background inventory; thus the response's component contains a background factor. Background supports in the original pattern are pairwise disjoint. Each resulting component consequently has a nonempty inventory, and the components' inventories partition the original flagged set. Conversely, a permitted disjoint collection of such connected components recovers the unique original set of background and response labels by union. The coefficients are unchanged products of the same factors. This bijection proves the identity term by term. For the singleton responses used here, $u_j^2 = 0$; hence $\prod_j (1+u_j) = \mathrm{Exp} \sum_j u_j$. This completes both the algebraic identity and its connected-component price estimate.

7.3 Finite transport inventory without changing the insertion

The passage to a finite transport inventory uses two different reductions. First, a countable family of local leaves with one fixed dictionary key can be summed in that key. Suppose its complex coefficients are r_i, its nonnegative leaf bounds are b_i, and $\sum_i |r_i| b_i e^{s c_B(S)} < \infty$, where all leaves have the same support S. Sum the real and imaginary coefficient families separately. Each real sum is a real leaf with bound at most $\sum_i |r_i| b_i$; locality, gauge invariance, measurability and its eligibility condition pass to the absolutely convergent sum. The transported quote also passes to the sum by the sound and tuned transport identities. Use charges (1) and i for the two resulting real leaves. They reproduce the original complex family both before transport and in every transported profile, at total marker price at most

$$2w_v \sum_i |r_i| b_i.$$

An insertion tale uses only finitely many keys. Group its countable leaves by key, apply this construction within each group, and sum over the finite set of keys. The support weight is constant on each group, so this regrouping preserves the summed support budget. It gives finitely many singleton leaves reproducing the primary function minus its scalar displacement, the original forced profiles, and the coarse forced profiles minus their attached response hooks. A tale with price A_v yields total support-weighted singleton price at most $2w_v A_v$.

Second, higher polynomial corrections can be placed in blank dictionary slots. For a nonzero coefficient function (h(U)) in mode A, write $h = \Re h + i\Im h$, using two real leaves with that mode and charges (1,i). If $|h(U)| \le b_A$, their combined price is $2t_w(A)b_A$. The locality, scope, gauge invariance and permission conditions are inherited coefficientwise. Only modes that are nonzero somewhere need be included. Summing their prices proves that a higher correction family of budget f becomes a finite leaf family of budget at most (2f).

The dictionary contains a dormant regular slot of zero original coefficient for every nonempty support, and a forced slot of every admissible nonempty inventory with its exact eligibility condition. Place an exponential correction in the first kind of slot, and a flagged correction in the second. A nonzero higher correction has at least two scoped markers, so its support is nonempty. These placements preserve the original correction functions exactly, while the two families together cost at most (4f). Their modes all have degree at least two. The existence of these dormant dictionary slots is part of the history construction; this reduction uses that property explicitly.

Combine the singleton leaves from the tales with these higher leaves. Their primary contributions sum to

$$\sum_v e_{\{v\}}(\Phi_v - \nu_v) + \sum_S H_S,$$

and their forced contributions to

$$\sum_v e_{\{v\}}\omega_v(q) + K_q.$$

The singleton transported quotes reproduce the prescribed coarse forced profile minus its response hook. Higher leaves have no singleton quote, since their modes have degree at least two. Occupancy converts the singleton root budgets into a sitewise budget $M\,\mathrm{siteFactor}(B)d$. Together with the (4f) correction budget this is the quantity T used in the numerical step.

The exponential at a regular dictionary key must still be represented exactly. If u_i are its finitely many tagged monomials, each has nonempty marker mode and $u_i^2 = 0$. Consequently

$$\mathrm{Exp}\left(\sum_i u_i\right) = \prod_i (1+u_i) = 1 + \sum_i u_i + \sum_{\substack{A:\,|A|\ge 2 \\ \text{disjoint allowed marker modes}}} \prod_{i\in A} u_i.$$

Products with overlapping marker modes vanish, and products outside the census vanish. Each remaining product is a real product leaf at the same dictionary key, charged by the product of charges and carrying the union mode. Its price is the product of the original prices. These extra composite leaves have degree at least two; they do not change any singleton quote. They are added only at regular keys, so forced-slot contributions are untouched. This finite enlargement replaces the exponential by (1+) the enlarged key sum without changing its value.

Expanding the represented activity gas now gives the finite transport expression. At each dictionary key choose either its old scalar contribution or one marked leaf; the ticket condition is that no two chosen leaves use the same key. Each term is the old unselected activity product multiplied by the chosen marked ingredients. The polynomial normalization of each marked leaf is exact: its positive fee times its normalized monomial is its original charge times its marker monomial and leaf bound. Thus no division by a zero total coefficient occurs. The original activity series is absolutely summable under the lawful-activity bounds; only finitely many keys receive marked perturbations. Their sum is still absolutely summable, so regrouping by barcodes and interchanging the finite ticket sum with the activity series are legitimate. The result is exactly the restored fine density with the scalar displacement factored out as $\operatorname{Exp}\sum_v e_{\{v\}}\nu_v$.

The additional composite-leaf price is bounded as follows. For a key k, let p_k be the sum of its unweighted leaf prices. The common support contains a marker pin, so the sitewise budget implies $0 \le p_k \le T$. The sum of all compatible composite prices at that key is bounded by the sum over all nonempty subsets:

$$\sum_{\text{composites at } k} \text{price} \le \prod_{i \text{ at } k} (1+f_i) - 1 \le e^{p_k} - 1 \le e^{p_k} p_k \le e^T p_k.$$

The penultimate inequality follows by integrating $e^t \le e^{p_k}$ on $0 \le t \le p_k$. All composites at this key have the same support as their factors. Multiplying by its support weight and summing over distinct regular keys therefore gives a sitewise composite budget at most $e^T T$: the disjoint key classes use disjoint sets of original leaf labels. Adding the original budget yields $T + e^T T$. This proves the inventory bound with its stated constant.

The normalized monomial used in the ticket expansion has fee $c = t_w(A)|q|$ and value $e_A(q/c)$. If $c > 0$, cancellation gives $c\, e_A(q/c) = e_A q$. If $c = 0$, positivity of $t_w(A)$ implies $q = 0$, so both sides vanish under the source's zero-inverse convention. Its weighted norm is at most one in either case. This explicitly handles zero charges. Each such monomial squares to zero because $A \neq \varnothing$, which also justifies its exponential expansion above. The remaining input at this point is the Mark-step reserve acting on this finite, fully priced inventory.

7.4 The repeated-occurrence rooted-tree bound

Repeated occurrences matter in the rooted-tree estimate: two vertices may carry the same activity label and still be distinct vertices of the expansion. Summing over distinct supports too early would omit those terms or lose their factorials. The parent-map proof below keeps occurrences numbered until its nonnegative sums and factorial estimates have been completed; only then is the passage to connected sets made.

Here is the tree estimate used in the cluster argument. Set $V_B = (2B+1)^4$, $t = V_B\gamma/\delta > 0$, and $F(Y) = \max\{\delta c_B(Y), 1\}$. For nonnegative vertex weights w_i obeying the sitewise budget at exponent $a + \delta$, define

$$K(Y,i) = \frac{\mathbf{1}_{X_i \cap Y \neq \varnothing} w_i e^{(a+\delta)c_B(X_i)}}{tF(Y)}.$$

Every finite row sum of this kernel is at most one: sum the sitewise bound over sites of Y to bound its numerator by $|Y|\gamma \le V_B c_B(Y)\gamma = t\delta c_B(Y) \le tF(Y)$. Its entries are nonnegative.

Represent a forest on m numbered internal occurrences by a parent map $p : \{1,\dots,m\} \to \{0,1,\dots,m\}$, where (0) is an external root with support Y. Only parent maps admitting ranks strictly decreasing toward each internal parent are allowed. The support of every child must meet that of its parent. No condition prevents repeated activity labels at different occurrences. For a

fixed allowed parent map, the sum over all label assignments of the product of the corresponding K-entries is at most one. To prove it, sum a vertex of maximal rank first. It has no remaining children, so its label occurs only in its own row factor. Its row sum is at most one. Remove it and repeat. Finite label restrictions suffice; arbitrary label sets follow by nonnegative summation.

Let k_j be the number of children of vertex j, including the root. The elementary exponential estimate gives

$$F(S)^{k_j} \le k_j! e^{\delta c_B(S)}.$$

For $\delta c_B(S) \ge 1$, this is $x^{k_j}/k_j! \le e^x$; for $x \le 1$, use $F(S) = 1$, $k_j! \ge 1$ and $e^x \ge 1$. Multiply over parents. Each internal vertex's extra factor $e^{\delta c_B(X_i)}$ cancels the extra support weight introduced into K; only the root factor remains. Thus the original overlap-constrained weight product at exponent a is at most

$$e^{\delta c_B(Y)} t^m \left(\prod_{j=0}^{m} k_j! \right) \prod_{i=1}^{m} K(S_{p(i)}, i).$$

After summing labels and dividing by the occurrence factorial (m!), its contribution is at most $e^{\delta c_B(Y)} t^m \prod_j k_j!/m!$.

All child factorials are paid together by an exact parent-map count. For maps from m numbered objects into r bins, let $A_m(r)$ be the sum of the products of their bin-size factorials. Appending one object to bin j multiplies the weight by $k_j + 1$; summing over bins gives (m+r). Hence $A_0(r) = 1$ and $A_{m+1}(r) = (m + r)A_m(r)$. With $r = m + 1$,

$$\sum_p \prod_{j=0}^{m} k_j! = \binom{2m}{m} m! \le 4^m m!.$$

This upper bound sums all parent maps, so it also bounds the allowed ones. The total forest weight is therefore at most

$$e^{\delta c_B(Y)} (4t)^m \le e^{\delta c_B(Y)} q^m.$$

The root may instead be an internal occurrence with nonnegative weight u_i. Suppose $\sum_i u_i e^{(a+\delta)c_B(X_i)} \le b$. Multiply the forest bound rooted at X_i by $u_i e^{a c_B(X_i)}$ and sum i. The result is bq^m, with factorial (m!) for the other m occurrences. This proves the designated-root estimate used for size $m + 1$ clusters. Summing the geometric size tail gives $bq^m/(1 - q)$. The argument covers $m = 0$, empty external root and arbitrary branching without changing the normalization.

The graph interval construction used to pass from connected graphs to these trees is explicit. Fix an injective numerical ordering of the occurrence vertices and a root (0). For a connected graph H, let (d(v)) be shortest-path distance from (0). For every $v \ne 0$, choose as parent the least ordered neighbor at distance (d(v)-1). The resulting parent edges form a ranked spanning tree. Its depth function is uniquely determined by its parent map: it is zero at the root and increases by one at every child. More generally any graph with this breadth-first parent has that same depth function, since every edge changes depth by at most one and the tree realizes the assigned depths.

For a fixed tree p arising in this way, let U_p be the subset of original graph edges satisfying both of the following requirements. Their endpoint depths differ by at most one. If an edge joins a vertex $v \ne 0$ to a preceding-level vertex u, the order of u is no smaller than the order of (p(v)). Then the graphs whose breadth-first parent is p are exactly

$$\mathrm{stem}(p) \subseteq H \subseteq U_p.$$

One direction follows from shortest-path distances and the least-parent choice. Conversely, containing the stem gives a path of assigned length from the root to every vertex, while the upper constraint prevents any shorter path; hence these are the graph's actual distances. The second constraint and injectivity of the order give the same least preceding-level neighbor. Thus the breadth-first parent is again p. These intervals partition all connected subgraphs of the original graph. If $U_p \setminus \mathrm{stem}(p)$ contains an edge, toggling it pairs every graph in the interval with one of opposite parity. The interval sum vanishes. If there is no such edge, its sum has absolute value one. Summing interval bounds proves precisely the rooted spanning-tree bound on the absolute connected coefficient.

7.5 Passing from numbered trees to distinct connected sets

The same forest estimates also bound distinct-set expansions. For a distinct label set D of size m, a parent map on D, and each bijection $\{1, \dots, m\} \to D$, relabel the parent map and its activities by that bijection. The resulting numbered forest determines all three objects: its label range is D, its label assignment determines the bijection, and transport back through that bijection recovers the parent map. This assignment is injective. There are exactly $m!$ bijections; all have the same product weight divided by $m!$. Their sum is therefore the unnumbered product weight, with no factorial left. Restricting the numbered forest sum to this injective image proves the distinct-parent-map estimate and its geometric size tail.

A finite overlap-connected set whose union meets x can be rooted at the external support (x). Connect the root to labels whose supports contain x, and use a breadth-first spanning tree of the resulting connected graph. The sum over all parent maps includes this one, so it dominates that set's nonnegative product weight. For the stronger sitewise prefactor γ, instead choose one label containing x as an internal root, charge its activity to the site's budget, and use the internal-root forest estimate for the other labels. Summing over all possible roots only increases the bound. Numbering the distinct nonroot labels cancels their factorial exactly as above; allowing the root label once more among nonroots in the majorant is harmless. Consequently the distinct connected-set sum rooted at x, with the actual union support weight, is bounded by $\gamma/(1-q)$ under the readiness assumptions. This is the distinct-set input used in the positive-multiplicity restoration argument.

7.6 Cluster logarithms with repeated occurrences

Let (R(i,j)) be a symmetric incompatibility relation with (R(i,i)) for every label. Let $X_i \neq \varnothing$ be its support, with incompatibility implying intersecting supports. For real activities w_i, the hard-core gas is

$$Z(w) = \sum_{A \text{ compatible}} \prod_{i \in A} w_i.$$

The empty subset contributes one. In contrast, its logarithm uses ordered occurrence tuples $z = (i_1, \dots, i_n)$, permitting repeated labels. Form the graph of incompatible occurrence pairs and define

$$\varphi(z) = \sum_{\substack{H \text{ connected} \\ H \subseteq \text{incompatibility graph}}} (-1)^{|E(H)|}, \qquad L(w) = \sum_{n \geq 1} \frac{1}{n!} \sum_{z \in I^n} \varphi(z) \prod_{j=1}^{n} w_{i_j}.$$

For one occurrence the connected graph has no edges and coefficient one. Equal labels at distinct occurrence positions are incompatible. This is essential to the logarithmic identity even though the gas itself is indexed by distinct-label subsets.

The absolute estimate uses cancellation inside φ. Breadth-first exploration partitions connected subgraphs into intervals with one fixed spanning tree as lower endpoint and an allowed edge set as upper endpoint. In each interval, if an optional edge exists, toggling that edge is a sign-reversing involution, so the alternating sum is zero. Otherwise the interval contains only its tree and its absolute contribution is one. Hence $|\varphi(z)|$ is bounded by the number of compatible rooted spanning trees. Opening the absolute value before the connected-graph sum would lose this estimate.

Here is the resulting root accounting. Put $c_B(S) = \mathrm{coverNumber}_B(S)$, take $a \geq 0, \delta > 0, \gamma > 0$, and suppose

$$\sup_x \sum_{i: x \in X_i} |w_i| e^{(a+\delta) c_B(X_i)} \leq \gamma, \qquad 0 < q < 1, \qquad \frac{4(2B+1)^4 \gamma}{\delta} \leq q.$$

The rooted-tree estimate bounds the size-n cluster sum with a designated root-activity budget b by bq^{n-1}, using the normalization (1/(n-1)!). In an unrooted cluster that contains a label with a chosen property, sum over all possible marked root positions and divide by n. Its (1/n!) normalization is then bounded by the same rooted expression. Relabeling tuple positions preserves the graph coefficient and product weights, so the estimate holds for each possible position, including repeated labels.

If $x \in \bigcup_j X_{i_j}$, at least one occurrence has support containing x. Use that property for the root, with $b = \gamma$. A nonzero connected coefficient gives an overlap-connected support union. As in the cover argument above,

$$e^{a c_B(\cup_j X_{i_j})} \leq \prod_j e^{a c_B(X_{i_j})}.$$

Thus, summing sizes $n \geq m + 1$, the absolute logarithmic terms whose support contains x, including their actual union support weight, have total at most

$$\frac{\gamma q^m}{1-q}.$$

Summing root bounds over all sites gives absolute convergence of L at fixed torus, with total at most $|\Lambda|\gamma/(1-q)$. The rooted-tree estimate and its breadth-first interval construction are the remaining combinatorial ingredients behind this bound; the occurrence/root normalization and transfer to the union support have been made explicit here.

The real logarithmic identity follows without assuming $Z(w) > 0$. First restrict to a finite alphabet of s labels. Let $P(t) = Z(tw)$, a polynomial of degree at most s, and write $L(t) = \sum_{n \geq 1} \ell_n t^n$ for the connected series. The finite graph decomposition into connected components gives the coefficient recurrence

$$n[P]_n = \sum_{k=1}^{n} k \ell_k [P]_{n-k},$$

obtained by designating one occurrence, choosing its connected component and then the remaining graph. The binomial choices of positions and the component factorials give precisely these coefficients. The absolute cluster bound gives $|\ell_n| \leq C q^{n-1}$. Choose $u > 1$ with $qu < 1$; both (L(t)) and its differentiated series converge uniformly on compact subsets of ((-u,u)), by the summable

majorant $Cn(qu)^{n-1}$. The recurrence implies $P' = L'P$. Therefore $(Pe^{-L})' = 0$ on this interval. Its value at zero is one, so $P(1) = e^{L(1)}$.

For an arbitrary alphabet, restrict to finite subsets directed by inclusion. Every fixed gas subset and every fixed occurrence tuple eventually occurs unchanged. The gas has an absolute majorant obtained by dropping exclusions, bounded by $\prod_i (1 + |w_i|) \le e^{\sum_i |w_i|}$; the sitewise budgets make that sum finite at fixed torus. The connected series has the absolute majorant above. Dominated convergence passes both finite-alphabet sides to their full sums. Thus

$$Z(w) = e^{L(w)} > 0, \qquad \log Z(w) = L(w).$$

This also applies after removing all labels incompatible with a specified forced set: the absolute activity budgets only decrease.

7.7 Primitive marked integrals and averaging over local tree choices

The primitive activity attached to a marked team is an actual integral. A team is a finite collection of real local leaves with distinct dictionary keys; its weight is the product of their bounds and its ingredient the product of their functions. For a fixed coarse configuration and local tree choice, first choose a permitted flag sector, a local block layout, and a finite list of warning and regular-exponential selections. Their union is the molecule address p. The marked integrand is the product of the team's paid ingredients and the selected local vertex factors; it is set to zero unless all of the team's keys occur in p. Its primitive coefficient is its integral over the tree-gauge variables and the Gaussian fluctuation variables. The Gaussian measure has the covariance and density already attached to that same local tree choice. Sector selectors and all eligibility tests stay inside this integral.

More explicitly, if X denotes that local data, k the sector, and $I_{X,k,p}(h, z)$ its selected integrand, set

$$a_{X,k}(p) = \mathbf{1}_{\mathrm{approved}(X,k)} \int I_{X,k,p}(h, z)\, d\nu_X(h)\, d\gamma_X(z).$$

The full primitive coefficient averages these values over the finite boundary choices for the layout of p; it is zero for a disallowed layout. Its scalar prefactor is the original local overhead. The local vertex and integral estimates are developed in the packet, non-packet and warning arguments below.

The expansion into molecule addresses is justified before the cluster expansion. In a permitted layout, warning factors are expanded by the finite identity $\prod_j (1 + d_j) = \sum_D \prod_{j \in D} d_j$. The regular exponential factors use their absolutely convergent subset expansion. Marked regular keys are dormant in the unmarked density: their original regular coefficient is zero, so their unmarked exponential increment $e^{-0} - 1$ is zero. Consequently one may insert those marked keys into every selected list and delete them again uniquely; they are supplied by the team ingredient rather than by an unmarked exponential increment. Forced marked keys must likewise be present in the sector. If they are absent, both the selected expansion and its marked integrand vanish.

The resulting address assignment records the layout, sector, warning list and regular list. It is injective on permitted data and has an explicit inverse obtained by reading these pieces from p. Thus each nonzero molecule occurs with exactly the specified coefficient, and disallowed molecule addresses contribute zero. The summable activity majorant bounds the integral norms of these terms by the team's weight times the product of the unselected primitive tariffs. This gives a summable series of integrable terms. Fubini and the theorem on integration of a series with summable integral norms therefore give

$$\sum_{p} \text{overhead}_X \; a_X(p) = \mathcal{P}(\text{team density})(U).$$

This is the primitive team expansion. In Gaussian coordinates the change of measure is explicit: $d\gamma_X(z) = \theta_X(z)\, dz$. The transformed integrand contains that same θ_X and its tribute prefactor, so the Gaussian integral agrees with the original local fluctuation integral by substitution of the density. No independent Gaussian approximation is made in this identity.

To obtain compatible products of primitive coefficients, average over independent local tree choices at coarse sites. A molecule's integral depends only on the tree choices in its enlarged ground support. To prove this locality, project the fluctuations to the finite coordinates needed by that ground. Two local data sets agreeing there have the same projected covariance. Their Gaussian pushforwards under the coordinate projection are therefore the same measure. The integrands depend only on those projected coordinates and agree on each fiber, by their local footprint formulas. Their integrals coincide. The same argument applies to the gauge variables through their local measure-preserving coordinate maps. Thus changing exterior tree choices leaves the molecule coefficient unchanged.

Compatible molecule grounds are disjoint. For functions f_p of independent finite choices, each depending only on its ground, finite product averaging gives

$$\text{Av} \prod_{p \in A} f_p = \prod_{p \in A} \text{Av}\, f_p.$$

This follows by factoring the finite sum over assignments on disjoint coordinate sets; unused coordinates contribute one. The local scalar overhead is invariant under the tree choice. Averaging the team expansion consequently yields the product of the averaged molecule coefficients with the original compatibility and inventory conditions.

Finally expand a finite leaf inventory into tickets with distinct keys. On a compatible set of molecules, every chosen leaf key lies in exactly one molecule; selecting it there gives a unique global ticket, and conversely a global ticket restricts uniquely to each molecule. Multiplication of the leaf amplitudes factors over this partition. Hence the product of primitive marked molecule sums is exactly the ticket sum of the team's compatible product. Finite linearity of the Mark integral then proves

$$\sum_{A \text{ compatible,permitted}} \text{overhead} \prod_{p \in A} \left(\sum_{D \text{ ticket}} \prod_{i \in D} v_i \, a_D(p) \right) = \mathcal{P} \left(\sum_{D} \prod_{i \in D} v_i \text{ teamDensity}_D \right) (U).$$

The primitive-integral bounds and support-hit estimates in the nonlinear and marked-integration chapters give these summable integral-norm and footprint bounds. The address regrouping, Gaussian marginal comparison, independent tree averaging and ticket factorization show how they produce the literal Mark expansion used in the gas/log argument.

7.8 Fusing attached logarithms into the forced dictionary

The dictionary fusion used twice in the coarse construction has an explicit formula. Let the old forced labels be $i \in \mathbb{N}$, with supports S_i, nonempty inventories $I_i \subseteq S_i$, and coefficients $K_i(U)$. Let attached logarithms $\ell_j(U)$ have nonempty supports T_j, and suppose $\ell_j(U) \neq 0$ only if T_j meets the configuration's flagged set. Both families have summable uniform envelopes. Form the disjoint atom alphabet of old forced labels and attachments, with factors

$$h_{\mathrm{f},i} = K_i(U), \qquad h_{\mathrm{a},j} = e^{\ell_j(U)} - 1.$$

A new label is a finite overlap-connected atom set containing at least one old forced label, with the old forced supports mutually disjoint. Its new support is the union of all atom supports, its inventory the union of the old forced inventories, and its coefficient the product of its atom factors, restricted to its exact inventory and original eligibility conditions. These labels form a countable set; it is infinite since each singleton old forced label is one. Choose an enumeration by $\mathbb{N}$. This choice only names the labels and does not change their functions or bounds.

Expand the old forced gas times $\exp\sum_j \ell_j$ using $\prod_j(1 + (e^{\ell_j} - 1))$. Absolute summability follows from the uniform logarithm envelopes and the old forced envelopes. In a nonzero term, every attachment meets a flagged site. The exact inventory coverage of its old forced set puts that site in an old forced inventory, hence also in its support. Therefore each connected component of the selected atom supports contains an old forced label; no free attachment component occurs. Components are precisely the new labels, with disjoint supports and inventories covering the original flagged set. Conversely a permitted set of new labels has disjoint atom sets, since every atom support is nonempty; their union recovers the original old forced set and attachment subset. The exact inventory condition also implies the original local inventory conditions: another disjoint forced support cannot supply a flag within the given support. Thus decomposition and union are inverse on all nonzero terms, with identical products. This proves

$$\mathrm{Gas}_{\mathrm{new}}(U) = \mathrm{Gas}_{\mathrm{old}}(U) \exp \sum_j \ell_j(U).$$

The regular dictionary and its original coefficients remain the regular part of the fused representation. The attached logarithm is consumed only through the displayed exponential and the new forced factors. Gauge invariance and measurability pass to finite products and measurable eligibility indicators. Locality uses the union support. Widening a dictionary beforehand is simply an injective placement of its old labels in a universal dictionary with zero coefficients at all unused indices. It preserves the gas and the regular-log sum by reindexing the absolutely convergent families.

The fusion estimate for switched coefficients follows from the distinct-component forest bound already proved. Suppose the old switched forced coefficients have paid sitewise budget B_0 at support exponent (s+1), and the fixed attached logarithms have envelope budget w_0 there. Choose $\kappa \geq 0$. Assign inflated envelope $e^{\kappa} b_i$ to an old forced label and envelope $2c_j$ to an attachment with $|\ell_j| \leq c_j \leq 1$. The latter bound is justified by

$$|e^x - 1| \leq 2|x| \quad (|x| \leq 1),$$

which is used on the source's smaller admitted logarithm range. The combined sitewise budget is

$$\gamma = e^{\kappa} B_0 + 2w_0.$$

Each new connected label contains an old forced atom. Undoing its inflation supplies at least $e^{-\kappa}$; every other undone factor is at most one. Thus its paid envelope is bounded by $e^{-\kappa}$ times the product of the inflated atom envelopes. Its eligibility indicator is independent of the switches and does not increase the price. If $V_B \gamma \leq 1$, the earlier connected-component forest estimate, spending one support-exponent unit, bounds the new sitewise paid budget at exponent s by

$$e^{-\kappa}(e^{\gamma}-1).$$

The selected payment inequalities require this quantity to fit the requested output allowance. This proves transfer of the paid switching family through each fusion. The family's unmarked value is the fused unmarked coefficient by the same formula. An exterior-switch silence identity passes through fusion term by term: if an atom's union support omits the required parent site of a switch, each factor is unchanged on erasing that switch; so is their product and its eligibility indicator. These facts supply the value, price and scope properties needed after the two actual coarse dictionary changes.

The scalar normalization needed after these fusions is an equality of exponents, not a comparison of partition functions. Let $E_0(U)$ be the intermediate frozen exponent, let $\mathcal{R}(U)$ be the coupling-reset receipt, let α be its coefficient, and let $L_1(U), L_2(U)$ be the attached-log sums consumed by the two fusions. The path reset gives

$$E_{\mathrm{final}}(U) + L_2(U) = E_0(U) - \alpha\mathcal{R}(U).$$

The unmarked local integral normalization gives

$$\exp\{E_0(U) + L_1(U) - \alpha\mathcal{R}(U)\} = \mathrm{overhead}(U)Z_0(U).$$

Combining these equations and the exponential addition law proves

$$e^{E_{\mathrm{final}}(U)}e^{L_1(U)}e^{L_2(U)} = \mathrm{overhead}(U)Z_0(U).$$

In particular both sides use the same unmarked local gas Z_0 that appears in the switched formula. Subtracting its logarithm in the switched regular increment cancels precisely this normalizer. The path-reset identity and unmarked local normalization proved in this chapter retain the stationary path and vacuum choices used here.

7.9 The coupling-reset exponent identity

The reset changes the represented coupling and compensates that change by explicit local counterterms. The total density is preserved by the displayed exponent identity. Its projected coefficient is read from the affine Hessian with the source's chosen normalization. A probe value is a different operation and is not silently identified with that projection.

We now specify the reset in the scalar identity. Let v be the source on a listing X, let R be the same path family used in the kinetic energy, and let p fix the threshold, support scale and Taylor order. All masks below are the actual healthy-field masks at that threshold. Write $\Delta = \delta(\mathrm{read}_X(v))$, and let α be the already committed loop coefficient. The new coupling and vacuum scalar are

$$\beta' = \beta + \alpha + \Delta, \qquad \tau' = \tau - \frac{1}{|\mathrm{Site}|}\sum_i \mathrm{logRead}_i(v;1).$$

The receipt $\mathcal{R}(U)$ is the averaged sum of the masked plain low path terms, with their original paths. It is not obtained by changing paths to match a later choice.

The regular-source decomposition separates each log read into its value at the identity, a local projected payout, a fault, and a masked rebate. Absolute summability of the original paid activities bounds these four series respectively by the original total bill and fixed multiples of it. Summing

the payouts first at their anchor sites gives $\Delta \sum_o \mathrm{heads}_p(o, U)$, by the calibrated projection identity. Consequently

$$\mathrm{additive}_p(\mathrm{read}_X(v); U) = \sum_i \mathrm{logRead}_i(v; 1) + \Delta \sum_o \mathrm{heads}_p(o, U) - \sum_i \mathrm{dust}_i(U) + \sum_i \mathrm{untagged}_i(U).$$

The projected-payout identity uses the simple-group projection hypothesis and the stipulated period $n > 20$.

There are also low and high path corrections. Denote their averaged sums by F_Δ and $H_{\alpha,\Delta}$, and the low loan sum by $L_{\alpha,\Delta}$. Their defining termwise identities, summed with the same masks, are

$$(\alpha + \Delta)K_R + F_\Delta + H_{\alpha,\Delta} = \alpha\mathcal{R} + \Delta \sum_o \mathrm{heads}_p(o, \cdot) + L_{\alpha,\Delta}.$$

Indeed the true kinetic base is the sum of live low terms and high summits. The high correction is minus $\alpha + \Delta$ times the summit, so those terms cancel. For each low term its defining loan identity reads

$$(\alpha + \Delta)\,\mathrm{live} + \mathrm{fade}_\Delta = \alpha\,\mathrm{maskedPlain} + \Delta\,\mathrm{sketch} + \mathrm{loan}_{\alpha,\Delta}.$$

The sketch sum is the head sum. The low live, masked plain, proxy and fade series are absolutely summable by the path bounds; this identity then also proves summability of the loans. The high summit series is absolutely summable as well. Averaging over the finitely many seats preserves these properties. Hence the cancellations are valid equalities of convergent sums.

The reset dictionary's additive sum is the untagged sum plus $F_\Delta + H_{\alpha,\Delta}$; its attached logarithm is the dust sum plus $L_{\alpha,\Delta}$. Substitution into the representation exponent $\tau'|\mathrm{Site}| - \mathrm{inChannels}(\beta') -$ additive, followed by the two displayed identities and the channel decomposition, gives

$$E_{\mathrm{reset}}(U) + \sum_i \log_i(U) = E_{\mathrm{frozen}}(U) - \alpha\mathcal{R}(U).$$

This is the path-reset identity used above, including its constant term and signs. Exponentiation gives the same identity for unnormalized evaluators. It holds even when the activity gas vanishes, since the gas multiplies both sides and is never divided out. The subsequent dictionary fusion absorbs the attached-log exponential by the already established fusion identity. The lower-level projection, live-base and loan formulas supply the mathematical content of this reset decomposition and must retain the same threshold and paths.

7.10 Gauge invariance at the flat germ

Let $F(u) = f(\mathrm{increment}(1, u))$. Regularity gives a twice continuously differentiable germ at $u = 0$. Constant gauge transformations act there by the linear map $u_e \mapsto \mathrm{Ad}_g\, u_e$, because exponentiation commutes with conjugation. Hence $DF(0)$, restricted to a single edge's color space, is an invariant linear form. Such a form vanishes. Indeed differentiating invariance along $\exp(tx)$ shows it kills every bracket $[x, y]$. Its kernel is an ideal; simplicity makes this kernel zero or the whole algebra. The zero case would make the form injective and force all brackets to vanish, contradicting non-abelianness. Summing single-edge directions proves $DF(0) = 0$. Twice differentiating the same constant-gauge identity proves adjoint invariance of $D^2F(0)$.

For a color-valued site field ϕ, let $(\nabla\phi)_{x,i} = \phi(x + e_i) - \phi(x)$. Gauge the link exponential of u by the site elements $\exp(-t\phi(x))$, and take its local logarithm. This defines a smooth chart $C(u, t)$ near zero with

$$C(0, 0) = 0, \qquad DC(0, 0)[u, t] = u + t\nabla\phi.$$

The derivative follows by differentiating the three factors on each link: the contributions are $-t\phi(x), u_{x,i}, +t\phi(x+e_i)$. The local logarithm recovers the gauged configuration near the identity. Gauge invariance gives $F(C(u,t)) = F(u)$. Differentiate this equality in directions $(0,1)$ and $(u,0)$. In the second chain rule the term involving D^2C is multiplied by $DF(0) = 0$, while the right side has no t-dependence. Thus

$$D^2F(0)[\nabla\phi, u] = 0.$$

Its symmetric counterpart follows from symmetry of the second derivative. Locality implies that Hessians agree when the tested direction fields agree on the support: first compare their diagonal exponential lines, then polarize. These arguments concern smooth unmasked germs throughout.

The color-bilinear assertion used below also admits a direct finite-dimensional proof. If B is adjoint invariant, differentiation gives $B([x,y],z) = -B(y,[x,z])$. Its antisymmetric part $A = B - B^{\mathsf{T}}$ satisfies $A([x,y],z) = A([x,z],y)$. Cycling this relation and using bracket antisymmetry makes this value equal to its negative, so it is zero. For fixed z, the linear form $A(\cdot,z)$ kills brackets, and the kernel argument above makes it zero. Hence B is symmetric.

Represent $B(x,y) = \langle Tx, y\rangle$ by a self-adjoint real operator T. Invariance and skew-adjointness of the adjoint action give $T[x,y] = [x,Ty]$. A nonzero real eigenspace of T exists by the finite-dimensional spectral theorem and is an ideal by this commutation identity. Simplicity makes it the whole color space. Thus $T = cI$ and $B = c\langle\cdot,\cdot\rangle$. Skew-adjointness here follows directly by writing the color inner product as the real matrix trace and cyclically moving the commutator factors. This proves the invariant-bilinear-form input without an assumption about a complex or infinite-dimensional color representation.

7.11 Reduction of affine tests to constant curl

Use centered lifted coordinates $z_i(x) = \mathrm{lift}(x-o)_i$. A linear jet is

$$A_i(x) = a_i + \sum_j z_j(x) b_{ij}, \qquad F_{ij} = b_{ji} - b_{ij}.$$

On a support satisfying 10 $\mathrm{dist}(x,o) < n$, with $n \geq 4$, the centered lift has its ordinary increment under one positive or negative lattice step. This follows from uniqueness of the centered representative: both the old lift plus that step and the new lift wrap to the same site and stay within the centered coordinate range. Thus the discrete curl of A on this support is F.

If $F = 0$, the slope matrix is symmetric. The color-valued polynomial

$$\phi(x) = \sum_i z_i(x)(a_i - \tfrac{1}{2}b_{ii}) + \tfrac{1}{2}\sum_{i,j} z_i(x) z_j(x) b_{ij}$$

has forward difference $\phi(x+e_i) - \phi(x) = A_i(x)$ there. The diagonal correction removes the extra half-diagonal term in the quadratic difference; symmetry combines the remaining cross terms. Hence every curl-free affine field agrees on the tested support with a gauge gradient. Locality of the germ transfers its Hessian to agreeing directions, and gauge invariance annihilates its pairing with that gradient. Two affine fields with equal curls differ by this construction, so the Hessian depends only on the two curls.

An explicit primitive for a flux is obtained by adding the wedge fields: for a positive plane $a = (i,j)$, its wedge has only component j, equal to $z_i w_a$. Its curl is w_a in plane a and zero in every other positive plane. Summing therefore gives a primitive with exactly the prescribed flux. Its threshold-scaled unit-color wedge is an admissible paid direction, since $|z_i(x)| \leq \mathrm{dist}(x,o)$ and the permitted direction growth dominates $1 + \mathrm{dist}(x,o)$. The average of two such scaled wedges is admissible by the triangle inequality.

For a signed permutation with signs ε_i and permutation π, pullback of an affine jet changes its slope to $b'_{ij} = \varepsilon_i\varepsilon_j b_{\pi i,\pi j}$. Its constant term is

$$a'_i = \varepsilon_i\big(a_{\pi i} - \mathbf{1}_{\{\text{axis } i \text{ reversed}\}} b_{\pi i,\pi i}\big).$$

The subtraction is necessary because a reversed positive link is read from the shifted base. Direct substitution of the lifted-coordinate increment proves this formula on the nonwrapping support. Its curl is the signed plane reframe of the old curl. Pullback of exponential link increments is a linear pullback of their direction fields; the Hessian chain rule at the critical identity therefore gives precisely the signed-plane covariance claimed for the projected form. No global identification of a nonperiodic affine polynomial on the torus is used.

7.12 Affine-Hessian coefficients and their bounds

Fix a unit color vector ξ, which exists because the simple gauge algebra is nonzero. For each positive coordinate plane a and color x, let $W_a(x)$ be the affine wedge field with constant curl x in that plane and zero curl in the other planes. It is linear in x and commutes with adjoint color transformations. For the regular germ f at the identity, define

$$c_{ab}(f, o) = D^2 f(1)[W_a(\xi), W_b(\xi)].$$

For fixed a, b, the bilinear form $B(x, y) = D^2 f(1)[W_a(x), W_b(y)]$ is invariant under simultaneous adjoint transformations, by gauge invariance of the germ. Differentiation makes it invariant under the infinitesimal adjoint action. The simple-algebra invariant-bilinear-form result therefore makes it a scalar multiple of the color inner product. Evaluating at ξ, ξ fixes the scalar, giving

$$D^2 f(1)[W_a(x), W_b(y)] = c_{ab}(f, o)\langle x, y\rangle.$$

Since the primitive of a flux is the sum of its wedge fields, bilinearity gives

$$D^2 f(1)[\text{primitive}(w), \text{primitive}(z)] = \sum_{a,b} c_{ab}(f, o)\langle w_a, z_b\rangle.$$

The invariant-bilinear-form argument and the reduction of affine Hessians to curls are the simple-algebra and gauge-calculus inputs of this formula.

The constants in the projection estimate follow from actual normalized directions. The direction $tW_a(\xi)$ is admissible in the paid second derivative bound; hence $t^2|c_{aa}| \le c_f$. So is $u = \frac{t}{2}(W_a(\xi) + W_b(\xi))$. At the critical identity, its repeated ordered left derivative equals its diagonal Hessian, and

$$4D^2 f(1)[u, u] = t^2(c_{aa} + c_{bb} + 2c_{ab}).$$

Bounding the derivative on the left by c_f and each diagonal term by c_f yields $t^2|c_{ab}| \le 3c_f$. These estimates use only the paid directions at the identity, rather than a limiting periodic Fourier test.

The actual projected matrix is c_{ab} if the enlarged anchored support is nonwrapping, and zero otherwise. Its entry bounds hold in both cases. On a nonwrapping support it has exactly the flux-Hessian formula above. This cutoff is preserved by signed rotations and translations, as is the enlarged radius-two support. The pullback formula for flux Hessians then gives covariance of its quadratic form under the signed plane action, including the zero case.

The paid calculus used for these statements bounds all ordered left derivatives up to p.order in the specified anchored direction class, on the closed healthy set. The underlying functions remain smooth on full neighborhoods of each such configuration. Thus it differentiates the unmasked germs,

not their discontinuous indicators. Constant multiplication multiplies the bound by the absolute constant; sums and differences add bounds. For an inner product the ordered Leibniz rule has at most $2^{p.\mathrm{order}}$ splittings, giving

$$\mathrm{pairingCost}(E,k) = 2^k \|\mathrm{innerSL}_E\|.$$

Linearity of the derivatives and Hessian under subtraction follows on these full neighborhoods; enlarging the healthy support only restricts the configurations tested. These facts justify the residual and clover bills, including boundary points of the healthy sets.

7.13 Isotropy of the summed coupling projection

Let A be the actual countable activity listing, with its regularity, paid bound and label permutations implementing every signed lattice motion. Define the matrix at anchor o by

$$T_o(a,b) = \sum_{\{i:\,\mathrm{anchor}(i)=o\}} \Pi_i(a,b).$$

These sums are absolutely convergent: the coefficient bound $t^2|\Pi_i(a,b)| \le 3c_i$ is dominated by the summable activity bill. Finite sums over plane types can therefore be interchanged with this series. Reindexing by the label permutation for a signed motion, using its exact anchor, support and value transformations, gives

$$Q_{T_{\mathrm{affine}(s,b)o}}(w,z) = Q_{T_o}(\mathrm{reframe}_s w, \mathrm{reframe}_s z).$$

In particular the nonwrapping cutoff in each projected entry has already been transported before the label sum is taken.

Test this equality with fluxes supported in single planes and equal there to the fixed unit color vector. Translation and a permutation taking the reference plane $(0,1)$ to any chosen plane show that every diagonal entry at every anchor equals $\Delta = T_0((0,1),(0,1))$. For two distinct positive planes, there is an axis occurring in one and not the other. Reflecting just that axis changes the sign of their mixed pairing and fixes their plane labels. Its coefficient is consequently its own negative and is zero. Therefore

$$T_o(a,b) = \Delta\,\mathbf{1}_{\{a=b\}}, \qquad \sum_{\mathrm{anchor}(i)=o} Q_{\Pi_i}(w,w) = \Delta \sum_a \|w_a\|^2.$$

Applying this with the actual clover flux, and the same singleton healthy indicator on every term at that anchor, proves the projected-payout identity used in the reset exponent. It would not justify replacing different individual support indicators by that singleton indicator; the faults above are precisely what account for that change.

The diagonal coefficient estimate gives the sharper bound $t^2|\Delta| \le c$, where c is the listing's anchored paid budget. Indeed the triangle inequality bounds the anchor-zero diagonal sum by its local mass, and that mass is at most c. The off-diagonal factor three is not spent here. Thus both the scalar and its bound are obtained at the original finite period from the original label orbits.

7.14 The local clover counterterm

The petal flux at an anchor o is the based clover: the signed average of the four oriented plaquette curvatures in each coordinate plane, all based at o. Every such oriented square touches o, and all its visited sites lie within distance two of o. Thus a healthy singleton anchor bounds each square chord

by t, and the petal flux reads only the radius-two clover. Its gauge transform is adjoint transport by the gauge value at o; signed lattice motions act by the corresponding signed action on plane indices.

For a word of length l, curvature is the linear color projector applied to its matrix word minus the identity. The projector norm is at most the inverse color scale. The paid word derivative bound therefore gives the curvature bill

$$t\,\mathrm{colorBill}(G,l,P,k), \qquad \mathrm{colorBill} = G.\mathrm{colorScale}^{-1}\mathrm{feeWalk}(G,l,P,k).$$

Adding the four clover terms with coefficients of absolute value $1/4$ gives this same bill with $l=4$ for each petal. The ordered product derivative rule for their inner product gives

$$\mathrm{pairBill}(G,p)t^2, \qquad \mathrm{pairBill} = \mathrm{pairingCost}(G.\mathrm{Color}, p.\mathrm{order})\,\mathrm{colorBill}(G,4,p.\mathrm{power},p.\mathrm{order})^2.$$

Here the pairing cost accounts for all ordered derivative splittings already used in the paid calculus.

For a real matrix q on plaquette types, the counterterm is

$$Q_q(U) = \tfrac{1}{2}\sum_{a,b} q_{ab}\langle \mathrm{petals}(U,o)_a, \mathrm{petals}(U,o)_b\rangle.$$

Its bill is at most $\frac{1}{2}\mathrm{pairBill}\,t^2\sum_{a,b}|q_{ab}|$. It is smooth, gauge invariant, supported on the radius-two clover, and linear in q.

Its normalization follows directly from affine tests. For the exponential link line $U(t)=\exp(tv)$, differentiating a based word at zero gives the signed link sum along that word. For a square this is the discrete curl

$$[v_j(x+e_i)-v_j(x)]-[v_i(x+e_j)-v_i(x)].$$

For the nonperiodic affine field associated to a linear jet, on a nonwrapping neighborhood this equals its constant curl. The signed clover average consequently has derivative equal to the jet's positive-plane flux w. The period $n>20$ ensures the clover neighborhood has the required affine-coordinate increments. Hence

$$Q_q(U(t))/t^2 \longrightarrow \tfrac{1}{2}\sum_{a,b} q_{ab}\langle w_a, w_b\rangle.$$

Gauge invariance for the simple group makes the identity a critical point of every regular germ. Taylor's formula then identifies the left limit with one half its diagonal logarithmic Hessian. Therefore the counterterm Hessian on affine directions is exactly the quadratic form q on their curls. Symmetry of the Hessian and polarization recover all mixed affine pairs from these diagonal tests. This establishes the counterterm normalization used in the residual decomposition.

7.15 Masked residual decomposition

For a regular term f anchored at o with support D, enlarge the support to $D^+ = D\cup\mathrm{halo}_2(o)$. Let q_f be the counterterm built from its projected affine Hessian and the local petal flux. Define

$$r_f = f - f(1) - q_f, \qquad d_f(U) = \begin{cases} 0, & D^+ \text{ is healthy},\\ f(1) + \mathbf{1}_{\{o\}\ \mathrm{healthy}} q_f(U) - \mathbf{1}_{D\ \mathrm{healthy}} f(U), & \text{otherwise.}\end{cases}$$

Then the exact identity is

$$\mathbf{1}_{D\ \mathrm{healthy}} f = f(1) + \mathbf{1}_{\{o\}\ \mathrm{healthy}} q_f - d_f + \mathbf{1}_{D^+\ \mathrm{healthy}} r_f.$$

On a healthy enlarged support all three smaller health tests hold and the equation reduces to the definition of r_f. On its complement the enlarged-support term vanishes and the equation reduces to the definition of d_f. This is why the constant $f(1)$ is retained even outside the original healthy support.

If f has paid supremum and derivative bound c_f, its projected matrix entries obey $|\Pi_f(a,b)|t^2 \le 3c_f$. The quadratic counterterm bill is thus at most $M_p c_f$, where

$$M_p = \tfrac{1}{2}q^2\, 3\,\mathrm{pairBill}(G,p).$$

Indeed its half-weighted double type sum has q^2 entries, and each petal pairing has bill $\mathrm{pairBill}(G,p)t^2$. It follows by the triangle inequality that both the rebate's paid bound and the fault's all-configuration absolute bound are at most $(2+M_p)c_f$. The value $f(1)$ costs at most c_f. The fault vanishes on a healthy enlarged support, and hence is a forced term with an actual failure witness.

The rebate is regular on D^+, vanishes at the identity and has zero affine Hessian: the projected counterterm has exactly the affine Hessian of f, while subtraction of $f(1)$ changes no derivative. This identification uses the nonwrapping support and period hypotheses of the affine-curl construction. Gauge invariance and measurability pass to the counterterm, rebate and fault. The fault uses the enlarged support's field data and its final-threshold flags, precisely as permitted for forced terms. Signed rotations and translations send the petal flux and the projected matrix covariantly, and send healthy sets by their actual affine images, so they commute with this decomposition. Projection and the counterterm are linear in f; consequently differences of rebates and faults are the rebate and fault of the difference. These facts justify applying the decomposition to the common parameterized source series in the reset.

7.16 Low-path atoms and the loan cancellation

The low indices are a plaquette type together with either one direct symbol or an ordered pair of path-row labels (i,j). For a direct index let its mark be $m = c.\mathrm{direct}$, its length zero, and its plain value the plaquette price. For a pair let

$$m = c_i c_j, \quad l = \mathrm{length}(s_i) + \mathrm{length}(s_j), \quad P(U) = \tfrac{1}{2}\langle s_i(U), s_j(U)\rangle.$$

Its analogue on a flux w is $\langle w_a, w_a\rangle/2$ in the direct case, and $\langle w_{\mathrm{plaq}(s_i)}, w_{\mathrm{plaq}(s_j)}\rangle/2$ in the pair case. The proxy $Q(U)$ is this analogue evaluated on the local petal flux at the anchor. The home H is the anchor's radius-B halo together with both true path stencils in the pair case; in the direct case it is just the halo.

Let h_H and h_o be the healthy indicators on this home and on the anchor singleton. The live value $J(U)$ is the clipped plaquette price for a direct index, and $\mathbf{1}_{\mathrm{kept}}\, mP(U)$ for a pair. The freshness and loan functions are exactly

$$F_\Delta = \Delta m(Q-P), \qquad L_{\alpha,\Delta} = \begin{cases} 0, & h_H = 1, \\ (\alpha+\Delta)J - \Delta m h_o Q, & h_H = 0. \end{cases}$$

The masked plain value is $h_H mP$. On a healthy home the selector is kept and the direct plaquette is below threshold, so $J = mP$. Also $h_o = 1$. Substitution gives

$$(\alpha+\Delta)J + h_H F_\Delta = \alpha h_H mP + \Delta m h_o Q + L_{\alpha,\Delta}.$$

On an unhealthy home the same equality is the defining formula for the loan, with both home-masked terms zero. This proves the low loan cancellation in all cases, including a healthy anchor inside an unhealthy home.

These functions have the required locality despite the selector's apparent infinite test. Outside the radius-B aperture its exponential threshold already exceeds the maximum chord two, by the decent-budget hypothesis. Inside it the plaquette corners lie in the aperture unless the test is automatically true. Thus the kept predicate is determined locally. Slot values read their true route and plaquette stencils, and proxy values read only the anchor's radius-two clover, contained in the home. Gauge conjugation preserves chords, prices, kept predicates and the inner products of the transported slots. All these functions are measurable; a loan is a finite combination of continuous functions and measurable healthy/kept indicators. Its possible extra flag read is exactly the final-threshold read allowed in the definition of the forced dictionary.

Write $t = p.\text{threshold}$ and $\rho = c.\text{relaxation}$. A kept slot of length l_i has norm at most $te^{l_i/\rho}$. Using the chord bound two and the clip kink inequality for the direct term, both cases give the all-configuration estimate

$$|J(U)| \le |m|\, 2t\, e^{l/\rho}.$$

On a healthy anchor the proxy has the derivative bill $\frac{1}{2}\text{pairBill}(G,p)t^2$. Consequently

$$|L_{\alpha,\Delta}(U)| \le |m| e^{l/\rho} \left(2t|\alpha+\Delta| + \tfrac{1}{2}|\Delta|\,\text{pairBill}(G,p)t^2\right).$$

A nonzero loan therefore has an actual unhealthy home, and its bound retains the path length and relaxation factor needed for its support charge. It is linear in (α,Δ), as is the freshness in Δ, so their difference estimates use the same formulas with parameter differences.

Finally the freshness has zero affine Hessian. Along an affine exponential link field, a slot divided by its scalar parameter tends to the corresponding constant plaquette curl. This follows by differentiating the finite link product; the two nonconstant affine edge terms cancel to the constant curl. On a nonwrapping home, both the plain value divided by the parameter squared and the proxy divided by that square tend to the same flux analogue. Their difference therefore has zero quadratic limit. The freshness is smooth in a healthy neighborhood, vanishes at the identity, and the affine Hessian limit identifies its quadratic coefficient with this limit; polarization gives zero mixed affine Hessians. The nonwrapping and period hypotheses are retained here, as in the reset projection argument.

7.17 Recovering the actual kinetic base from the path atoms

The low pair sum is the square of the original path row, with all ordered labels retained. If $f_i = c_i s_i(U)$ and $\sum_i \|f_i\| < \infty$, then

$$\sum_{i,j} \tfrac{1}{2}\langle f_i, f_j\rangle = \tfrac{1}{2}\Big\|\sum_i f_i\Big\|^2.$$

Absolute convergence follows by domination by $\frac{1}{2}(\sum_i \|f_i\|)^2$; continuity of the inner product then permits the two sum exchanges. Applying this to each row proves that the total live low terms are exactly the clipped direct term plus the kept shortening half-square. The pair selector is common to its whole row, so the assertion holds on both selector branches.

For flux analogues the same square identity gives half the squared norm of the actual zero-momentum symbol row. If $S_{ab} = R.\text{symbol0}(a,b)$, the path normalization is

$$c.\text{direct}\,\delta_{bi} + \sum_a S_{ab}S_{ai} = \delta_{bi}.$$

Expanding the symbol-row norm squares and applying this matrix identity proves

$$\sum_{\text{low } d} m_d \,\text{analogue}_d(w) = \tfrac{1}{2} \sum_a \|w_a\|^2.$$

All series converge absolutely by the weighted path coefficient bounds. Multiplying by the common healthy singleton indicator and evaluating on the petal flux is therefore legitimate.

The layer operation is the sum over translated signed views divided by the number of signed orientations; translations are summed, not probability averaged. For each view-dependent term, translation covariance changes its anchor-zero value to its value at the translated anchor. Reindexing the finite sites by the signed permutation then proves that the low live layer plus the cubic and quartic masked layers is exactly the stipulated kinetic energy. It also proves that the low proxy layer is $\sum_o \text{heads}_p(o, U)$: signed reframe preserves the sum of squared plane norms, and the orientation count cancels its inverse. Thus both identities used in the reset's path cancellation come from the original symbol normalization and actual path rows.

On an everywhere healthy configuration every masked plain low atom is unmasked. Their layer is then the original unclipped second term, namely the orientation average of $c.\text{direct price} + \|\ell\|^2/2$. This specialization does not remove masks in a configuration with any flags.

The high freshness terms retain the actual cubic and quartic atoms on their own footprints and have coefficient $-(\alpha + \Delta)$. An atom of arity $r = 3$ or 4 has paid bill its startup coefficient times t^r; since $0 < t \le 1$, both are bounded using t^3. Along any exponential link line every slot is first order, so the atom divided by the line parameter to the power r has a finite limit. Dividing instead by its square leaves a factor tending to zero. Taylor's formula at the critical identity then gives zero diagonal Hessian for each high freshness term, and polarization gives zero mixed Hessian in arbitrary directions. This proves the required vanishing without changing the cubic/quartic masks or attaching an extra boundary term.

7.18 The unmarked local scalar

The normalization in the reset uses the same local pool $C(V) > 0$ as the density expansion. Its trained connected-log series has sum $\log C(V)$ and is absolutely convergent by the paid log estimate. The incoming-to-current label equivalence preserves the actual union support. For each of these fixed labels use minus its stream as the regular coefficient on that support. The support-mask identity for the trained stream says that this masked coefficient is minus the corresponding trained log term. Reindexing its absolutely convergent series therefore gives the regular sum $-\log C(V)$.

Adjoin the literal path-score terms with their original supports. Their masked sum is

$$-\gamma \mathcal{R}(V) - \beta\big(\text{bill}_{\text{cubic}}(V) + \text{bill}_{\text{quartic}}(V)\big) - \text{bill}_{\text{loop}}(V).$$

Thus the combined regular dictionary has additive sum equal to this expression minus $\log C(V)$. No anchor is selected from an evaluated field: the anchor for each stream support is chosen once from that nonempty support, and its address is retained along the label equivalence.

All these complete sums are invariant under a signed translated view. For the path bills, reindex the finite signed-seat average. For the pool logarithm, reindex the trained connected labels by the exact incoming-label motion; its support and eligibility data move with it. The regular dictionary now uses a separate full-seat average: each pulled coefficient is multiplied by the reciprocal of the number of seats, including translations. The healthy mask is pulled with its actual support. Each view has the same complete sum, so summing these averaged coefficients preserves that sum exactly. This full-seat normalization differs from the orientation-normalized layer operation on kinetic paths.

The field-independent scalar token is the original sector prefactor times the incoming vacuum exponential times the fixed deducted chart factor. It is positive. Set $\tau = |\mathrm{Site}_c|^{-1} \log(\mathrm{token})$. The actual scalar overhead is then

$$O(V) = \exp\left\{ \tau |\mathrm{Site}_c| - \beta \, \mathrm{average}_{\mathrm{orientations}} \sum_i \mathrm{baseTerm}(V, i) \right\}.$$

This follows by separating the token from the exact clipped-base exponential in the overhead; the threshold in that base is the same one in the dictionary.

Insert the combined regular sum into the frozen exponent, whose channel decomposition contains exactly that base, the cubic and quartic bills at price β, and the loop bill at its stated price. The latter three cancel the path-score terms. The receipt correction $-\gamma\mathcal{R}$ cancels the receipt in the regular sum. The remaining exponent is $\log O(V) + \log C(V)$. Therefore

$$\exp\{E_{\mathrm{frozen}}(V) - \gamma\mathcal{R}(V)\} = O(V)C(V).$$

If an earlier dictionary fusion has already been performed, its attached-log sum is added to the exponent on the left, giving exactly the scalar normalization formula stated above. Positivity of $C(V)$, supplied by the convergent real log identity, justifies $\exp(\log C(V)) = C(V)$.

The forced gas in this same frozen representation is the canopy-pool sum. Its convergent density identity is $O(V)C(V)\,\mathrm{gas}(V) = \mathrm{push}(R; V)$. Multiplying the scalar equality by this gas proves the early evaluator identity for every field, including fields where the gas vanishes. This identifies the unnormalized push without dividing by its gas or its value at a selected field.

7.19 Pool averaging and the literal score series

The regular pool uses integrated averages of the single-forest activities. If every forest value has absolute value at most a common tariff, the probability average has the same bound, by the triangle inequality and cancellation of the finite averaging cardinality. Thus the original thin and thick tariffs also bound the averaged entries and deposits. The renewed reserve gives the connected-log convergence hypotheses for these averaged entries, with the same supports and compatibility relation. Applying the already proved real gas/log identity gives

$$\log C(V) = \sum_{\text{connected labels } q} \mathrm{Gas.term}(\mathrm{entriesPool}(V), q),$$

$$C(V) = \exp\left(\sum_q \mathrm{Gas.term}(\mathrm{entriesPool}(V), q) \right) > 0.$$

It is this averaged gas, not an average of logarithms of separate forest gases, that occurs in the scalar proof.

For the full density, split each nonzero molecule family into components with nonempty flags and components with empty flags, retaining the requirement of nonempty ground. The two subtype maps have an injective union map and recover every nonzero family. Thus the integrated full gas reindexes as a sum over thick families and thin families. With a thick family fixed, excluding the thin supports it meets amounts to setting those thin entries to zero. Their absolute bounds can only decrease, so the same log theorem applies. The omitted-cluster restoration identity factors this excluded gas as

$$C(V) \exp\left(- \sum_q \mathrm{missing}_A(q, V) \right).$$

Multiplying by the banked thick-family factor and summing gives the pool density identity used above. Throughout this step the forest averaging is performed at the level justified by disjoint-ground factorization, rather than assuming arbitrary products commute with a common forest average.

The literal path score can also be read directly term by term. A low score label consists of a signed translated seat and a low atom. Its coefficient is $-\gamma/|\mathrm{Signed}|$ times the atom's marked plain value in that view; its support is the view of that atom's full home. Pulling its healthy mask through the view gives exactly the corresponding summand of $-\gamma\mathcal{R}$. A cubic, quartic or loop score label consists of a seat and an atom label, with its actual transformed footprint. Its coefficient is respectively $-\beta/|\mathrm{Signed}|$, $-\beta/|\mathrm{Signed}|$ or $-1/|\mathrm{Signed}|$ times the unmasked atom. Applying the footprint mask produces its placed atom.

The absolute atom budgets justify summing these labels. Reindexing the seat's translation by its signed action identifies the seat sum with the orientation average of the sum over all anchors. This proves the exact score total used above, with no change of masks. Finally the channel decomposition is simply the expansion of its definition:

$$\mathrm{inChannels}(\beta) = \beta\,\mathrm{average}\sum \mathrm{baseTerm} + \beta(\mathrm{bill}_{\mathrm{cubic}} + \mathrm{bill}_{\mathrm{quartic}}) + \mathrm{bill}_{\mathrm{loop}}.$$

These are convergent equalities before any smallness estimate or cancellation of cubic coefficients is imposed.

7.20 Why a nonzero detector forces a warning

The quiet-watch condition at o says every fluctuation coordinate with anchor within B of o has $|z_i| \le \xi$ t/g. Since the lifted field is μ+gz, each color coordinate of its difference from μ is at most ξt. If the color dimension is d and $\|\mu_e\| \le Mt$, the elementary Euclidean bound by the sum of coordinate absolute values gives

$$\|(\mu + gz)_e - \mu_e\| \le d\xi t,$$
$$\|(\mu + gz)_e\| \le (M + d\xi)t$$

on those links. A watch on radius 2B gives watches of radius B at nearby sites, by the distance triangle inequality.

Choose a positive error margin τ and the fine/noise plateau bounds

$$2/\sqrt{c_{\mathrm{direct}}}(A/L^2 + \tau) < 1/4,$$
$$2/\sqrt{c_{\mathrm{penalty}}}(A + \tau) < W.$$

These choices are possible in the required order. First take $\tau = \sqrt{c_{\mathrm{direct}}}/32$; next choose L large enough that $A/L^2 \le \tau$; finally choose W larger than $2(A+\tau)/\sqrt{c_{\mathrm{penalty}}}$. The stationary maximum A is chosen before this block-length choice.

The physical local path estimate has three contributions to its discrepancy bill: the effect of the fluctuation away from μ, the protected mean bill, and the nonlinear path-approximation remainder. Choose $\xi > 0$ with $d\xi \le 1$ and with the first contribution at most $\tau t/3$. The source's coefficient bound makes this possible by taking ξ no larger than τ divided by one plus three times its displayed linear-response coefficient. The mean and approximation errors have reference-coupling order -1 and are eventually each at most $\tau t/3$: multiplying either by sqrt(reference coupling) tends to zero, while t is at least the inverse square-root coupling. Their sum is therefore at most τt.

Away from the mandatory bad-seed region, the clean-patch and grafted-field estimates apply to these same links. They place the likelihood site and pin tags on their plateau and identify

the entry-cost rows with the actual fine curvature and prediction-mismatch curvatures. Add the background guide bound $A/L^2 \cdot t$ to a fine curvature row and A·t to a noise row. The curvature comparison inequality then bounds the plaquette chord by

$$2/\sqrt{c_{\mathrm{direct}}}(A/L^2 + \tau)t < t/4$$

and the prediction mismatch by

$$2/\sqrt{c_{\mathrm{penalty}}}(A + \tau)t < Wt.$$

Hence the fine cutoff, noise cutoff, site tag and pin tag all equal one. The forced-charge detector also equals one there: any putative flag would be witnessed by a nearby plaquette with chord greater than t, contradicting the t/4 bound. Each nonmandatory detector is its bounded auxiliary cutoff times (one of these quantities minus one), so it is zero under the radius-2B watch. A mandatory detector is defined to be zero already. The active-link cutoff complement likewise vanishes when the lifted link lies within its plateau, using the mean bound and $d\xi \le 1$.

This proves the detector-to-warning implication from the local protected-mean, physical path-accuracy, clean-patch tag and curvature-comparison estimates, with their three error contributions kept separate. Those local geometric estimates remain dependencies of this paragraph; the Gaussian probability argument does not supply them.

7.21 The joint Gaussian warning estimate

Let z be a centered finite-dimensional Gaussian vector of covariance $A \ge 0$, and suppose $v^{\mathsf{T}}Av \le M\|v\|^2$ for every real v. No independence of its coordinates is required. For a coordinate set D and a sign choice $\varepsilon_i \in -1, 1$, let v have those signs on D and zero elsewhere. Then $\|v\|^2 = |D|$ and

$$E\exp(sv \cdot z) = \exp(s^2 v^{\mathsf{T}}Av/2) \le \exp(Ms^2|D|/2).$$

If $\varepsilon_i\ z_i \ge t$ at every $i \in D$, then $v \cdot z \ge t|D|$. Markov's inequality, for $s \ge 0$, bounds this orthant event by $\exp((-st + Ms^2/2)|D|)$. The union over its $2^{|D|}$ sign choices gives

$$\begin{aligned}&\mathbb{P}(|z_i| \ge t \text{ for all } i \in D)\\ &\le 2^{|D|}\exp((-st + Ms^2/2)|D|).\end{aligned}$$

This also covers a degenerate covariance and an empty coordinate set.

Warnings are located at sites o_j, Gaussian coordinates at sites g_i. Suppose at most m warning addresses and at most p coordinate addresses lie over any one site. A warning at o_j is witnessed by some coordinate g_i within distance R whose absolute value is at least t. For a finite warning set W, choose one such coordinate for every warning. There are at most $(V_R p)^{|W|}$ witness assignments, where $V_R = (2R+1)^4$ bounds the number of sites in a native torus ball. For an assignment let D be its distinct coordinate image. Then $|D| \le |W|$. Each coordinate can witness at most V_R m warnings, so $|W| \le V_R$ m|D|.

Choose $H \ge 0$ with $M/2 + HV_R m \le t$ and apply the preceding rectangle estimate with $s = 1$. Its exponential factor is at most $\exp(-HV_R m|D|) \le \exp(-H|W|)$, and $2^{|D|} \le 2^{|W|}$. A union bound over witness assignments proves

$$\begin{aligned}&\mathbb{P}(\text{every warning in } W \text{ has a witness})\\ &\le (2V_R p\exp(-H))^{|W|}.\end{aligned}$$

Thus shared witnesses and overlapping warning neighborhoods have been counted explicitly. They are not counted as independent Gaussian events.

For the actual warning family $R = 2B$, where B is the support width. Its address census is the sum of the fine plaquette types, two coarse bond types, coarse and fine site types, and fine-link itinerary types. With block length L it is bounded by

$$m = L^4|PlaquetteAxes| + |Axis| + |Axis| + 1 + L^4 \\ + itineraryDensity(L).$$

The Gaussian coordinate census p is the tubular coordinate count of that same block. The covariance upper bound M comes from its actual fluctuation covariance estimate. The warning event uses threshold $t = \xi\ t_{\text{field}}/g$, where $g = \beta^{-1/2}$. Band comparability gives t_{field}/g at least one half of the fine monitor P. Put $H = \varepsilon\sqrt{P}$, $\varepsilon = \xi/4$. Eventually $V_{2B}m \le \sqrt{P}$ and $M/2 \le \varepsilon P$, because the monitor dominates the fixed covariance cost and the polynomial width cost. Hence $M/2 + HV_{2B}m \le 2\varepsilon P \le t$. The preceding bound gives warning probability at most $E^{|W|}$, with

$$E = 2V_{2B}p\exp(-\varepsilon\sqrt{P}).$$

The monitor's growth makes this E smaller than every prescribed reference-coupling power, after the finite block and stage parameters are fixed. This is the rapidly decaying warning allowance used by the activity prices. Its scheduling uses the original monitor and covariance bound, with no reselection of a Gaussian for a given molecule.

The event is connected to the detector factors by a pointwise implication. A nonzero detector means the corresponding local quiet-watch condition fails. Unfolding that condition supplies a coordinate within distance 2B with $|z_i|$ at least $\xi\ t_{\text{field}}/g$. Thus a product with all its warning detectors nonzero lies in the joint event just bounded. The preceding plateau calculation proves this implication: under the quiet watch the curvature is at most one quarter of its threshold, the prediction mismatch is within its plateau, no nearby flag can occur, and every nonmandatory detector is a cutoff times a vanishing plateau complement. A mandatory detector is identically zero. Taking the contrapositive therefore gives precisely the witness used in the probability bound.

7.22 The primitive integrand and its local extension

The primitive integrand is a function of actual group links, chart variables, masks and insertion factors. On the protected interior its logarithm has the stated relation to the group-valued readout. Ordered group multiplication is retained throughout. In particular, from $\exp(\ell) = UB^{-1}$ one obtains $\exp(\ell)BU^{-1} = I$; there is no step that commutes B with U. The extension away from the protected interior retains its cutoff and discrepancy rather than claiming this identity on its full domain.

Here are the finite objects to which the integration estimates apply. Fix a coarse-to-fine step, its fine tree F, a seed inventory, a coarse configuration V, local discrete choices k, and the off-tree group variables Y. The active off-tree links are determined by the tree and seed inventory. For a color array u, the fine configuration has link $\exp(u_e)(\text{replica}\,V)_e$ on an active link, the value supplied by Y on an inactive off-tree link, and the fixed tree value on a tree link. The logarithmic signal equals u on active links, zero on tree links, and the chosen cutoff logarithm of $U_e(\text{replica}\,V)_e^{-1}$ on inactive off-tree links. Thus the group configuration and the logarithmic signal are different objects. Their identification by exponentiation is asserted only on the interior domain where the cutoff logarithm agrees with the genuine chart inverse.

A letter is of one of four disjoint types: a block packet A, an old forced index q, a regular logarithmic index λ, or a warning index w. A molecule p is a finite set of letters; repeated copies of a letter are not allowed. Its forced inventory is the set of its forced indices. Its seed inventory is the union of the seeds of its packet letters. The ground $H(p)$ is the union of the enlarged read footprints of all its letters. In particular it contains every seed and every forced read. Admission requires all of the following: the overlap graph of letter footprints is connected, the packet letters form the permitted layout of their combined seed inventory, the forced inventory satisfies its sealing and incidence conditions with those seeds, and

$$\text{flags}(V) \cap H(p) = \text{flags}(\text{seed inventory of } p). \tag{PRI1}$$

We spell out the discrete tests in this assertion. A seed is a member of the disjoint union of fine plaquettes, coarse bonds and coarse sites. These types remain distinct even when their spatial positions coincide. A store A is the triple of finite sets $(A_{\text{plaq}}, A_{\text{bond}}, A_{mflag})$; its number of seeds is the sum of their cardinalities. The read set of a plaquette seed is its coarse projected start set, that of a coarse-bond seed its two pin endpoints, and that of a site seed its singleton. The place of A is the union of these read sets; its dome $A^{[r]}$ is their coarse radius-r halo. Give the seed set the graph in which two seeds are adjacent when their radius-lB halos intersect. A connected packet means a nonempty connected set in this graph.

A permitted layout $\mathcal{L}$ for A at parameters l, B is a finite set of stores with the following three properties:

$$\begin{gathered} \text{each } A' \in \mathcal{L} \text{ is a connected packet,} \\ (A')^{[lB]} \cap (A'')^{[lB]} = \varnothing \quad (A' \neq A''), \\ \text{seeds}(A) = \bigcup_{A' \in \mathcal{L}} \text{seeds}(A'). \end{gathered} \tag{PRI6}$$

It exists by taking the connected components of the seed graph and converting each component back to its store. Conversely any permitted layout is this one: an edge cannot connect distinct members because their halos are disjoint, and each member is connected. Thus the packet layout is determined by the combined seed inventory. It does not introduce a multiplicity of choices or a choice of a graph witness.

For a finite set I of old forced indices write P_i for its actual fine support and $J_i \subset P_i$ for its declared inventory. Sealing means $P_i \cap P_j = \varnothing$ for distinct $i, j \in I$. Incidence means

$$\forall i \in I,\ \forall y \in J_i, \quad \exists p \in A_{\text{plaq}} \text{ with } y \in \text{touched}(p). \tag{PRI7}$$

Here touched(p) is the fine set inspected for that plaquette seed in the original sector partition. This condition uses the inventory, whereas sealing uses the full supports. Their roles cannot be interchanged.

The selector used in the packet integrand is defined at each fine site y by

$$c_I(U; y) = \begin{cases} 0, & y \in \text{flags}_\vartheta(U) \setminus \bigcup_{i \in I} P_i, \\ 1, & \text{otherwise.} \end{cases} \qquad \mathcal{O}_{I,Z}(U) = \prod_{y \in Z} c_I(U; y). \tag{PRI8}$$

For the packet, Z consists of fine sites whose parents are inside its $300B$ dome in the prescribed sense. For the full gas take all fine sites. If a product of old forced factors is nonzero, their individual permission identities give $J_i = P_i \cap \text{flags}_\vartheta(U)$. On a sealed family, the full selector is one exactly when the union of those inventories is all flags. Otherwise it is zero. This proves the exact rewriting of the old gas occurrence as the sealed indicator times (PRI8) times its old forced product. If the product is zero, both expressions are zero without using a permission identity for a vanishing factor.

The footprint of a packet letter is its enlarged dome; that of every other letter is its prescribed inspected, projected halo. Define two letters to be adjacent when these coarse footprints intersect. The molecule ground is their union, as in (PRI1). Its admission is now completely specified by nonempty connectedness in this graph, the layout (PRI6) of its packet letters, sealing and incidence (PRI7) of its forced indices, and (PRI1). The empty seed inventory is allowed when there are no packet letters; its permitted layout is empty.

Compatible molecules have disjoint grounds. For such a finite family, admission of its union is equivalent to admission separately in every member, together with the global flag coverage. To prove the nontrivial localization of incidence, suppose a forced inventory site in one molecule is touched by a seed attributed to another. The parent of that site lies in the forced footprint of the first molecule and in the seed footprint of the second. Disjointness of grounds forces the two molecules to coincide. The same parent argument localizes the sealing test. Layout restrictions localize because the enlarged packet domes are contained in their grounds; joining the local layouts gives a global layout. Equation (PRI1) localizes by intersecting the global flag set with disjoint grounds. Finally every finite set of letters has a unique partition into connected components of its footprint graph. These statements prove that expansion into finite letter sets and regrouping into compatible connected molecules are inverse operations, with multiplicity one. Absolute summability proved below permits this finite combinatorial identity to pass to the countable alphabets. All these admission tests remain discrete masks and are not differentiated.

For the common lifted fluctuation $u = \mu + \beta^{-1/2} z$, the four vertex values are, respectively,

$$\begin{aligned} v_A &= 1_{\{\text{packet connected}\}} 1_{\{\text{packet eligible}\}} \mathcal{S}_A(u), \\ v_q = K_q(U(u))\, \mathcal{E}_{2a,q}(u), &\qquad v_\lambda = e^{L_\lambda(u,z)} - 1, \qquad v_w = D_w(u). \end{aligned} \tag{PRI2}$$

Here $\mathcal{E}_{2a,q}$ is the product of fluctuation cutoffs on the entire forced read footprint. Packet eligibility requires its place to lie in the chosen seed place and its declared flags to agree with the actual coarse flags throughout its $300B$ neighborhood. To define the profile factors, let χ be the fixed smooth cutoff and put $h_s(r) = \chi((r/s)^2)$. For a tested fine plaquette use $h_{t/4}(\mathrm{toll}(U,p))$ if $p \notin A_{\mathrm{plaq}}$, and its complement if $p \in A_{\mathrm{plaq}}$. For a tested coarse bond use $h_{Wt}(\mathrm{wire}(\mathrm{prediction}(U,k), V; b))$ if $b \notin A_{\mathrm{bond}}$, and its complement otherwise. Multiply these factors once over every plaquette whose projected start set meets $A^{[300B]}$ and every bond whose pin set meets that dome. These are the profiles in $[0,1]$ in the following formula.

Let n_F, n_P be the cardinalities of the finite frame-tag and predictor-point sets. At a site x, the ticket factor is n_F times the actual frame-choice probability at x. At a bond it is n_P times the actual predictor-choice probability for its connector. Multiply the site factors over $A^{[300B]}$ and the bond factors over all pin sets meeting it. They are nonnegative and bounded respectively by n_F, n_P, rather than by one. With the uniform average over the complete choice space,

$$\mathrm{likelihood}(U,k) = \frac{1}{|\mathcal{K}|} \prod_x \mathrm{siteMass}(U,k,x) \prod_b \mathrm{pinMass}(U,k,b). \tag{PRI9}$$

Indeed $|\mathcal{K}| = n_F^{|\Lambda_c|} n_P^{|\mathrm{Bond}_c|}$, which cancels the inserted cardinality factors. This identity also includes zero-likelihood choices; their ticket product is zero. The frame probability is uniform over the admissible minimizing tags when that set is nonempty, with the prescribed uniform fallback; the predictor probability is its prescribed nonnegative normalized weight with its corresponding fallback. Thus the ticket retains the actual choice denominators, including the fallback branch.

The packet value is explicitly

$$\begin{aligned} \mathcal{S}_A(u) = e^{-\beta(Q_{\mathrm{near}}(A,u)+Q_{\mathrm{far}}(A))+\sum_{\mathrm{visit}_{200B}\,A} c_e+\sum_{\mathrm{active}_{304B}\,A} j_e(u)-N_A \log m_\beta} \\ \times \mathcal{O}_A(U)\, \mathcal{T}_A(U,k)\, \mathcal{P}_A(U,k,V)\, \mathcal{E}_{a,A}(u), \end{aligned} \tag{PRI3}$$

where m_β = chart scale $(\sqrt{2\pi}\,\beta^{-1/2})^{\dim \mathfrak{g}}$, and N_A counts the off-tree frozen links in the $10B$ visit set. The terms c_e are the link condensation terms; j_e are the local logarithmic Haar and cutoff corrections. The far term is the signed-orientation average of the extracted nonnegative row costs on the $200B$ posted set. The remaining four factors are the forced-order selector, the finite layout ticket, the product of sector profile factors, and the active cutoff product on the $304B$ neighborhood. The selector, profiles and cutoffs are in $[0, 1]$; the ticket has its finite combinatorial toll bound. The normalization deficit in (PRI3) is per frozen link, and is retained before that toll is estimated.

The regular logarithms L_λ comprise the old regular activity and cubic, quartic, loop and quadratic-tail energy pieces, the remote row and local chart terms, and the finite ledger terms from Gaussian completion and covariance reweighting. Their signs are the signs in the exponent of the original density: in particular an old regular occurrence contributes its negative value. A warning vertex is zero when mandatory; otherwise it is its fluctuation cutoff times the corresponding sector-profile complement. The Gaussian warnings are also retained in this warning alphabet. All factors in (PRI2) use this single configuration, signal and Gaussian variable.

Put $b_p(Y, z) = \prod_{\alpha\in p} v_\alpha(Y, z)$. The primitive value before the finite choice average is the literal integral

$$a_p = \int b_p(Y, z)\, d\nu_F(Y)\, d\gamma_C(z), \tag{PRI4}$$

with the product tree-gauge probability measure and the one covariance C prepared for this sector. There is no normalization by the mass of an admissible event. Local discrete choices are subsequently averaged with weight $1/|\mathcal{K}|$; their probabilities do not replace any of the profile or ticket factors.

For extension to coarse configurations outside the healthy domain, use the same finite formulas with the extended local row expressions and retain the admission masks at their actual enlarged output footprints. For a chosen ground H, retain only the finite Gaussian coordinates $I(H)$ read there. The submatrix C_I is positive definite at an actual admissible sector and has density

$$(2\pi)^{-|I|/2} \exp\Big(-\tfrac{1}{2}\log\det C_I - \tfrac{1}{2} z^T C_I^{-1} z\Big)\, dz. \tag{PRI5}$$

Integrate the extended product against this density and the tree-gauge measure, then take the uniform finite choice average. Set it to zero if the structural connectedness, layout or incidence tests fail. The marginal extension has its own fixed support and is not presumed positive or energetic at other coarse configurations.

On a healthy ground this agrees with (PRI4). The integrand is constant on fibers of the coordinate projection to $I(H)$; its Gaussian pushforward has covariance C_I. The remaining local data agree on the ground. Integrating over these fibers and using (PRI5) proves the identity. The same finite-choice fiber argument proves independence of choices outside H. For disjoint grounds the finite choice space is a product; writing its cardinality as the product of the cardinalities of the factors proves that the average of their product is the product of their averages. The tree variables have the analogous product-measure identity. No independence of overlapping Gaussian warning events is used.

The later continuity passage uses these same primitive integrals. Along an ultrafilter all finite seed stores and discrete tags can be fixed on one member of the ultrafilter. Keep the tree-gauge input Y fixed, and compare the two histories on this fixed finite data. Their covariance densities and vertex factors have the asserted pointwise limits. The primitive bounds below give a common integrable majorant, so dominated convergence applies to each molecule. Since this holds along every ultrafilter tending to infinity, the difference of the primitive values tends to zero in the ordinary sequential sense. This explains why the fixed finite data in the continuity argument can be selected

after fixing Y, without claiming that a continuously varying configuration has a globally constant sector.

7.23 Integrating the primitive vertex bounds

The primitive molecule body is a finite product of four kinds of vertices: connected block packets, old forced factors, regular exponential increments, and warning detectors. Fix a molecule p and let $B(p)$ be its warning subset. The prepared estimates give nonnegative pointwise tariffs f_i for each vertex and integrated tariffs t_i, with

$$\prod_{i\in p} t_i = h^{|B(p)|} \prod_{i\in p} f_i.$$

Only warning tariffs receive the factor h. The packet tariff is zero unless the packet is connected in the specified enlarged geometry. Otherwise it is its displayed block-packet income. Forced tariffs are the old forced envelopes, regular tariffs bound both the local log and exp(log)-1, and warning pointwise tariffs are the fixed storm bound. The source proves these four pointwise inequalities on the same local data and sector.

Write $F_p = \prod_{i\in p} f_i$. The triangle inequality for a finite product gives $|\operatorname{body}_p(Y,z)| \le F_p$ for every tree-gauge variable Y and Gaussian fluctuation z. If $\operatorname{body}_p(Y,z)$ is nonzero, every warning detector in $B(p)$ is nonzero. The warning estimate then places z in the joint catastrophe event $C_{B(p)}$, whose Gaussian probability is at most h^$|B(p)|$. The event and its covariance are common to all the factors; no factorization of its probability is used. Therefore pointwise

$$|body_p(Y,z)| \le F_p 1_{C_{B(p)}} z.$$

The body is measurable, and the uniform finite bound makes it integrable against the product of the probability tree-gauge measure and the reference Gaussian measure. Integrating the displayed indicator majorant and then integrating Y proves

$$\begin{aligned}&\int |body_p(Y,z)| d\nu Y d\gamma z\\ &\le F_p \gamma C_{B(p)} \le \prod_{i\in p} t_i.\end{aligned}$$

This proves the primitive integrated tariff from the four vertex bounds and the joint catastrophe estimate. In particular an empty warning set costs no probability factor. Correlated warnings are paid by their actual joint event.

For a marked team, remove the team's selected keys from p and multiply that remaining body by its paid ingredient product. Each paid leaf ingredient has absolute value bounded by its assigned leaf weight, so their product is bounded by the team weight W_M. If the team keys are not all in p, define the term to be zero. Otherwise the preceding argument gives the marked bound

$$\int |markedBody_{M,p}| \le W_M \prod_{i\in p\setminus \mathrm{keys}(M)} t_i.$$

The corresponding uniform pointwise bound uses the f_i instead of t_i. Taking these products over all finite molecule subsets is summable whenever the primitive tariff family is summable: the sum is bounded by W_M times the product of $1+t_i$, hence by $W_M \exp(\sum t_i)$, with finitely many compulsory keys held present. This supplies the summable integral-norm bound used in the primitive Mark expansion. Its remaining analytic inputs are the block-packet bound, old forced and regular-factor envelopes, warning storm bound and joint Gaussian catastrophe estimate.

7.24 Factoring the regular gas and restoring forced exclusions

At one Mark step split primitive activity labels into regular labels with empty flag inventory and forced labels with nonempty flag inventory. Fix a compatible forced set A, and let Y_A be the union of its enlarged supports. The allowed regular gas in the presence of A removes the regular labels meeting the forced exclusion region. Denote the full regular gas by Z_0, and its convergent cluster logarithm by L_0. The removed logarithmic terms are precisely regular occurrence clusters whose union meets that region. If their sum is M_A, the restricted gas is

$$Z_{0,A} = \exp(L_0 - M_A) = Z_0 e^{-M_A}.$$

This is an equality of absolutely convergent series: cutting the regular label family decreases its absolute budgets, and subtracting its cluster sum from the full cluster sum leaves exactly the indicated missing clusters. It retains their original occurrence indices and factorials.

Multiply by the forced product and sum the compatible forced sets with the required total inventory. Expand e^{-M_A} by positive multiplicities of missing-cluster addresses. The resulting term consists of the forced set and a finite multiplicity family of negative regular clusters; its weight is the product of the forced activities and the cluster factors $(-h)^r/r!$. Each selected missing cluster intersects the forced exclusion region. Absolute summability follows by the restoration bound, or at fixed volume by the summable exponential of the absolute missing-cluster series. Therefore grouping these terms does not change their sum.

Decompose every such term into overlap-connected components of its atom supports. Every component containing a missing cluster also contains a forced atom: that cluster meets the union of the selected forced supports. A component consisting only of forced atoms is already allowed. The resulting components have pairwise disjoint supports; their forced inventories partition the original inventory. Conversely, union of any permitted disjoint collection of these components recovers the original forced set and each missing cluster with its multiplicity. The multiplicity factorials factor over components because no atom address belongs to two distinct components. Thus the map is bijective and preserves the product weights.

The original gas, multiplied by its external scalar prefactor, is consequently Z_0 times the gas of these connected restored forced components. This proves the regular-log factorization and forced regrouping identity once the primitive Mark expansion has been supplied. It uses exact inventories and actual overlap exclusions; the more permissive repetition count used for the absolute bound does not alter this equality.

For switched leaves the same argument holds at each real switch value $0 \le x \le 1$, with the same supports, threshold masks and inventory tests. The paid bounds established above apply uniformly there. Two subsequent dictionary fusions put the restored components into the final coarse dictionary. Each fusion preserves the gas up to the exponential of its recorded unmarked regular log contribution. These extra logarithms do not depend on the switch selection. The unmarked normalization identity says that the final coarse exponent, its fusion logarithms, and $L_0(0, \varnothing)$ reproduce the original scalar prefactor. Substitution into the factored switched gas identity yields

$$e^{E_A(U)} e^{L_0(x,R) - L_0(0,\varnothing)} \operatorname{Gas}_A(K(x,R);U) = \mathcal{P}(\operatorname{Buy}(x,R))(U).$$

All divisions in the regular-gas ratio are legitimate because the real gas/log identity proves $Z_0 > 0$, including the unmarked branch. The primitive Mark expansion, dictionary fusion and unmarked scalar normalization proved in this chapter establish the scalar equality, with the stated exclusion restoration and component regrouping.

7.25 Paid regular and forced cluster terms

A switched primitive activity is a finite ticket sum. Each ticket carries switch dependence $x^{|D|}\mathbf{1}_{D\subseteq R}$ and a scalar factor bounded by the product of its marked leaf weights and all unmarked activity weights not selected by the ticket. Its switching model is concentrated at D. Summing tickets gives a paid bound equal to the product of the augmented activity tariffs. The selected mixed-activity estimate enlarges this bound by its displayed support and polynomial-moment weights; it sets it to zero when the activity's eligibility conditions fail. These are pointwise paid bounds uniform in configurations.

A regular cluster term is the product of its n switched primitive activities, multiplied by $\varphi(z)/n!$. The product rule for paid models gives the bound

$$\frac{|\varphi(z)|}{n!}\prod_{j=1}^{n} u_{i_j},$$

where u_i are the selected primitive paid bounds. The rooted cluster estimate just proved gives a support-weighted sitewise sum at most $P_0/(1-\rho)$ for their regular cluster envelopes, when the selected primitive budgets are P_0 and the tree ratio is $\rho<1$. Subtracting each cluster's unmarked value doubles its paid bound. These differences are the regular log increments in the switched identity; they are not independently selected regular activities.

The forced restoration terms combine forced primitive activities and negatives of regular cluster terms, each with its explicit occurrence multiplicity. For an atom occurring $r\geq 1$ times the factor is $h^r/r!$; different atom factors are multiplied. The paid product and power bounds therefore give exactly the corresponding product of envelope powers divided by factorials. Restrict to connected terms with at least one forced atom and the required flag inventory. These restrictions depend on configuration and on the fixed combinatorial term, not on the switching parameter, so multiplying by their indicators preserves the paid bounds. The omitted-cluster restoration and primitive-integral arguments in this chapter control their support-weighted sum.

7.26 Restoring omitted clusters with positive multiplicities

The restoration exponential uses a different indexing from the occurrence tuples in the logarithm. Its index is a finite set D of atom addresses together with a positive multiplicity $r_i\geq 1$ at every $i\in D$. Its weight and total size are

$$W_h(D,r)=\prod_{i\in D}\frac{h_i^{r_i}}{r_i!},\qquad N(D,r)=\sum_{i\in D} r_i.$$

An absent atom is absent from D, rather than being included with multiplicity zero. For nonnegative $h_i\leq\sigma<1$, summing each positive multiplicity gives

$$\sum_{r_i\geq 1}\frac{h_i^{r_i}}{r_i!}\leq\frac{h_i}{1-\sigma}.$$

Indeed $r_i!\geq 1$ and $h_i^{r_i}\leq h_i\sigma^{r_i-1}$. Multiplying these inequalities over a fixed distinct set D reduces the repeated-atom bound to a distinct-set bound with inflated vertex weights $h_i/(1-\sigma)$. The sitewise smallness hypothesis guarantees each required $h_i\leq\sigma$: choose a site in its nonempty support and retain its single summand in that site's budget. All extra support factors and $(1-\sigma)^{-1}$ are at least one.

The distinct connected-set tree estimate then bounds the sum of repeated connected terms whose union contains a fixed site by $\eta/(1-q')$, provided the inflated vertex activities satisfy the readiness condition with budget $\eta \le \sigma$ and tree ratio $q' < 1$. The actual union support weight is paid by the product of its vertex support weights: connected cover subadditivity applies to D, and each multiplicity $r_i \ge 1$ only increases the available exponent sum. Nonnegative summation justifies all these successive sums.

For omitted-cluster restoration take two kinds of atoms. Forced atoms have nonnegative weights d_i, nonempty supports Y_i, and sitewise budget c at exponent $A + \delta'$. Regular-log atoms are occurrence-indexed clusters, with their absolute cluster weights and union supports. The preceding logarithmic tail estimate gives their sitewise budget $\gamma/(1-q)$ at that exponent. Their combined repeated products bound the absolute value of every signed restoration product; forgetting original exclusions and retaining additional repeated forced atoms only enlarges the upper bound.

To obtain both a size tail and a compulsory-forced-atom gain, choose $t \ge 1$ and $\mu \ge 0$. Multiply every atom weight by t, and every forced atom weight additionally by e^{μ}. After the geometric multiplicity allowance, the combined sitewise budget is at most

$$\frac{te^{\mu}c}{1-\sigma} + \frac{t\gamma}{(1-\sigma)(1-q)}.$$

Require this to be at most $\eta \le \sigma$, with $0 \le \sigma < 1$, $\delta' > 0$, and $4(2B+1)^4\eta/\delta' \le q' < 1$. Apply the preceding repeated connected estimate to these augmented activities. A term of total size $N \ge m$ acquires the factor $t^N \ge t^m$. If it contains a forced atom, the exponent summed with multiplicities acquires at least μ, so its additional factor is at least e^{μ}. Dividing the augmented bound by these factors proves

$$\sum_{\substack{\text{connected } (D,r),\ N(D,r)\ge m \\ \text{at least one forced atom} \\ x\in\cup_{i\in D}X_i}} e^{Ac_B(\cup_{i\in D}X_i)} W_h(D,r) \le t^{-m}e^{-\mu}\frac{\eta}{1-q'}.$$

This is the omitted-cluster restoration bound used for the forced paid terms. Its applications use $t = 2$, $\sigma = 1/2$, and $q' = 1/2$; at $m = 0$ the right side becomes $2e^{-\mu}\eta$. The compulsory atom is charged at least once, regardless of its multiplicity. The remaining input in this argument is the distinct connected-set tree estimate under the stated readiness condition, which belongs to the same rooted-tree bounds as the regular cluster estimate.

7.27 Switching real source formulas into the marker algebra

The nilpotent substitution in the marked expansion is defined through real switching values, not by a formal differentiation of an unbounded series. For a finite switch set I, a real coefficient array $g_x(A)$, $A \subseteq I$, gives switch values

$$f(x,R) = \sum_{A\subseteq R} x^{|A|} g_x(A), \qquad 0 \le x \le 1.$$

At $x > 0$ the array is uniquely determined by these values. Induct on (|R|): subtract all proper-subset terms from (f(x,R)), then divide by $x^{|R|} > 0$. Equivalently the inverse is the finite Möbius formula

$$g_x(A) = x^{-|A|} \sum_{R\subseteq A} (-1)^{|A|-|R|} f(x,R).$$

A model for f is such an array with each $g_x(A) \to g_0(A)$ as $x \downarrow 0$. The switching jet is $J_f(A) = g_0(A)$, independently of the model by uniqueness for $x > 0$. A paid model includes nonnegative bounds (b(A)) with $|g_x(A)| \le b(A)$ for $0 \le x \le 1$ and $\sum_A b(A) \le c$, uniformly over any additional configurations. In particular $|f(x, R)| \le c$ on that interval.

For square-zero elements z_i of the finite marker algebra, define

$$\mathcal{S}_z(f) = \sum_{A \subseteq I} J_f(A) \prod_{i \in A} z_i.$$

This substitution is additive on modeled families. To prove multiplicativity, multiply the two switch-value formulas. Their coefficient convolution is

$$(g *_x h)(A) = \sum_{B \cup C = A} x^{|B \cap C|} g(B) h(C).$$

At $x = 0$ only disjoint pairs survive. On substitution, every overlapping pair also vanishes because some z_i occurs twice. Hence $\mathcal{S}_z(fh) = \mathcal{S}_z(f)\mathcal{S}_z(h)$. Constants substitute to their scalar polynomial. Paid bounds multiply by the corresponding product bound, and addition adds their bounds.

Countable sums pass through this operation when their paid bounds c_n are summable. For each subset A, the corresponding coefficient arrays are dominated by summable bounds, so dominated convergence for series gives their limit at $x = 0$ and commutes it with summation. There are only finitely many A; applying the finite linear substitution proves

$$\mathcal{S}_z\left(\sum_n f_n\right) = \sum_n \mathcal{S}_z(f_n).$$

Apply this to the exponential series. The n-th term has paid bound $c^n/n!$, which is summable. Multiplicativity gives

$$\mathcal{S}_z(e^f) = \exp\{\mathcal{S}_z(f)\}.$$

If the scalar mode of $\mathcal{S}_z(f)$ is zero, the right side is its finite nilpotent exponential. This establishes the exponential identity without dividing by a density.

The price estimate is coefficientwise. If $|z_i|_w \le 1$, define the envelope of mode A to be the sum, over switch subsets D, of (b(D)) times the absolute coefficient of $\prod_{i \in D} z_i$ in that mode. The triangle inequality bounds the substituted coefficient by this envelope. Summing its weighted price and using submultiplicativity gives at most $\sum_D b(D) \le c$. Thus a paid scalar switching model produces the local polynomial price required by restoration.

When every z_i has zero scalar mode, the scalar mode of $\mathcal{S}_z(f)$ is $f(0, \varnothing)$. Its singleton marker coefficient receives contributions only from switch subsets of size one; products of two or more zero-scalar factors cannot have singleton degree. These two facts prove the scalar and singleton coefficient readings used below.

Locality also passes at the coefficient level. If changing switch i leaves a scalar switching family unchanged, its Möbius coefficients indexed by subsets containing i vanish. If the output support omits a site in the parent support of leaf i, the local silence identity gives exactly that switch independence. A nonzero coefficient can therefore involve only leaves whose required parent supports lie in the output support. Their scoped marker monomials then have all their pins there. The finite substitution preserves gauge invariance, locality and measurability. If every scalar switching value vanishes off a prescribed healthy region or inventory event, the Möbius coefficients and their limits

vanish there as well. These statements supply the support and permission properties of the marked corrections.

7.28 Reading the marked Mark-step output

The reserve used by a retained step is constructed for the very finite leaf inventory just described. Give each leaf its fee and normalized square-zero monomial z_i. The parent support contains all pins in the monomial's scope, and $|z_i|_w \le 1$. The marked activity expansion produces a family $K(t, A)$ indexed by subsets of the finite leaf inventory. Its mixed jet in the leaves is evaluated on products of the z_i; because $z_i^2 = 0$, only distinct-leaf derivatives survive. Thus the marked polynomial is a finite substitution in these jets. The coefficient bounds on the mixed jets yield local polynomial bounds by the triangle inequality and $\|\prod_i z_i\|_w \le 1$.

The output has an exponential correction and a forced activity correction. The latter is grouped by the coarse support-and-inventory barcode. Its scalar mode remains the unmarked background profile. For a singleton marker v, only single-leaf jets can contribute; the single-leaf jet identity is

$$\partial_i K_j = \mathrm{fee}_i \, D_j(M_i),$$

where $D_j(M_i)$ is the transported hinge of that same leaf. Multiplication by its normalized monomial cancels the fee, including the zero-charge case established above. Hence the singleton coefficient of the grouped forced correction is exactly the sum of the leaves' quoted transported profiles. Absolute summability of the hinge families follows from the sound-door estimate: each hinge has a forced envelope whose total weighted mass is bounded by a fixed multiple of its leaf load. Summing over finitely many leaves and over the barcode fiber therefore commutes with coefficient extraction. This proves the singleton-profile identity used in the restoration step.

The coarse exact density identity is likewise an evaluation identity for the marked expansion: the Mark push of the priced ticket expression equals the exponential output times the coarse field polynomial. It is multiplied by the factored scalar displacement, which commutes with the Mark integral since it is a fixed finite-dimensional polynomial. The exponential output has zero scalar mode. The coefficient bounds and scoped locality give its served envelope; the forced mixed-jet bounds give the field envelope. These are the reserve's two output budgets. The same construction is additive under summable grouped leaves because the underlying hinge map respects those sums. The switched marker expansion, primitive integral and connected-cluster arguments in this chapter construct the marked activity jets and prove their exact Mark identity.

The exact marked identity follows from the real switching identity once its paid bounds have been established. At a coarse configuration U, let $a = e^{E_A(U)}$, let $L(x, R)$ be the sum of the regular log increments relative to the unmarked branch, and let $G(x, R)$ be the gas of the switched forced activities. The real identity is

$$a\, e^{L(x,R)} G(x, R) = \mathcal{P}\big(\mathrm{Buy}(x, R)\big)(U), \qquad 0 \le x \le 1,$$

with $\mathcal{P}$ the actual Mark push. The paid regular increments have a summable total budget. The forced-activity bounds give a summable paid expansion of the gas. Apply $\mathcal{S}_z$ to this identity, commute it with those summable series, and use its product and exponential laws. Since the regular log increment is zero at the unmarked switch, its substituted scalar mode is zero. The left side becomes

$$a\,\mathrm{Exp}\left(\sum_S H_S(U)\right)\mathrm{Field}(K(U), F_A(U)).$$

For the right side, expand Buy into finitely many tickets. A ticket with leaf set D has switch dependence $x^{|D|}\mathbf{1}_{D\subseteq R}$ times its scalar coefficient; its switching jet is concentrated at D. Its substitution is therefore that coefficient times $\prod_{i\in D} z_i$. Each scalar ticket coefficient is integrable against the Mark kernel. Finite linearity of integration proves that substitution commutes with the push of this ticket sum. Its value is exactly the push of the priced polynomial inventory. This proves the marked Mark identity from its real counterpart and paid envelopes, with all infinite summations justified by their paid series bounds.

The shipping bound ensures that this finite inventory is admissible to the activity estimates. If a fine leaf has cover cost $v = c_{B_1}(S)$, its enlarged coarse birth support obeys

$$(q+1)c_{B_*}(\mathrm{birth}(S)) \le (q+1)C_{\mathrm{foot}} + \tfrac{5}{4}\varepsilon v.$$

The tariff also spends $2\varepsilon v$. With $s_1 \ge 100$, the selected relative width and exponent satisfy $\frac{13}{4}\varepsilon \le s_1 - 4$. Thus the original fine support weight pays both costs, leaving only the fixed multiplier $e^{(q+1)C_{\mathrm{foot}}}$. Charge each leaf to one of its scoped pins, use the parent-pin occupancy bound M, and convert the rooted budgets to a sitewise bound. The remaining numerical demand is

$$M\,\mathrm{siteFactor}(B_*)e^{(q+1)C_{\mathrm{foot}}}\,v' \le \mathrm{ribbon},$$

where v' is the inventory's input budget. This is the shipping price. It pays the aggregate tariff after grouping equal activity labels, since those groups simply sum nonnegative leaf prices. Adding the unmarked activity reserve gives the admissible marked tariff. The geometric birth-support inequality and the scalar connected-activity estimates providing the ribbon, paid log increments and switched gas identity remain the lower-level inputs to this construction.

7.29 The numerical form of a retained insertion step

Write $s_1, s_0 \ge 100$ for fine and coarse support exponents, B_1, B_0 for their support widths, and M for the occupancy bound. Choose $0 < d \le 1$. First replace fine marker weights I_v by dI_v. Since every correction has degree at least two, its new monomial weight is at most d times its old weight, so its unit price budget becomes at most d. Reduce the support exponent to $s_1 - 4$. The finite transport inventory then has cost controlled by

$$T + e^T T, \qquad T = M\,\mathrm{siteFactor}(B_1)d + 4d.$$

Purchase a Mark-step reserve whose input budget is at least this quantity, and whose output field budget is u. The response transport bounds must satisfy $2I_v a_v \le 1$ and $I_v b_v \le 1$, where a_v, b_v are its two local envelope prices. After shrinking, these become $2dI_v a_v \le d$ and $dI_v b_v \le d$, precisely the bounds required by the reserve.

At the coarse level put $Q = u + M\,\mathrm{siteFactor}(B_0)d$. The numerical requirements are

$$(2B_0+1)^4 Q \le \tfrac{1}{2}, \qquad e^Q - 1 \le 1.$$

The first pays the closure loss in support exponent, and the second bounds its final correction price by one. The matching condition on coarse weights J_v pays another half unit of support exponent when passing from dI_v to J_v, with price multiplier at most one. The reserve supplies the coarse exponential exponent $s_0 + \frac{1}{2}$ and field exponent $s_0 + 1$; these two half-unit expenditures leave at least s_0. The singleton feed identity and the restoration argument above identify the coarse primary and forced insertions. This proves the retained unit-price served conclusion conditional on the finite transport inventory and Mark-step reserve. The finite transport inventory and marked reserve constructed in this chapter provide these inputs.

7.30 Separation pays for complex inflation

The large analytic disk in the omission argument is a geometric consequence of separated markers. Let $B \geq 8$, let $A \subseteq V$ have at least two elements, and suppose distinct pins in A have torus max-distance at least B^2. If a support S contains these pins, then

$$B|A| \leq 64 c_B(S).$$

To prove this, choose a minimizing cover with $D = c_B(S)$ centers. Its halo admits a walk visiting the centers in at most (D(4B+2)) unit steps. For each pin choose a visit time whose center is within B of it. Two chosen times differ by at least $g = B^2 - 2B > 0$, by the distance triangle inequality and the unit-step walk bound. Thus their integer quotients by g are distinct and $(|A| - 1)g \leq N$, where N is the walk length. Since $|A| \geq 2$ and $B \geq 8$,

$$|A|B^2 \leq 4(|A| - 1)g \leq 4N \leq 4D(4B + 2) \leq 64DB.$$

Division by $B > 0$ gives the assertion. Notice that the walk visits covering centers; visiting the pins themselves is unnecessary.

Every nonzero higher correction has a nonzero mode with at least two markers, all pinned in its support. Hence $c_B(S) \geq B/64$, and

$$S_B := e^{B/64} \leq e^{c_B(S)}.$$

For $|z| \leq S_B$, multiply a nonzero correction envelope a_i by (|z|), and set the new envelope to zero for identically zero corrections. Then

$$|z| a_i e^{(e-2)c_B(S_i)} \leq a_i e^{(e-1)c_B(S_i)} \leq a_i e^{e c_B(S_i)}.$$

Thus multiplying every higher correction by z preserves the price budget while reducing the support exponent by two. Locality, gauge invariance, measurability, scope and permission conditions survive scalar multiplication, and no scalar or singleton mode is introduced. This proves precisely the inflation estimate needed for holomorphy and retained transport in the omission comparison. It leaves the primary and forced singleton insertions at their literal values.

7.31 A positive denominator and an explicit vacuum bill

Let μ_G be normalized Haar measure and d_M the real dimension of the ambient matrix space of the chosen representation. Normalize the matrix chord by the constant matrix norm $\gamma > 0$, so $\mathrm{chord}(g) = |g - I|/\gamma$. For $0 < r \leq 1$,

$$\mu_G\{g : \mathrm{chord}(g) \leq r\} \geq (r/4)^{d_M}.$$

Here is an ambient-space proof, requiring no coordinates on the group. A collection of group matrices with mutual distances greater than $r\gamma$ gives disjoint ambient balls of radius $r\gamma/2$, all inside the ball of radius 2γ. Comparing Euclidean volumes bounds its cardinality q by $q(r/4)^{d_M} \leq 1$. A finite maximal such collection exists: otherwise successive insertion would violate this cardinality bound. Its radius-$r\gamma$ closed balls cover the group. By unitary invariance of the matrix norm, their intersections with the group are left translates of the chord shell. Haar invariance and the union bound give $1 \leq q\mu_G(\mathrm{shell}(r))$. Combining the inequalities proves the assertion.

On a finite four-dimensional torus Λ, take the shelter of configurations with every positive oriented link in the chord shell of radius $s/4$, where $0 < s \leq 1$. Product Haar measure gives this shelter measure at least

$$(s/16)^{d_M|E|}, \qquad |E| = 4|\Lambda|.$$

The chord of a product is at most the sum of the chords, and inversion preserves the chord. Every plaquette in the shelter consequently has chord at most s. If $s \le \theta$, the field threshold, there are no flags there; its activity gas is exactly one. Signed translations preserve the shelter as they only permute links and invert some of them.

Write C for the path bound, Λ_0 for the loop cap, δ for the regular-activity allowance, κ for the forced allowance, $\beta \ge 0$ for the coupling, and C_{dir} for the direct quadratic clipping coefficient. Define

$$\ell = \beta(C\theta^3 + C\theta^4) + \Lambda_0\theta^2, \qquad b(s) = |\mathcal{P}| \left(\frac{C_{\mathrm{dir}} s^2}{2} + \frac{C^2 s^2}{2} \right), \qquad e_0 = \delta + \ell,$$

where $\mathcal{P}$ is the finite set of plaquette-axis pairs. Each healthy atom of degree k is bounded by its coefficient magnitude times θ^k: every transported curvature leg has norm at most the threshold, since its footprint is healthy and parallel transport preserves the norm. Absolute summation with the atom budgets bounds the cubic, quartic and loop remainders per site by $C\theta^3, C\theta^4, \Lambda_0\theta^2$, respectively. The quadratic base is nonnegative. Thus the channel-averaged action is bounded below by $-|\Lambda|\ell$.

On the shelter, every curvature leg has norm at most s. Summing the absolute path coefficients gives norm at most (Cs) for the linear path field, so the quadratic base per site is bounded by (b(s)). The same remainder estimate bounds the channel action above by $|\Lambda|(\ell + \beta b(s))$. The regular activity exponent has absolute value at most $|\Lambda|\delta$: charge each occurrence's weighted budget to its designated root site and sum those root budgets. If v_0 is the scalar vacuum exponent, the full exponent E_R therefore satisfies

$$E_R(U) \le v_0 + |\Lambda| e_0 \quad \text{everywhere}, \qquad E_R(U) \ge v_0 - |\Lambda|(e_0 + \beta b(s)) \quad \text{on the shelter}.$$

The represented density is nonnegative. At the finest scale this follows from the positive bare Wilson density; downward induction preserves it because each Mark kernel is positive. Integrate over the shelter, where the gas is one, and use nonnegativity outside it. The resulting denominator bound is

$$Z_R := \int R(U)\, d\mu(U) \ge \exp\{v_0 - |\Lambda|(e_0 + \beta b(s)) - 4|\Lambda| d_M \log(16/s)\} > 0.$$

Combining this with the insertion norm bound proves the explicit normalized estimate

$$t_w(A) \frac{|\int [X(U)]_A \, d\mu(U)|}{Z_R} \le \exp\left\{ |\Lambda| [2e_0 + \kappa + \beta b(s) + 4 d_M \log(16/s)] + |V|(x + y + 2f) \right\}.$$

The scalar v_0 cancels. This bound applies to the full denominator, with no change of law or restriction of the normalized numerator to the shelter.

For a selected history use $s = \widehat{g}_k = \widehat{\beta}_k^{-1/2}$, with $\beta \le 2\widehat{\beta}_k$. The loop price is monotone in nonnegative β, as is $\beta b(s)$. Hence the coefficient of $|\Lambda_k|$ in this bound is at most the explicit treasury

$$\tau_k = 2\{\delta_k + \ell_k(2\widehat{\beta}_k)\} + \kappa_k + 2\widehat{\beta}_k b_k(\widehat{g}_k) + 4 d_M \log(16/\widehat{g}_k).$$

This establishes the root estimate used in the omission argument. Its extensivity is explicit; no volume factor is silently absorbed in an error tending to zero.

The treasury also has a useful stagewise bound. Set $P(s) = \exp\{(\log\log s)^2\}$ and $\Theta(s) = s^{-1/2}P(s)$. For every fixed b and every $p > 0$, eventually $P(s)^b \le s^p$, because $b(\log\log s)^2/\log s \to 0$. Consequently $s\Theta(s)^3 \le 1$, $s\Theta(s)^4 \le 1$, and $2P(s)^v\Theta(s)^2 \le 1$ uniformly for $v \le 1$, eventually. Substitution into the actual thresholds and loop caps gives $\ell_k(2s) \le 4C_k + 1$, while $\delta_k, \kappa_k \le 1$. Also $s\widehat{g}_k^2 = 1$. Therefore, eventually,

$$\tau_k \le 2\{1 + (4C_k + 1)\} + 1 + |\mathcal{P}|(C_{\mathrm{dir},k} + C_k^2) + 4d_M\{\log 16 + \tfrac{1}{2}\log\widehat{\beta}_k\}.$$

Stage parameters are kept visible in this inequality; it asserts logarithmic growth in the reference coupling only after those parameters are fixed.

8 Convergence of the interacting scale histories

Convergence requires one common numerical schedule for all fixed finite demands. The order is: choose the finite stage profile, then its block length, then its post-length constants and coupling-entry delay, and only afterwards the terminal cutoff and physical period. Increasing a delay does not reselect a previously fixed block. This prevents circular parameter choices in the stationary estimates, shooting construction and actual coupling running.

8.1 A common analytic stationary graph and its response estimates

The stationary graph solves the vertical critical-point equation in a common analytic neighborhood. Positivity of its vertical Hessian gives the inverse used in the response estimates. Retained coordinates and vertical coordinates have different norms; the explicit color-coordinate maps and their bounds below justify each passage between them. No extensive scalar action is assigned a volume-independent zeroth-order norm.

We specify first the actual score to which the analytic argument is applied, at a fixed derivative order $n \ge 2$. A retained coordinate is a coarse oriented bond together with a real color index. A vertical coordinate is a free fine bond of the fixed forest together with such an index. Their disjoint union is the joint coordinate set. Arrays carry the maximum norm. Splitting the union is an isometry; replacing real color coordinates by their Euclidean color vector defines a linear isomorphism paint. Its norm is at most $1 + \dim\mathfrak{g}$, and its inverse has norm at most one. Moreover the sum of squared vertical real coordinates is exactly the sum of squared Euclidean color norms. Denote the resulting retained and vertical vectors by $(\mathrm{tint}\, x, \mathrm{tint}\, v)$. The actual scalar score in these coordinates is

$$F(x, v) = \mathrm{tested}(\mathrm{tint}\, x, \mathrm{tint}\, v). \tag{AS1}$$

Thus it is the main fine action plus the prediction penalty, with the fine and outer configurations constructed from these same variables. The connector is not an independent variable. A retained coordinate is addressed at its embedded coarse site, and a vertical coordinate at its fine site. Each address fiber has at most $h = |\mathrm{TypeFace}_1 \times \mathrm{ColorIndex}| + |\mathrm{Axis} \times \mathrm{ColorIndex}|$ members: orientation and color distinguish the members, and the coarse embedding is injective.

Here is the local-to-global estimate, including the counting loss. For addressed coordinate slots $q_1, \dots, q_k$, put $S(y, q) = \sum_i \mathrm{dist}(y, \mathrm{address}(q_i))$. An anchor estimate of price M and width a means

$$|D^k f[e_{q_1}, \dots, e_{q_k}]| \le Me^{-aS(y,q)} \quad (k \le n). \tag{AS2}$$

If a term depends only on coordinates within distance R of y, and its operator derivatives through order n are bounded by M_0, a coordinate derivative vanishes whenever one of its slots is outside that set. Otherwise $S \le nR$. Thus (AS2) holds with $M = M_0 e^{anR}$. For path series the common exponentially weighted length budget dominates these derivative sums absolutely; differentiation term by term follows on a smaller common analytic neighborhood. This proves the same estimate for the cubic and quartic series, not merely for finite path sums.

Sum such terms over anchors with multiplicity at most h_0. Fix a marked slot q_j. The triangle inequality gives $\sum_i \operatorname{dist}(\operatorname{address}(q_j), \operatorname{address}(q_i)) \le (n+1)S(y,q)$ and $\operatorname{dist}(y, \operatorname{address}(q_j)) \le S(y,q)$. Split the exponential in half and use the lattice counting estimate

$$c(b) = 16\,6! e^{2b} b^{-6}.$$

The summed derivative has price $Mh_0 c(a/2)$ and width $a/[2(n+1)]$, measured from any marked slot. This is the Scatter bound used here. At order zero there is no marked slot; no volume-independent bound for the total scalar action is being asserted.

For the actual main score take

$$a = \frac{\mu}{4(n+1)(1+L)}.$$

At a signed site its density is the sum over planes of $p_{\mathrm{direct}} \operatorname{price}(U_\square) + \|\ell(U)_\square\|^2/2$, together with the cubic and quartic path series. The plaquette has length four and locality radius $5 + L$. The matrix product rule costs 2^n per product; inversion costs one for unitary matrices. In particular, writing $\mathcal{C}_n(K, C) = n! K (1 + C)^n$, a walk of length at most r has derivative price

$$\mathcal{T}_n(r, C) = \mathcal{C}_n(e\|1\|[1 + 2^n(1 + r)]^n, C).$$

The Wilson price has bound $\mathcal{P}(C) = 1 + \|\mathrm{traceCharge}\| C$, where $\mathrm{traceCharge}(M) = -\operatorname{Re}\operatorname{tr}(M)/m$. If o is the purchased outer-link jet bound, the fine-link bound is $C_f = H + 2^n H^2$, $H = o + \|1\|$, by the same product rule. These are ordinary finite-dimensional derivative bounds on the unit real coordinate ball.

To make the path prices explicit, let r_μ be the common path analyticity radius at width $\mu/2$, and set $J_n = \mathcal{C}_n([1 + 2^n/r_\mu]^n, C_f)$. If s_G is the color-slot coefficient and c_G the real color projection, the row and atom prices are $R_n = J_n |s_G| \|c_G\| e^\mu$ and $A_{n,k} = J_n |s_G|^k e^\mu$. These are the convergent weighted path bounds obtained by spending half the exponential length width. Choose $Z' > 0$ bounding all derivatives through order n of $w \mapsto \|w\|^2/2$ on $\|w\| \le CR_n$; any number exceeding $1 + CR_n + (CR_n)^2/2$ suffices. Put $N_n = \sum_{i=0}^n |\mathrm{OrderedFinpartition}(i)|$. The per-plane and per-site prices are

$$Q_n = |p_{\mathrm{direct}}| \mathcal{P}(\mathcal{T}_n(4, C_f)) e^{an(5+L)} + N_n Z' (1 + CR_n)^n, \qquad M_n = 6Q_n + CA_{n,3} + CA_{n,4}. \quad \text{(AS3)}$$

The second term is the ordered-partition chain rule for the squared row. Signed site maps are injective, so their anchor multiplicity is one. Averaging over the signed symmetries cancels their number. The main-action Scatter price is therefore $M_n c(a/2)$.

For the prediction term choose derivative inputs through order $n+1$ first. If C_f^+ is the fine-link price at that order, put

$$F_* = 2\,\mathrm{alignCost}\,\mathrm{pathChange} + L\,\mathrm{variation}, \quad \rho = \min\left\{\rho_{\mathrm{frame}}, \frac{\sqrt{\mathrm{predictorCut}}}{F_*}, \frac{r}{\mathrm{colorScale}\,F_*}\right\}. \quad \text{(AS4)}$$

Here ρ_{frame} is the positive local minimizer radius for the frame score with the purchased $(n+1)$-jets and half its root gap. The constants pathChange and variation bound the actual connector paths under a unit coordinate displacement; alignCost converts the two endpoint frames. The three inequalities in (AS4) keep the frame minimizer in its implicit chart, keep the focus below its predictor cutoff, and keep its group logarithm inside radius r, respectively. Consequently the selected connector agrees there with the analytic minimizer. One differentiates that analytic expression, rather than the discrete selection tags.

Here are the local frame constants in this formula. Let N_P be the number of frame points, with the central frame fixed to the identity. The blank score's quadratic form is the average over ordered point pairs of squared differences of their logarithmic frame coordinates. The pairs with one endpoint at the fixed center contribute twice the sum of squared free coordinates. Consequently its counting gap is $g_P = 2/N_P^2$. Put $e_n = \mathcal{C}_n(K, 1)$, and, for input link jets bounded by C, set

$$S_n = \mathcal{P}\left(2^n[2^n e_n \mathcal{T}_n(8R, \mathcal{C}_n(C, 1))]e_n\right), \quad R_H = \min\left(1, \frac{g_P}{2(1+S_n)}\right). \tag{AS4a}$$

There are two frame factors and an intervening route of length at most $8R$; their product gives exactly this jet bound for the frame score. The mean-value estimate for its Hessian on the radius-R_H ball changes the blank Hessian by at most $S_n R_H \le g_P/2$. Its inverse there has counting-norm bound $2/g_P$. Let $J = 1 + 8RC/\mathrm{colorScale}$, the path chord price, and let $A = \mathrm{alignCost}\, J$. For the purchased logarithm derivative bound D, valid inside chord radius δ, the frame radius is

$$\rho_{\mathrm{frame}} = \min\left\{R_H, \frac{\mathrm{frameCut}}{3J}, \frac{\delta}{A}, \frac{R_H}{N_P D A}\right\}. \tag{AS4b}$$

The last term puts the logarithmic coordinate of the selected minimizing frame inside the Hilbert ball; the preceding two ensure its branch and logarithm agree with the chosen analytic chart. These are the finite local frame and alignment constructions in the nonlinear geometry section. Applying them with order $n+1$ is why the penalty estimate purchases that extra order before choosing (AS4).

The implicit derivative polynomial used here is also explicit. For branch value bound R_0, equation jet bound M, and inverse linearized-equation bound C_0, define

$$I_0 = R_0, \qquad I_{a+1} = I_a + C_0 M\{1 + |\mathrm{OrderedFinpartition}(a+1)|(1+I_a)^{a+1}\}. \tag{AS4c}$$

Differentiate the implicit equation $a+1$ times. Its single-block term is the vertical linearized equation applied to the unknown highest branch derivative. Every other ordered-partition term uses only branch derivatives of orders at most a, and is bounded by $M(1+I_a)^{a+1}$. Multiplication by the inverse bound proves the recursion; the added I_a retains all lower jets. In the frame application $R_0 = 1$, M is the frame-score price and $C_0 = 2/g_P$. This defines the polynomial used in the next paragraph without an unpriced implicit differentiation step.

More explicitly, let I_n be the implicit-derivative polynomial for the local frame score, evaluated with its score-jet price and inverse half-root-gap. The frame transport price is $H_n = \mathcal{C}_n(K, I_n)$. Then the connector price is $K_{\mathrm{conn}} = 2^n[2^n H_n \mathcal{T}_n(L, C_f^+)]H_n$, by its two endpoint factors and intervening path. Here K bounds the local exponential chart jets, and W below bounds the focus-map jets on the radius kept by (AS4). Also $o_{n+1} = \mathcal{C}_{n+1}(Z, 1 + \dim \mathfrak{g})$, where Z is the purchased group-chart jet price. Each prediction bond is supported within radius $3L + 16R$, where R is the patch radius. Its anchor price is explicitly

$$N_{\mathrm{pred}} = \mathcal{P}\left(2^n o_{n+1} \mathcal{C}_n(W, K_{\mathrm{conn}})\right) e^{an(3L+16R)}. \tag{AS5}$$

There are four coarse bond orientations per embedded anchor, so summation has multiplicity four. Thus the penalty Scatter price is $4|p_{\mathrm{penalty}}|N_{\mathrm{pred}}c(a/2)$. Combining the two terms gives the actual price

$$C_* = 1 + |M_n c(a/2) + 4|p_{\mathrm{penalty}}|N_{\mathrm{pred}}c(a/2)|\,, \quad a_* = a/[2(n+1)]. \tag{AS6}$$

All constants in (AS3)–(AS6) are fixed from the group, stage, path budgets, local frame jets and patch geometry before the torus period. In particular no number of lattice sites occurs in them.

Finally the origin slope and gap are consequences of the actual score expansion. The quadratic line limit is one half of the signed average of $\sum_{\square}[p_{\mathrm{direct}}\|dz\|^2 + \|\ell_{\mathrm{lin}}z\|^2]$ plus one half of the penalty times the squared linear prediction residuals. Cubic and quartic terms start in their respective orders. Analytic Taylor expansion therefore gives zero first derivative at zero and this quadratic form as the Hessian. Restrict to vertical vectors. Keeping the identity signed orientation gives coefficients $b = p_{\mathrm{direct}}/|\mathrm{Signed}|$ and p_{penalty} in front of the curl and prediction-residual energies. The forest counting estimate reconstructs a vertical field from these energies, with loss $K = \mathrm{curlFactor}(L, \mathrm{box}, \mathrm{linkCost}) + \mathrm{probeFactor}(L)$. Its tree edges vanish; the harness link estimate bounds the remaining reconstructed links. Hence

$$D_v F(0,0) = 0, \qquad D_v^2 F(0,0)[v,v] \geq \frac{\min(b, p_{\mathrm{penalty}})}{1+K} \sum_i v_i^2. \tag{AS7}$$

For an explicit count, let X be the finite cover box of radius macroR, let $V_L = L^4$, and let $c_1, c_2 > 0$ be the two local harness constants. With kernel predictor cutoff $q_* > 0$, put

$$\alpha = 2\max(1, q_*^{-1})(2c_1 + c_2), \qquad E_L = (8L+1)c_2 L + 2c_1 L^2 + \alpha/L^2.$$

The local link reconstruction bounds a nonnegative link norm by E_L times the sum of curvature norms in its translated X-box, plus its coarse prediction-residual norm. Squaring uses $(a+b)^2 \leq 2a^2 + 2b^2$, and the curvature sum uses Cauchy–Schwarz with $6|X|$ terms. Summing fine bonds, reindexing their coarse parent and residue, introduces $4V_L$ positive-axis residue bonds; a fixed translated curvature can occur at most $|X|$ times. A residual occurs V_L times. Thus

$$K_{mcurl} = 2E_L^2 V_L\, 4\,(6|X|)\,|X|, \qquad K_{mprobe} = 2V_L, \qquad K = K_{mcurl} + K_{mprobe}. \tag{AS8}$$

This derives the finite denominator in (AS7); it may grow with the blocking ratio and fixed patch, but never with the macroscopic torus period. The positive numerator is part of the actual step hypotheses. The equality of color and coordinate squared sums above gives precisely this counting norm, with no dimension conversion. Also $\mathrm{tint}(e_i)$ is the corresponding physical color frame, so the vertical Hessian block is the physical one. Invertibility of tint shows that the zero vertical derivative in (AS1) is exactly the original stationary equation. Thus (AS4), (AS6), the address multiplicity h, and the gap in (AS7) provide the actual initial radius, width, price, multiplicity and gap for (US1)–(US8). Their stationary graph, inverse-Hessian and determinant estimates consequently apply on the same common domain used in stationary integration and the reference chart.

We make explicit the uniform implicit construction used in the stationary integrations. Split the finite coordinate set into retained coordinates ι and vertical coordinates τ, assign every coordinate a lattice site by l, and assume at most K coordinates have any given site. Arrays carry their maximum norm. Write $f(x,v)$ for the real analytic cost, $S(x,v) = D_v f(x,v)$ for its vertical gradient, and $H(x,v) = D_v^2 f(x,v)$. The analytic scatter bounds established above say that its coordinate derivatives, through a purchased order $N \geq 3$, are bounded by a constant $C > 0$ times the exponential connecting weights with width $a > 0$, on the maximum-norm ball of radius $r > 0$.

Assume $S(0,0)=0$ and the counting-measure gap

$$\sum_{ij} u_i H(0,0)_{ij} u_j \geq m \sum_i u_i^2, \qquad m>0. \tag{US1}$$

All these inputs were obtained for the actual cost; a bound on the extensive value $f(0,0)$ is unnecessary.

Take the explicit four-dimensional counting bound $c(a)=16\cdot 6! e^{2a} a^{-6}$. A distance shell of radius k has at most $(2k+1)^4$ sites. The exponential-series inequality $e^{a(k+2)} \geq [a(k+2)]^6/6!$ bounds its contribution by $c(a)/[(k+1)(k+2)]$. Summing the telescoping series gives $\sum_y e^{-a\,\mathrm{dist}(z,y)} \leq c(a)$, uniformly in the period. Bounded multiplicity changes this to $Kc(a)$. Put $D=Kc(a)$. The scatter third-derivative bound and integration along a segment give

$$\|H(z)-H(0)\|_{\infty\to\infty} \leq \|z\|_\infty C D^2. \tag{US2}$$

The same row and column sums bound the corresponding quadratic form: apply $2|u_i u_j| \leq u_i^2+u_j^2$ to each entry. Thus the gap remains at least $m/2$ on the ball of radius

$$R_H = \min\{r/2, m/[2(1+CD^2)]\}. \tag{US3}$$

We need an inverse in the maximum norm with a price independent of the number of sites. Suppose a matrix B has entries bounded by $Ce^{-a\,\mathrm{dist}(l_i,l_j)}$ and counting gap $M>0$. For a fixed site z, conjugate it by the diagonal weights $w_i = e^{b\,\mathrm{dist}(z,l_i)}$. The triangle inequality and $|e^t-1| \leq |t|e^{|t|}$ show that, if $0<b\leq a/4$, the difference from B has entries bounded by

$$Cb(4/a)e^{-(a/2)\,\mathrm{dist}(l_i,l_j)}.$$

Choose

$$b = \min\left\{a/4, \frac{M}{2\{1+C(4/a)Kc(a/2)\}}\right\}. \tag{US4}$$

The quadratic-form price of the perturbation is at most $M/2$; hence the conjugated matrix has counting coercivity $M/2$. It need not be symmetric. To estimate its inverse column, solve $Bu=e_j$, take $z=l_j$, and put $v_i=w_i u_i$. Coercivity gives

$$(M/2)\sum_i v_i^2 \leq v_j.$$

In particular $(M/2)v_j^2 \leq v_j$, so $v_j \leq 2/M$, and then $|v_i| \leq 2/M$ for every i. Undoing the weight gives

$$|(B^{-1})_{ij}| \leq (2/M)e^{-b\,\mathrm{dist}(l_i,l_j)}. \tag{US5}$$

Summing a row yields a maximum-norm inverse price $(2/M)Kc(b)$. This argument uses the true counting gap and exponential summability throughout, and introduces no volume-dependent equivalence between the maximum and Euclidean norms.

Apply this with $M=m/2$ on (US3), and denote the resulting width by b. Set

$$J = 1+(4/m)Kc(b), \quad R_V = \min\{R_H, [4J(1+CD^2)]^{-1}\},$$

$$R_X = \min\{R_V, R_V/[4J(1+CD)]\}. \tag{US6}$$

For $\|x\| < R_X$, consider the map $Y_x(v) = v - H(0,0)^{-1}S(x,v)$ on $\|v\| \leq R_V$. Its derivative is $H(0,0)^{-1}\{H(0,0)-H(x,v)\}$. Equations (US2) and (US6) give its norm at most $1/4$. Integrating

the mixed Hessian in the retained variable gives $\|S(x,0)\| \le CD\|x\|$, so $\|Y_x(0)\| \le R_V/4$. Therefore $\|Y_x(v)\| \le R_V/2$ on that closed ball. Iteration is Cauchy with successive differences reduced by $1/4$; its limit $g(x)$ is its unique fixed point. Since $H(0,0)$ is invertible, being a fixed point is exactly $S(x,g(x)) = 0$. We have

$$g(0) = 0, \quad \|g(x)\| \le R_V/2, \quad S(x,v) = 0,\ \|v\| \le R_V \implies v = g(x). \tag{US7}$$

At every point of this graph the vertical Hessian is invertible. The local analytic implicit function theorem gives an analytic branch there; its continuity keeps it inside the same vertical ball nearby, and uniqueness in (US7) identifies it with g. These local identifications prove analyticity of the single graph on the whole retained ball. This also identifies any retained critical germ with this graph wherever both satisfy the corridor bounds.

For completeness, its higher response estimates follow directly from the critical equation, rather than from a dimension-dependent implicit-function norm. Let $Z(x) = (x, g(x))$. Differentiate $S \circ Z = 0$ $k+1$ times. In the ordered-partition chain formula the one-block term containing $D^{k+1}g$ is $H(Z)D^{k+1}g$; every other term contains graph derivatives of order at most k. At a fixed purchased order use the common width

$$b_* = \min\{\mathrm{weaker}_N(a), \mathrm{weaker}_N(b)\} > 0,$$

where $\mathrm{weaker}_N(t) = t/[2(N+1)]$. Combining the joining-tree weights loses only this width once for the whole induction: the span of any one resulting tree is at most the sum of its constituent spans, and splitting each available exponent between this connecting span and the summation variables leaves the remaining half for the counting sums. At most $N+1$ vertices are joined at the purchased order. If P_k prices graph derivatives through order k, the other terms have price at most

$$F_k = C + |\,\mathrm{OrderedPart}(k+1)|[C\{1 + P_k K c(a/2)\}]^{k+1}.$$

Multiply by the inverse kernel (US5) and sum its output coordinate. The recursion

$$P_0 = 1, \qquad P_{k+1} = P_k + F_k(4/m)Kc(b/2) \tag{US8}$$

therefore bounds every coordinate response through order $k+1$, with its connecting-tree exponential of width b_*. Every constant in (US3)–(US8) depends on the purchased finite orders and local input prices, but not the period. Differentiating $f \circ Z$ uses the same ordered partitions; its terms containing a vertical first derivative of f vanish by stationarity. This gives the on-shell action bounds used in the nonlinear integral and the reference estimates. The separate stronger strip estimates for generated precision kernels retain their own widths; (US8) does not replace them.

8.2 The Fourier normalization and the kinetic readout

The two-field determinant film has an internal momentum and two opposite external momenta. Its internal average is normalized Haar average, namely $\mathrm{stock}\, F = \int_{[-\pi,\pi]^4} F(k)d^4k/(2\pi)^4$; its finite-period average is $\mathrm{shop}_n\, F = n^{-4}\sum_k F(2\pi k/n)$. The curvature picture called $\mathrm{profitable}\, F$ is one half of the internal average with the two external link and color indices in their specified order. The factor one half is the quadratic Taylor factorial. Translation invariance makes the actual Hessian entry zero unless its two Fourier momenta sum to zero, and otherwise that entry is the site count times the finite-period average of the film.

To compare the two averages, expand the internal variable in its absolutely convergent Laurent series. Finite averaging retains exactly the coefficients whose four indices are divisible by the

period; Haar averaging retains only the zero coefficient. Every remaining index has length at least the period. Spending half the common strip for summation and a quarter for the tail gives a fixed strip price times $e^{-\sigma n/4}$. Multiplication by the site count gives the recorded-Hessian error. Thus this internal integral represents the finite Hessian with a quantitative error and a fixed normalization.

Compose that picture with the momentum chart for real initial links. The readout $\mathrm{leapTip}(\mathcal{P})(e_0, 1, (a, a))$ is the divided coefficient $D^2 P_{(1,a),(1,a)}(0)[e_0, e_0]/2!$. Reality makes this diagonal coefficient real. Color invariance and simplicity make its color bilinear form a scalar multiple of the fixed invariant form, so its value is independent of the normalized color. The factor two in (DP1) is additional to the factor one half in the internal film average; both are retained in that definition.

Here is why this one transverse entry determines the leading kinetic tensor. Subtract its multiple of the normalized Wilson tensor, whose readout is one. Gauge invariance gives transversality. Signed covariance includes the phases from reversing initial links; at the first nonzero homogeneous order those phases have constant term one. The constant and linear homogeneous terms vanish by transversality and permutation covariance. Let $Q_{uv}(p, q)$ be the symmetric polarized degree-two term of the difference. Reflecting one axis shows that a component can survive only when its momentum occurrences and link occurrences have the same parity on every axis. Permuting axes moves the zero transverse diagonal readout to all transverse diagonal entries on single momentum axes. Longitudinal entries vanish by transversality. Reflection parity then kills mixed momentum entries in a diagonal link component and same momentum entries in an off-diagonal link component. The only remaining possibility has its two distinct momentum axes equal to its two link axes. Transversality at the sum of those two coordinate vectors makes twice that entry zero, since the diagonal entries have already vanished. Bilinearity gives $Q = 0$. The degree-three term vanishes under total inversion, whose momentum sign is odd and whose two link signs multiply to one. The difference thus starts at degree four. This proves the one-axis kinetic identification used below, including the Taylor normalization and reflected-edge convention.

8.3 Analytic exterior division with the origin jets retained

Division retains the harmonic obstruction Exterior division first solves a contraction equation on the regular momentum set. Analytic extension requires the stated vanishing along the relevant harmonic coordinate planes, not merely a value zero at the single full origin. The division formula below keeps the origin jets and the constraints from the other momentum blocks. These data are needed by the two-leg correction that follows; deleting a harmonic component instead would change the supplied stationary identity.

We give the analytic division used in issuing the loop coefficients. For one momentum coordinate set $q_i(p) = e^{ip_i} - 1$. If x is a Laurent index, define

$$L_{i,x}(p) = \begin{cases} \sum_{a=0}^{x_i - 1} e^{ip\cdot(x[i:=a])}, & x_i \geq 0, \\ -\sum_{a=0}^{-x_i - 1} e^{ip\cdot(x[i:=-a-1])}, & x_i < 0. \end{cases}$$

The empty sum is zero. Multiplication and telescoping give the entire identity

$$q_i L_{i,x} = e^{ip\cdot x} - e^{ip[i:=0]\cdot x}. \tag{ED1}$$

If F is periodic analytic and bounded by A on a strip of width σ, its Fourier coefficient bound on an intermediate strip makes $G(p) = \sum_x \widehat{F}(x) L_{i,x}(p)$ absolutely and locally uniformly convergent on width $\sigma/4$. Indeed the finite sum has at most $|x|$ terms and size at most $|x| e^{(\sigma/4)|x|}$; the spare

width absorbs $|x|$. Summing the geometric coefficient bounds in each momentum coordinate gives the explicit price

$$\|G\|_{\sigma/4} \le A\frac{4}{\sigma}\left[2\frac{1+\sigma/4}{\sigma/4}\right]^{v}, \tag{ED2}$$

where v is the number of momentum coordinates. Thus $q_iG = F - F|_{p_i=0}$, including at every zero of q_i; no singular quotient is defined on a mesh. If the numerator has zero jets below order $s+1$, comparison of its analytic Taylor series with $q_i = ip_i + O(p_i^2)$ gives zero jets of G below order s. Equivalently evaluate on nearby directions with nonzero i-component, divide their one-variable series, and extend the resulting polynomial-jet identity to all directions by continuity.

For exterior coefficients let C be contraction by the vector q, and I_i exterior multiplication by the i-th coordinate vector with its orientation sign. The coordinate anticommutation rule gives $CI_i + I_iC = q_i\,\mathrm{id}$. Apply (ED1) componentwise to a cycle u, and call the divided difference D_iu. Substituting in this anticommutation rule shows that

$$CI_iD_iu = u - \mathrm{comb}_i\, u,$$

where $\mathrm{comb}_i\, u$ evaluates at $p_i = 0$ and deletes every exterior component containing i. The identity is first algebraic where $q_i \neq 0$, and extends to the zeros by analyticity. The residual is again a cycle. Indeed, after separating components containing i, its contraction at momentum p equals its contraction with the i-th symbol set to zero, since every i-containing component has been removed; the original cycle equation at $p_i = 0$ supplies the remaining equations. Evaluation and deletion preserve the original strip, coefficient bound and Taylor vanishing order.

Repeat on the residuals in the four axis directions, adding the four lifted differences. Every residual still has the original width σ, price A, and order s. Each new lift is therefore computed directly at width $\sigma/4$, price at most A times (ED2), and order at least $s-1$. There is no iteration of four strip shrinkages or four losses of Taylor order. Summing gives width $\sigma/4$, price four times (ED2), and order $s-1$. The contraction of this sum telescopes to the original cycle minus the final residual. In positive exterior degree the latter is zero because every orientation contains a processed axis. In degree zero it is the original coefficient with all coordinates of this block set to zero; the required hypothesis is its vanishing on that entire coordinate plane. Positive order at the full origin implies this hypothesis only when this is the sole momentum block. This is why the exact harmonic-set check in (TEL2) is needed.

Suppose additionally that each coefficient vector lies in a prescribed linear subspace depending only on the other momentum blocks. Evaluation in a divided coordinate preserves that subspace. The divided Fourier sum also does so: where $q_i \neq 0$ it is the scalar multiple $(u - u|_{p_i=0})/q_i$, and at its zeros continuity and closedness of the finite-dimensional subspace give the same conclusion. Exterior multiplication is a signed selection of these coefficient vectors. Hence every residual and lift retains the other block constraints. For a dependent momentum equal to minus the sum of the independent ones, the linear shear exchanging that block with the dependent block is involutive. If there are $t+1$ independent blocks, its missing coordinate is the negative sum of their coordinates. A strip of width $\sigma/(t+2)$ is mapped inside width σ, since $(t+1)\sigma/(t+2) < \sigma$. Apply the preceding coordinate construction on that narrowed strip and shear back with the same narrowing factor. This gives the dependent-block width and price used in (TEL3).

8.4 The two-leg correction with both momentum blocks retained

Here is the precise two-leg use of the exterior division. There is one independent four-momentum p; the other is $-p$. Write $E^r = \Lambda^r\mathbb{C}^4$, with its ordered plane basis, and let $C_-^r : E^{r+1} \to E^r$ be

contraction by the lattice symbol at $-p$. Thus $C_-^r C_-^{r+1} = 0$. A coefficient with one remaining potential index belongs to $E^r \otimes E^1$; one with no such index belongs to E^r. These are coefficient-valued periodic analytic functions of the same complexified momentum. Let $H : E^r \otimes E^1 \to E^r$ contract the second index by the symbol at p. Since these maps act on different tensor slots,

$$HC_-^r = C_-^r H. \tag{TEL1}$$

There is no exterior sign in this convention. A coefficient has order a if all its real Taylor coefficients of total degree below a at zero vanish. Each contraction increases this order by one.

The given two-potential tensor is $u \in E^1 \otimes E^1$, of order $s+2$, satisfying both gauge contractions $C_-^0 u = 0$ and $Hu = 0$. Apply the dependent-block division to obtain

$$X \in E^2 \otimes E^1, \quad C_-^1 X = u, \quad \operatorname{ord} X \geq s+1.$$

The coefficient $HX \in E^2$ has order at least $s+2$, and is a dependent cycle because $C_-^1 HX = Hu = 0$. A second dependent-block division gives

$$Y \in E^3, \quad C_-^2 Y = HX, \quad \operatorname{ord} Y \geq s+1.$$

Next lift the independent contraction: choose $D \in E^3 \otimes E^1$ with $HD = Y$ and order at least s. The degree-zero obstruction for this last division must be checked on its actual harmonic set. There is just one independent momentum block, so that set is exactly $p = 0$, rather than an unspecified collision locus. Since $s + 1 \geq 1$, every coefficient of Y vanishes there. This proves the required obstruction is zero. In the general many-leg construction the other independent contractions define coefficient subspaces that remain constant when this block varies; the coordinate homotopy preserves them. For two legs there are no other independent indices, so this preservation condition is vacuous.

Now put

$$V = X - C_-^2 D \in E^2 \otimes E^1. \tag{TEL2}$$

It still has order at least $s+1$. Its represented tensor is unchanged, since $C_-^1 V = u - C_-^1 C_-^2 D = u$. Its remaining contraction is removed, since $HV = HX - C_-^2 HD = HX - C_-^2 Y = 0$. Apply the independent-block division separately to the six E^2 coefficients of V. This gives $g \in E^2 \otimes E^2$, of order at least s, whose two curl contractions equal u. All identities hold at zero momentum as analytic identities, as well as at nonzero real momentum.

For clarity the width and norm losses can also be followed without any existence convention. Put $\operatorname{slim}_t(a) = a/(t+2)$, and, for a coordinate set of size m,

$$T_m(a) = \frac{16}{a}\left[\frac{2(1+a/4)}{a/4}\right]^m, \quad P(a) = T_4(a/2), \quad E(a) = a/16, \quad B_r(a) = \binom{4}{r+1} 4(e^a+1).$$

Here T_m is four times the coordinate-division price in (ED2); B_r bounds the finite sum defining a curl contraction on the strip. If u has strip width $\sigma > 0$ and coefficient bound $A \geq 0$, set

$$\tau = E(\sigma), \quad w = E(\tau), \quad \mu = w/16, \quad b = P(\sigma), \quad c = bB_0(\tau), \quad d = cP(\tau), \quad e = dT_4(w).$$

The bounds for X, HX, Y, D are respectively Ab, Ac, Ad, Ae; their widths are $\tau, \tau, w, w/4$. Restricting to μ, the bound for V is

$$A\{b + eB_2(w/8)\}. \tag{TEL3}$$

For the final independent division there is exactly one remaining leg. Apply the coordinate-wall lift (ED1)–(ED2) with the four independent coordinates to each E^2 component of V. The other

coefficient indices are passive, so taking the maximum of their coefficient bounds introduces no extra factor six. The construction gives explicitly

$$\eta = \mu/4, \qquad \|g\|_\eta \le A\{b + eB_2(w/8)\}T_4(\mu), \qquad \operatorname{ord} g \ge s. \tag{TEL4}$$

To see the contraction identity componentwise, denote the lift of the independent index by $W_a(b,p)$, where a is the dependent two-form index and b the independent two-form index. Then $\sum_b \operatorname{symbol}_1(p)_{b,i} W_a(b,p) = V_{a,i}(p)$. Contracting this equality with $\operatorname{symbol}_1(-p)_{a,h}$ and summing over a gives $u_{h,i}$ by (TEL2). This is precisely the two-curl representation, with no omitted polarization or identification of the two momentum slots. More generally repeated independent divisions obey the recursion $\eta_0(\sigma) = \sigma$, $K_0 = 1$, $\eta_{t+1}(\sigma) = \eta_t(\sigma/4)$, $K_{t+1}(\sigma) = T_m(\sigma)K_t(\sigma/4)$, where m is the ambient coordinate count. Each step preserves the contractions in the other independent blocks because their symbols do not vary with the divided block. For this application only $t = 1, m = 4$ is needed. These prices depend on the input strip, not on the torus period or on s. The tensor spaces, the harmonic-set check and the two identities in (TEL2) give the two-leg construction used in both loop-moment applications below.

8.5 Production of the fresh loop moments

The moment normalization of the fresh determinant loop is obtained before its path representation is issued. Let f be the common analytic two-field picture of the determinant and Haar term, R the normalized coarse metric row with $R(0) = 1$, and g the inherited-loop picture. The latter vanishes to order four in real external momentum. Set α equal to the kinetic readout of f and subtract α times the normalized metric picture. The one-axis identification proved above shows that this difference, and therefore

$$h = f - \alpha \operatorname{slots}(R) + g,$$

vanishes to order four. It retains its two gauge-contraction identities, reality and color invariance. Its analytic norm is bounded by $A + |\alpha| \operatorname{bedCharge}(\sigma, M) + B$, where A, B bound f, g and M bounds the metric row. Thus this normalization does not discard the inherited loop price.

Apply (TEL1)–(TEL3) with $s = 2$: the input has zero jets below order four, and its two-curvature coefficient has zero jets below order two. In particular the second-contraction correction uses the harmonic set $p = 0$ checked there; the coefficient norm and positive strip width have the period-independent prices in (TEL3).

Write the resulting coefficients, for fixed plaquette and color types, as a real Laurent family $P(z)$. Taking real parts replaces an analytic coefficient on real momenta by the average of it and its conjugate at opposite momentum. This preserves the represented real tensor and its vanishing order. Project its finite color coefficient vector onto invariant tensors; this is a fixed real linear projection, of finite absolute matrix-sum price, which leaves the already invariant represented tensor unchanged. It commutes with reading the absolutely convergent Laurent series and with its derivatives. Thus

$$\widehat{P}(p) = \sum_{z \in \mathbb{Z}^4} P(z) e^{ip\cdot z}, \qquad \widehat{P}(0) = 0, \quad D\widehat{P}(0) = 0, \qquad \sum_z |P(z)| e^{\eta|z|} < \infty \tag{LM0}$$

for a fixed positive output width η.

Exponential summability also gives $\sum |z_i P(z)| < \infty$, by spending part of η. The series may therefore be differentiated on the real momentum axes. Its value and first derivative at zero are $\sum P(z)$ and $i \sum z_i P(z)$. Equation (LM0) proves both sums zero, separately for each plaquette and color choice. The actual path issue enumerates these types and offsets. Its first leg is based at zero, its second at z, and each leg includes the four plaquette edges. Its length is $|z| + 8$, so the

path price differs from the Laurent price only by the fixed factor $e^{8\eta}$ and the finite type and color counts. The displacement in the moment is exactly z. Reindexing by this enumeration proves the summable zeroth and first unwrapped displacement moments of the issued atoms. These are the fresh-loop identities needed in the next subsection and in (TF17)–(TF19).

8.6 Production and transport of the inherited loop moments

The full-family identities (TF17)–(TF19) require absolute summation before truncation. We give the construction that transports these identities across a block. For a two-curvature atom a, let q_a be its coefficient with its plaquette and color types fixed, and let $\delta_a \in \mathbb{Z}^4$ be the displacement from its first leg to its second. Its length bounds $|\delta_a|$. A charged exponential bound implies

$$\sum_a |q_a| < \infty, \qquad \sum_a |q_a \delta_a^i| < \infty, \tag{LM1}$$

because $s \le C_\eta e^{\eta s}$. Thus all changes of summation below are absolute, including those with a displacement factor.

For one atom translate its first leg to the origin in the finite torus sum. Along a fixed route from zero to δ_a, telescope the second leg. Subtract its value and its first differences at the origin. The remaining discrete second-order Taylor difference can be written as a finite sum of products of differences, one on each curvature factor, by summation by parts along successive route edges. Explicitly the removed terms are

$$q_a \sum_x F_a(x) G_a(x), \qquad q_a \delta_a^i \sum_x F_a(x) \nabla_i G_a(x). \tag{LM2}$$

The discrete Taylor identity is exact on the torus: each link telescope is a translated finite difference and the sum of a translated product can be shifted by the same lattice translation. Repeating this along the two route indices generates at most the square of the route length in remainder terms. The exponential budget absorbs this square after any fixed positive width discount. If the full family has zero zeroth and first displacement moments, summing (LM2) gives zero separately for each type and color. Its total action therefore equals the sum of the two-difference remainder patterns. This proves the representation used to inherit a quiet loop; the subtraction is performed on the whole summable family, not on an arbitrary finite prefix.

Let $\psi_i(r, b; z)$ be the installed fine-gradient column for a coarse residue r, curvature type b, direction i, and unwrapped displacement z. These columns are the product of the fine frame inverse, the differentiated linear tower, and the coarse frame. Whenever the stationary-column construction uses two curvature legs, its dependent-momentum join is precisely (TEL1)–(TEL3), with the retained Taylor orders stated there. On harmonic input the unmodified tower has residue-independent total columns. Install it on a one-site coarse torus. Its output is then constant on the fine torus; a fine difference of that output is zero. The intertwining identity between the differentiated column and the fine difference consequently gives

$$\sum_z \psi_i(r, b; z) = 0. \tag{LM3}$$

The frame factors preserve the constant-output statement by their translation convolution identities. Odd blocking is the hypothesis that makes this centered-gradient readout agree with the installed fine difference. Exponential column bounds justify passing from the installed periodic identity to the absolutely summed cover columns.

After substituting these columns in a two-difference pattern, one fixed choice of its two output types has coefficient $u f_0(z_0) f_1(z_1)$, where $\sum f_0 = \sum f_1 = 0$ by (LM3). Its zeroth total is

$u(\sum f_0)(\sum f_1) = 0$. Its displacement is $(c_1 + z_1) - (c_0 + z_0)$, with fixed carried offsets c_0, c_1. Splitting its first moment gives

$$u\left\{(\sum f_0)\sum_z (c_1+z)^i f_1(z) - \sum_z (c_0+z)^i f_0(z)(\sum f_1)\right\} = 0. \tag{LM4}$$

The absolute first moments exist by the column width bounds. Summing over residues, types and patterns is justified by (LM1) and the charged bound on the full substituted family. Thus inherited loops have the exact two moment identities, with their block price L^{-2}, rather than merely small moments.

Joining two families means enumerating their disjoint union. Absolute sums add, so both moment identities are preserved. Blank initial loops have zero coefficients and satisfy them. The coefficientwise limiting argument (SH1)–(SH3) preserves them as well: first pay the common exponentially weighted tail, then pass to the limit on a finite list. This supplies the full-family justification used to turn retained moments into omitted-tail moments in (TF18). Any fresh loop must satisfy the same identities before it is joined; they cannot be inferred from gauge invariance or the quadratic-potential normalization alone.

8.7 The one-loop drift used by the finite histories

The loop calculation evaluates the drift of the very finite histories used later in shooting. The fresh and inherited moments are kept separate until their transport identities are combined. The coefficient $s = 11C_A/(24\pi^2)$ refers to the physical Lie-algebra and Fourier normalizations fixed earlier, and the error allowance is bought before the final cutoff. The subsequent calendar turns these local drift estimates into uniform running for the actual represented couplings.

The finite coupling drift comes from the committed path coefficient and the bounded scalar reset, not directly from a formal beta function. At a transition with block length L, put $\mu = L^{-1}$. The determinant calculation on the retained physical path defines a coefficient $\alpha_{N,i}$ with sign appropriate to descent from level $i+1$ to i. We explain the logarithm and its uniform remainder.

For a fixed color direction, the continuum pole of the relevant second external-momentum coefficient is

$$p_G(k) = 4C_A\left(\frac{k_1^2}{|k|^6} - 4\frac{k_1^2 k_0^2}{|k|^8} - \frac{1}{|k|^4}\right). \tag{DR1}$$

The coefficient includes the vector determinant with weight $1/2$ and the gauge determinant with weight -1. These are the same two determinant weights kept in the massive-symbol calculation of the trace. The finite symbol coefficient $f_\mu(k)$, on the angular cube $[-\pi,\pi]^4$, satisfies

$$|f_\mu(k)| \le \frac{A}{(|k|+\mu)^4}, \qquad |f_\mu(k) - p_G(k)| \le \frac{D}{|k|^3} + \frac{D\mu^2}{|k|^6} \quad (|k| \ge \mu), \tag{DR2}$$

where $A, D \ge 0$ depend on the fixed analytic widths, coefficient and ellipticity bounds, and not on L, the torus period or its terminal cutoff. The input paths have their fixed quadratic normalization and the allowed higher soft jets; the latter preserve the continuum pole.

Here is the matrix estimate yielding (DR2). At scale $T = |k| + \mu$, the zero-field massive matrix, one-field vertex, and two-field vertex have sizes $T^2, T, 1$. Coercivity bounds the inverse by CT^{-2}. Differentiate the identity $AA^{-1} = 1$ to obtain $D(A^{-1}) = -A^{-1}(DA)A^{-1}$ and its second derivative, the sum of the two products with two first vertices minus the product with one second vertex. Along a momentum direction scaled by T, all jets of each inverse still have size T^{-2}. The lower determinant bubble therefore has size T^{-2}, and its second unscaled external-momentum coefficient

has size T^{-4}. The upper insertion obeys the same rough bound. Summing the finitely many matrix entries gives the first bound in (DR2).

For $|k| \geq \mu$, take $T = |k|$. The zero-field symbol differs from its continuum quadratic part by $O(T^3 + \mu^2)$, and the first vertex differs from its continuum linear part by $O(T^2 + \mu^2/T)$. Use $A^{-1} - B^{-1} = A^{-1}(B - A)B^{-1}$ and telescope the factors of the determinant bubbles. Their relative error has size $T + \mu^2/T^2$; their second momentum coefficient consequently has error $C(T + \mu^2/T^2)T^{-4}$, which is the second bound in (DR2). The upper insertion contributes bounds of these same or smaller orders. The origin-jet cancellations, analytic derivative bounds, inverse formulas and explicit continuum vertex contraction are developed in the weighted massive-symbol and logarithmic momentum calculations below. This argument uses them before integration, so its error bound is absolute.

To integrate, split the angular cube into $|k| \leq \mu$, $\mu < |k| \leq 1$, and its part outside the unit ball. The first bound in (DR2) gives

$$\left| \int_{|k| \leq \mu} f_\mu(k)\, dk \right| \leq A\mu^{-4} \frac{\pi^2 \mu^4}{2} = \frac{\pi^2 A}{2}.$$

Outside the unit ball it gives a bound A times the volume of the angular cube. On the annulus the second bound gives

$$\begin{aligned} \int_{\mu < |k| \leq 1} \frac{D}{|k|^3}\, dk &= 2\pi^2 D(1 - \mu) \leq 2\pi^2 D, \\ \int_{\mu < |k| \leq 1} \frac{D\mu^2}{|k|^6}\, dk &= \pi^2 D(1 - \mu^2) \leq \pi^2 D. \end{aligned} \tag{DR3}$$

All these errors are bounded independently of $0 < \mu \leq 1$.

The pole integral is exact. Rotational symmetry gives, on this annulus,

$$\int \frac{k_1^2}{|k|^6}\, dk = \tfrac{1}{4} I_\mu, \qquad \int \frac{k_1^2 k_0^2}{|k|^8}\, dk = \tfrac{1}{24} I_\mu, \qquad I_\mu = \int |k|^{-4}\, dk = 2\pi^2 \log(1/\mu). \tag{DR4}$$

For the mixed fourth moment, symmetry makes all off-diagonal moments equal. Expanding the fourth powers of $k_0 + k_1$ and $k_0 - k_1$, and rotating these directions to a coordinate direction, shows that the diagonal fourth moment is three times the mixed one. Summing over the four coordinates then gives the factor $1/24$. Substitution of (DR4) into (DR1), with Fourier weight $(2\pi)^{-4}$, gives

$$(2\pi)^{-4} \int_{\mu < |k| \leq 1} p_G(k)\, dk = -\frac{11 C_A}{24\pi^2} \log(1/\mu) = -s \log L. \tag{DR5}$$

Equations (DR2)–(DR5) bound the difference between the infinite-volume coefficient and this logarithm by a pre-block constant.

We now connect this integral to the coefficient of the actual stationary path. It is important to distinguish three determinant films: M_{act}, the film of the actual stationary integration; M_{core}, its core film before the stationary carrier change; and M_{rank}, the auxiliary rank film. Each film is a matrix of analytic functions of the external momentum, with indices consisting of a direction and a color. Their scalar reading is

$$\alpha_a(M) = 2 \operatorname{Re} \mathrm{leapTip}(\mathrm{profitable}\, M)(e_0, 1, (a, a)). \tag{DP1}$$

The operation leapTip is the coefficient of degree two in that external-momentum direction, with its Taylor factorial; profitable passes from the determinant film to its curvature tensor using its

analytic contraction identities. Thus α_a is linear on films satisfying the common bounds, and is the specific normalized kinetic coefficient, not the full film norm.

Here are the quantitative comparisons used in the transition. Fix an analytic width $h > 0$, with $h \le \sigma \le 1/128$, a block patch of stage J, its penalty k, a positive stationary width, and all the common metric, covariance and path bounds. The required hypotheses are: the input path is attached to its quadratic metric and has its prescribed analytic bounds; that metric has the uniform floor on every torus; the coarse metric has its analytic strip bound; the stationary graph is critical, guessed by the installed tower, and satisfies the gap and heat identities; the same forest and tag are used in its determinant realization; and the relevant torus sides form an auxiliary sequence tending to infinity. These are precisely the linear-horizon and stationary-graph data already constructed. The input loop-tensor norm may have its own bound Λ, whereas the comparison constants below depend only on the fixed analytic and matrix bounds specified before the block length. They do not grow with Λ.

Let $F \ge 0$ bound the inverse matrices in the determinant calculation, let $M_* \ge 0$ bound the common zero-, one- and two-field film factors, and let $K_* \ge 0$ bound their carrier difference. All three are obtained from the fixed stage, kernel, penalty and common analytic bounds before L is chosen. More explicitly the common factor bound is the sum of the absolute bounds for the leg, core, rank and frozen metric factors; the carrier bound is the common bound for the difference between the core and rank factors. Put

$$D_* = FK_*F,\ E_* \ = D_*M_* + FK_*,\ R_* = E_* + (E_*F + FM_*D_*)M_* + FM_*FK_*. \tag{DP2}$$

This is the source's exact reserve polynomial. It arises by subtracting inverse factors with $A^{-1} - B^{-1} = A^{-1}(B - A)B^{-1}$, then telescoping the zero-, one- and two-field determinant products. For example the first inverse difference costs D_*; multiplying by the common first factor and adding the factor difference costs E_*; the remaining two placements and their cross product give the last line. The trace estimate is a relative rank estimate, so this subtraction introduces no count of all L^4 block sites.

In detail, if $r = |\operatorname{Loc}_0| + |\operatorname{Loc}_1|$, the individual films are bounded on width $h/2$ by

$$L^4\, 2r(FM_* + FM_*FM_*). \tag{DP3}$$

The separate difference estimate is

$$M_{\mathrm{core}} - M_{mrank} \quad \text{bounded on width } h/2 \text{ by } R_*. \tag{DP4}$$

Here is why (DP4) has no block-volume factor. The matrix weights are invertible analytic frames $S_X(p), S_Y(q)$. A bound for a matrix U means a Euclidean operator-norm bound for $\widetilde{U} = S_X(p)^{-1} U S_Y(q)$. For the difference estimates the stronger bound is used:

$$|\operatorname{tr}(\widetilde{U}V)| \le C\|V\|_{2\to 2} \qquad \text{for every compatible } V. \tag{DP10}$$

It is subadditive, is invariant under transposition with dual frames, and multiplies with an ordinary operator bound on either side. At equal source and target frames, taking $V = 1$ bounds the actual trace by C. A general matrix with operator norm C only has price $\dim(X)C$ in (DP10); using that conversion would incur the unwanted L^4.

The exact difference between the core and rank quadratic films cancels their common sector and pure Gram terms. Write S for the shear row, T for its localized emitter correction, and B for the fixed block inclusion. Expanding $(S+T)^*(S+T) - S^*S = T^*(S+T) + S^*T$ gives its reflection part. The complete remaining difference is the sum of the local house film and this reflection part, minus

$B^{*}\mathcal{O}B$ and the finite motive film; $\mathcal{O}$ is the localized orbit correction. This identity holds separately for the zero-, one- and two-field matrices, with all their Leibniz placements. The common sector cancels before any trace is bounded.

We derive the uniform trace-ideal estimates before using them in the difference. For a finite matrix U, set

$$\|U\|_{\mathrm{tr},*} = \sup_{\|V\|_{2\to 2}\le 1} |\operatorname{tr}(UV)|.$$

For $U = uv^{*}$, the trace is $v^{*}Vu$, so its size is at most $\|u\|\|v\|$. Decomposition into columns bounds the size by the sum of their norms. A matrix with a fixed n_0-dimensional source consequently has price $n_0\|U\|_{2\to 2}$. Multiplication on either side by a bounded matrix multiplies this price by its operator norm, by trace cyclicity. Apply these statements after the analytic changes of frame in (DP10).

One further factor has a source dimension growing with L, and needs an actual summability argument. Put $\mu = L^{-1}$, let ω be the lattice Laplacian symbol, and define

$$d_\mu(k) = \frac{\mu^2}{\omega(k)+\mu^2}, \qquad J = 16\sum_{n=0}^{\infty}(n+1)^{-2} < \infty.$$

On the strip of width $1/128$, choose the centered real-momentum representative before enumerating aliases. The massive lower bound gives

$$|d_\mu(k_l)| \le \frac{2}{(1+\operatorname{pdist}(l))^2}.$$

At $l = 0$ this is the origin massive bound. At $l \neq 0$, centered alias separation bounds the symbol scale below by $(1+\operatorname{pdist}(l))\mu$; strip ellipticity bounds the denominator below by half the square of that scale. Division gives the inequality. Changing the centered representative permutes the aliases. There are at most $(2n+1)^4$ aliases at distance n, because their lifts inject into $[-n,n]^4$. Therefore

$$\sum_{l} |d_\mu(k_l)^3| \le 8\sum_{n\ge 0}\frac{(2n+1)^4}{(n+1)^6} \le 8J.$$

The alias-diagonal matrix $d_\mu^3 I$ has trace-ideal price $8n_0J$, with fixed $n_0 = |\operatorname{Loc}_0|$. The Fourier frame and inverse frame each cost at most $e^{4/128}$, leaving the block-independent price

$$Q = e^{4/128}(8n_0J)e^{4/128}.$$

This origin estimate controls the zero-, one- and two-field matrices of the regularizing ninth power. Write $P = \mu^2(\mathcal{A}+\mu^2 I)^{-1}$, where $\mathcal{A}$ is the sector Gram matrix, and $U = P^3$. The preceding estimate applies to U at zero field. If C_3 bounds its three ordinary film matrices, the corresponding three matrices of U^3 have trace-ideal price at most $20C_3^2Q$. In fact every zero-field product has an unmarked U, every first-field product has two, and every second-field product has at least one. Put one unmarked factor in the trace ideal and the other two in operator norm. The second-field Leibniz formula has three placements of a second-field matrix and six ordered placements of two first-field matrices. Their sum and the lower orders are covered by 20. This never converts an L^4-dimensional bound into a trace bound. The regularizer is μ^2P^9, so its price is μ^2B, where $B = 20C_3^2Q$.

We now price each localized composition. We specify how the derivative prices used here are obtained. If an analytic finite-dimensional map F is evaluated at z, define, for a derivative bound $c \ge 0$,

$$E_m(F,z,c) = \sum_{n=0}^{m}\sum_{\pi\in\mathcal{P}_n}\|D^{|\pi|}F(z)\|c^{|\pi|}, \tag{DP12}$$

where $\mathcal{P}_n$ is the finite set of partitions in the iterated chain rule, with its multiplicities, and $\mathcal{P}_0$ contributes the value bound. If an inner map takes the value z at the origin and its derivatives through order m have the prescribed tested bound c, its composite is bounded by (DP12): differentiate successively, grouping the directions according to the block that enters each derivative of the inner map, and sum their norm products. The source's apollo and encore are precisely these finite sums. For a product use $F(u,v) = uv$ and inner price $c_1 + c_2$; inversion of a group matrix is its adjoint and preserves these derivative bounds. The same construction applies to the matrix exponential at zero, the chosen analytic group logarithm at the identity, and the analytic nearest-group projection in its fixed neighborhood.

These rules generate the charter and emission prices from the actual stencil. A link perturbation has its explicit factor L^{-1}. On a length-L path, the product rule chooses the differentiated links. If q distinct links are differentiated, there are at most L^q choices and their derivative bounds contribute L^{-q}; all undifferentiated group factors have norm one. The remaining assignments are the finite partitions of the derivative directions. Thus every fixed derivative order has a common bound independent of L. Patch averages are probability averages. Moving from the padded patch to its translated subpatch costs at most its fixed stage-dependent cardinality ratio, since both counts are bounded multiples of L^4. The frame estimates already proved give the corresponding finite derivative bounds for the recentered gauge transformations. Form the connector by multiplying the two endpoint frames and the path transport, average it, and apply the nearest-group projection; (DP12) pays all these operations. Multiply by the coarse-link factors and the endpoint frames, and apply the logarithm. This constructs the finite positive prices denoted springs for the charter and rocket$_f$ *are* for the emission in FrameEst/Stencils. Every logarithm is evaluated at the identity in this centered calculation. None of these steps takes a supremum over a number of arbitrary independent site values.

Here is the last conversion into row norms, including its block powers. Let $n \le 2$ be the number of inserted field directions and let c bound their fixed Fourier-frame readings. Put $c' = \mathrm{muralRoom}(J)c$, including the fixed stage enlargement of those readings. The scalar directional estimate for the charter is at most

$$2^{n+1} |\operatorname{springs}(J, n+1, c')|;$$

the corresponding emission estimate, after changing from the patch probability average to the block-site probability average, is at most

$$\mathrm{portrait}(J) 2^{n+1} |\operatorname{rocket_fare}(J, n+1, c')|.$$

The factor 2^{n+1} pays the real-and-imaginary expansions of the n inserted directions and the one tested input direction. Conversion of the finitely many color coordinates to their Euclidean norm multiplies these by $\sqrt{n_1}$ and $\sqrt{n_0}$, respectively. For the emission the probability-average estimate is a sum of squares divided by L^4; removing that denominator and taking a square root contributes exactly L^2. Taking the maximum, or the larger finite sum, over $n = 0, 1, 2$ gives prices a and h independent of L, with charter price a and emission price $L^2 h$. These are RealGap.CharterR.costs and Chromatic.costs. The same finite derivative and Fourier-frame bounds give the sector price p; the positive massive inverse and its differentiated inverse identities give the regularizer price $\mu^2 g$. The house kernel has fixed dimension and the already constructed inverse bound, giving its ordinary price b.

Each of a, h, p, g, b thus bounds the zero-, one- and two-field matrices on the same strip. The house kernel has fixed source dimension $n_1 = |\operatorname{Loc}_1|$, so its trace-ideal price is $n_1 b$. Film multiplication costs at most four times the product, covering its four second-field Leibniz terms. The

house price is consequently

$$C_H = 4(4an_1b)a.$$

Composition of the regularizer with the site-to-color row cancels $\mu^2 L^2 = 1$. Its trace-ideal price is $r = 4Bh$ and its ordinary price $4gh$. The shear row has ordinary price $4p$, so the corrected row has price $t = 4p + 4gh$. The two reflection products have combined trace-ideal price

$$C_R = 4rt + 4(4p)r.$$

For the orbit term let s_i bound the fixed stage row in direction i, and put $S(z) = \sum_{i=0}^{3} 4(4s_i z)s_i$. Its unweighted trace price is $\mu^2 S(B)$. The two changes to the dual massive frames cost at most $e^4\mu^{-1}$ each, yielding

$$C_O = e^4 S(B) e^4.$$

The motive film factors through the fixed n_0-dimensional space, with inclusions of ordinary price one; its trace-ideal price is $4n_0$. All widths and derivative orders have been fixed before L, and each cancellation of a block power has been displayed. These four estimates supply (DP10). Composing a difference factor with an ordinary bounded factor keeps (DP10). In particular the house film has price C_H; the two reflection products have combined price C_R; $B^{*}\mathcal{O}B$ has price $16C_O$, including its two film-composition factors; and the motive film has price $4|\operatorname{Loc}_0|$. Thus one may use exactly

$$K_* = C_H + C_R + 16C_O + 4|\operatorname{Loc}_0|. \tag{DP11}$$

Each price here is the absolute weighted column sum of its stated local factor with its zero-, one- and two-field placements, not a norm of a full trace. The analytic row and localized column estimates bound those sums by the fixed stage and strip prices. They use the same $\sigma \le 1/128$, branch bounds and common stationary rows as (DP3), and are uniform in L. The inverse identity gives difference price $FK_*F = D_*$. For a two-field determinant frame, the direct term difference costs $D_*M_* + FK_* = E_*$. Telescoping its three-factor bubble gives $(E_*F + FM_*D_*)M_* + FM_*FK_*$. Taking its trace through (DP10) gives precisely (DP2) and (DP4). The factor $1/2$ in the core-minus-rank determinant may be discarded when increasing this bound. This proves the relative projection estimate while retaining the L^4 factor in the unrelated individual estimates (DP3). The stationary carrier change has zero kinetic reading:

$$\alpha_a(M_{\mathrm{act}} - M_{\mathrm{core}}) = 0. \tag{DP5}$$

To see this, write its diagonal film as $\frac{1}{2}(W(p) + W(-p) - V(p)/2)$. The stationary tax row has zero jets below order two. Its mixed carrier row has zero jets below order one, by the stationary graph identity. Hence each summand of W, a product of these two rows, vanishes to order at least three. Every summand of V contains two tax rows and vanishes to order at least four. These films are analytic on the common strip, so their external degree-two coefficient is zero. Linearity proves (DP5).

For an analytic scalar function bounded by A on width r_0, two successive Cauchy derivative estimates give

$$\|D^2 f(0)\| \le \frac{A}{r_0/2}\frac{1}{r_0/4}.$$

For the normalized reading (DP1), the source keeps the slightly weaker explicit price

$$\mathcal{B}(r_0, A) = \frac{A}{r_0/4}\frac{1}{r_0/8}\|\operatorname{axes}(e_0)\|^2. \tag{DP6}$$

The factor two in (DP1) cancels the Taylor factorial. Apply this estimate to (DP4), with $r_0 = h/2$, and use (DP5). The exact projection bound is

$$|\alpha_a(M_{\mathrm{act}}) - \alpha_a(M_{mrank})| \le \mathcal{B}(h/2, R_*). \tag{DP7}$$

No inverse power of the lattice period and no L^4 multiplier occurs in (DP7).

There is one more comparison, which is an equality rather than a bounded error. Let Π_μ be the analytic film of the massive pole test on the fine lattice, with $\mu = 1/L$. Its finite-period analytic tensors approximate that test with error $\mathrm{cents}(n)e_n$, where $e_n \to 0$, and the rank film approximates the same test pulled back through the stationary uniform map with error $\mathrm{cents}(m)y_m$, where $y_m = Ce^{-cm}$ for fixed L. The constants and analytic strip in these approximation statements may depend on L. They are used for identification only, not as the uniform bound in (DP7).

Fourier evaluation of the stationary pullback gives the exact alias formula

$$H(p) = L^4 \sum_{\ell} S((p + 2\pi\ell)/L)[n_-^\ell(p), n_+^\ell(p)]. \tag{DP8}$$

Here $S = \Pi_\mu$, and $n_\pm^\ell$ are the incoming and outgoing stationary response columns. The analytic contraction identities of S remove longitudinal columns. For $\ell \neq 0$, the remaining response columns each vanish to order two at $p = 0$; their product therefore has no degree-two coefficient. For $\ell = 0$, the primal and second-order stationary identities identify its contribution, through degree two, with $\mu^2 S(\mu p)$. The discrepancy vanishes to order three. Taking two momentum derivatives introduces another μ^2, so the coefficient multiplier is $L^4\mu^2\mu^2 = 1$. Consequently

$$\alpha_a(M_{mrank}) = \alpha_a(\Pi_\mu). \tag{DP9}$$

The finite-period approximation just described proves that the analytic pullback in (DP8) is the rank film: let the auxiliary coarse and fine sides tend to infinity at fixed L, keeping the same stationary graph and tags. The vanishing approximation errors identify their Fourier tensors and derivatives. The reality, gauge and signed symmetries of the actual pole test give its contraction identities and zero lower jets. No assumption about a momentum continuation between different lattice meshes is used.

Combining (DP7), (DP9), and (DR2)–(DR5) now gives the exact allowance used in the transition:

$$d = |\mathcal{B}(h/2, R_*)| + D_{\mathrm{ann}} + 1, \qquad |\alpha_{N,i} + s\log L_i| + 1 \le d, \tag{DR6}$$

where $D_{\mathrm{ann}} \ge 0$ is the bound for the massive-symbol integral remainder. The stage construction fixes $h, F, M_*, K_*, D_{\mathrm{ann}}$, and hence d, before the block length. Afterwards it imposes one lower bound on odd L, and post-block bounds on the analytic norm and loop allowance of the output path. In the transition notation these latter bounds have the forms $\Theta \ge \beta_*(L)$, $c' \le \gamma_*$, and $\Lambda' \ge \Lambda_0(L)$, $\Lambda' \ge \Lambda/2$. They may depend on L; they do not enter d. The stationary tensor is produced using arbitrary growing torus sequences satisfying the stated floor and stationary identities; after coefficient identification it is one cover path and works on every realized torus satisfying the construction's patch and support margins. Thus the auxiliary large-volume limit imposes no extra period-dependent error on (DR6), and the estimate is uniform over the terminal cutoffs whose paths have those same pre-block bounds. The scalar reset has $|\Delta| \le 1$. Since its actual coarse coupling is $\beta_{N,p,i} = \beta_{N,p,i+1} + \alpha_{N,i} + \Delta$, we obtain

$$\begin{aligned} |\beta_{N,p,i+1} - \beta_{N,p,i} - s\log L_i| &= |\alpha_{N,i} + \Delta + s\log L_i| \\ &\le |\alpha_{N,i} + s\log L_i| + 1 \le d_{i+1}. \end{aligned} \tag{DR7}$$

Thus the one-step drift used in the shooting and running arguments includes both the calculated path coefficient and the actual nonlinear reset. The choices of large block lengths and late coupling entry below pay this drift simultaneously with all the local activity and insertion bounds.

8.8 Massive auxiliary symbols and the uniform alias trace bound

The massive comparison used in the preceding determinant argument is fixed as follows. With the lattice exterior symbol $q_a(p) = e^{ip_a} - 1$, put

$$s(p) = \sum_a q_a(-p)q_a(p), \quad d_\mu(p) = \frac{\mu^2}{s(p) + \mu^2}, \quad c_\mu(p) = \mu^2 d_\mu(p)^9, \quad \mu = L^{-1}. \tag{MA1}$$

On real momentum $s(p)$ is the nonnegative squared curl frequency. The scalar massive covariant-site operator at a group field is $D_U^* D_U + \mu^2$; its normalized inverse is $\mu^2(D_U^* D_U + \mu^2)^{-1}$. The auxiliary regulator is μ^2 times its ninth power, and the comparison site operator is $D_U^* D_U$ plus that regulator. These are ordinary finite positive matrices when $\mu > 0$. Their symbols at the identity are respectively $s + \mu^2, d_\mu, c_\mu, s + c_\mu$. The ninth power is part of the construction, not a variable exponent sent to infinity.

For link fields, copy the site regulator into each of the four link directions at the link's initial site. Add it to the physical quadratic Hessian and to $D_U D_U^*$. This defines the positive auxiliary link matrix. The associated scalar functional is one half of its logarithmic determinant minus the logarithmic determinant of the site comparison operator. The copying operation respects gauges and unsigned permutations of axes. Reversed links move the initial site, so this auxiliary functional is not assigned a signed symmetry by that copying argument. The physical signed laws are established separately.

The symbol at the identity is the physical transverse quadratic symbol plus the longitudinal symbol $qq(-p)^T$ plus $c_\mu(p)$ times the identity. On real momentum the gap estimate bounds its quadratic form below by

$$\frac{\min(\epsilon, 1)}{2^{10}}(s(p) + \mu^2)\,\|v\|_2^2, \tag{MA2}$$

where $\epsilon > 0$ is the physical transverse floor. The scalar companion has the same bound without the transverse factor. Indeed the physical and longitudinal pieces bound their respective curl and gradient components; in the small-frequency region $s \le \mu^2$, $d_\mu \ge 1/2$, so $c_\mu \ge \mu^2/2^9$; in the other region the frequency term already pays a fixed multiple of $s + \mu^2$. This proves the required positive floor including zero momentum.

To see the uniform trace summation, aliases are

$$p_l = L^{-1}p + 2\pi L^{-1}\widetilde{l}, \qquad l \in (\mathbb{Z}/L\mathbb{Z})^4, \tag{MA3}$$

using the fixed centered integer lift $\widetilde{l}$. Write $r(l) = \max_a |\widetilde{l}_a|$. There are at most $(2n+1)^4$ indices with $r(l) = n$, by injection into the integer box. Consequently, uniformly in L,

$$\sum_l (1 + r(l))^{-6} \le J_0, \qquad J_0 = 16 \sum_{n \ge 0} (n+1)^{-2} < \infty. \tag{MA4}$$

This bound uses a box estimate for each shell; its excess over the actual shell size is harmless.

For odd L, on the strip of imaginary width $1/128$, first choose the real part of the coarse momentum in its fixed reciprocal fundamental cell. The lattice sine lower estimate on every nonzero alias and the complex perturbation estimate for $s + \mu^2$ give

$$|d_\mu(p_l)| \le \frac{2}{(1 + r(l))^2}. \tag{MA5}$$

At the zero alias the massive denominator directly gives the bound two. For a nonzero alias, its real frequency is at least $(r(l)+1)\mu$ in the norm used by the sine estimate, and $\frac{1}{2}((r(l)+1)\mu)^2 \le |s(p_l)+\mu^2|$; divide μ^2 by this lower bound. A general coarse momentum differs by a reciprocal integer translate; this permutes the alias set. Thus summing (MA5) cubed gives, on that whole strip,

$$\sum_l |d_\mu(p_l)^3| \le 8J_0. \tag{MA6}$$

In particular the diagonal alias matrix $d_\mu^3 \otimes 1_{\mathbb{C}^h}$ has nuclear norm at most $8hJ_0$. This is the three-massive-arc trace bound used inside the ninth-power regulator.

Here are the frames in which it is used. Let $T_L(p)$ have entries $L^{-2}e^{ip_l\cdot\widetilde{x}}\delta_{\alpha\beta}$, with row x,α and column l,β. The finite character sum shows that it is unitary at real p; its explicit inverse is the opposite-phase character matrix. On the fixed complex strip the two matrices have operator norms at most $e^{4/128}$, by separating their real unitary Fourier factor from the diagonal complex phases. Furnishing an alias symbol means conjugating its diagonal alias blocks by these two matrices. Thus its three-arc nuclear bound is at most

$$e^{4/128}\,(8hJ_0)\,e^{4/128}. \tag{MA7}$$

The weighted frames additionally use

$$h_l(p)=\sqrt{\mathrm{pebWeight}(p_l,\mu)}, \qquad A_L(p)=T_L(p)\,\mathrm{diag}(h_l(p)^{-1})T_L(0)^{-1}. \tag{MA8}$$

Here $\mathrm{pebWeight}(p,\mu)=\sum_a |e^{i\,\mathrm{Re}\,p_a}-1|^2+\mu^2$. Thus the positive weight uses the real part of momentum even when p is complex. Such a weight is a norm-comparison device; it is not asserted holomorphic. For complex differentiation, freeze the weights at the real anchor and use the analytic $T_L(p)$ factors. Comparing the original and frozen weights at nearby momenta gives the fixed height allowance. The link-plus-site frame is the direct sum of (MA8) on link components and the unweighted site frame.

This separation justifies extending the real floor to a complex strip without a hidden dimension factor. Test the conjugated matrix against two fixed vectors and apply the one-variable Cauchy estimate to the resulting scalar analytic function. If its operator norm is bounded by M on a ball of radius R, its derivative is bounded by $2M/R$ on the half ball. Integrate along the segment from the real anchor to a nearby complex momentum. Choose that distance so that the perturbation norm is less than half the real lower bound in (MA2). The Neumann series then inverts the perturbed matrix with at most twice the anchored inverse bound. Apply this in the frozen frames, then restore the compared positive weights. All constants are fixed before L; none comes from summing matrix entries to estimate an operator norm.

The derivative films used above are the actual zeroth, first and second coefficient matrices of the local analytic rows. A row's real directional derivative bound is first converted componentwise, and then its color-coordinate square sum is taken. For the link row this costs exactly $\sqrt{|\mathrm{Loc}(G,1)|}$; for the site row it costs $\sqrt{|\mathrm{Loc}(G,0)|}$ and the displayed local sampling-density factor. Pulling the site row into the block adds L^2, since the normalized site square sum has L^4 terms. These are the link and site row allowances used in the weighted operator bounds. The rows are differentiated through third order when constructing a second-order film of their first derivatives; that shift of order is retained.

For the regulator film, form three consecutive cubes of the normalized inverse film in (MA1), giving $d_\mu^3 d_\mu^3 d_\mu^3 = d_\mu^9$, and then multiply by μ^2. At the common center this is exactly c_μ. Its trace-ideal bound uses one unsplit cube as (MA7), while the other factors receive operator bounds. In

a first or second product derivative, choose the corresponding three-arc occurrence after applying the exact inverse and product rules; the column placements in (DP10)–(DP12) give the listed alternatives. The resulting finite sum is bounded by the positive sum of their tariffs.

Finally the house term is the physical Hessian film sandwiched between its two row films; the orbit term is the sum over four initial-site copies of the regulator film. Sandwiching a trace-ideal factor by operator-bounded factors preserves its nuclear bound, with the product of the two operator allowances. Sum the four copied terms and conjugate to the common frame; the fixed Fourier norms cost at most e^4 on each side at the wider intermediate allowance. The reflection correction is the sum of its two cross products with the regulated site row, and the fixed extra block contributes $4|\operatorname{Loc}(G,0)|$. Addition and the same two-sided multiplication rule give the full trace toll in (DP10)–(DP12). Thus the cancellation there is backed by the massive symbol (MA1), the alias sum (MA6), the real and complex floors, and the actual local row derivative bounds. It does not require treating an arbitrary large block matrix as uniformly trace class.

The complex denominator estimate used in (MA5) can be proved before alias summation. Put $x = \operatorname{Re} p$, $u = \|q(x)\|_2$, and $t = \|p - x\|_\infty \le 1$. Integrating the derivative of the exponential gives $\|q(p) - q(x)\|_2 \le 4t$, and the same estimate at $-p$. Expanding their scalar product yields

$$|s(p) - s(x)| \le 4t(u + 4t) + 4tu = 8tu + 16t^2.$$

For $t \le \mu/128$, the right side is at most $(u^2 + \mu^2)/2$: use $2u\mu \le u^2 + \mu^2$. Since $s(x) = u^2$, the reverse triangle inequality proves

$$|s(p) + \mu^2| \ge \tfrac{1}{2}(u^2 + \mu^2) > 0, \quad |d_\mu(p)| \le 2, \quad |c_\mu(p)| \le 2^9\mu^2. \tag{MA9}$$

The coarse strip has width $1/128$; each alias divides its imaginary part by L, so its fine momentum has exactly the width $\mu/128$ required here. In particular the inverse denominator is analytic throughout the strip.

For completeness the weights in (MA8) obey

$$h_l(p) \ge \mu, \qquad \|q(p_l)\|_2 \le 5h_l(p)$$

on the unit coarse strip. At coarse points of distance at most one, the preceding exponential estimate gives $\|q(\operatorname{Re} p_l) - q(\operatorname{Re} r_l)\|_2 \le 4\mu$. Using $h_l \ge \mu$ shows that the two positive weights have ratio bounded by a fixed constant in both directions. Freeze h_l at a real anchor while retaining $T_L(p)$ and its explicit analytic inverse. These are the analytic frames used in the Cauchy estimate. In ordinary unweighted norm their bounds are

$$\|A_L(p)\| \le e^{4\sigma}\mu^{-1}, \qquad \|A_L(p)^{-1}\| \le e^{4\sigma} \sup_l h_l(p). \tag{MA10}$$

For a real anchor, $h_l^2 \le 16 + \mu^2 \le 17$. These estimates display the mass powers; they must not be silently replaced by two mass-independent unweighted bounds.

Here is an explicit form of the strip-inverse argument. Suppose a matrix K, in the two specified weighted frames, has norm at most A, its real lower singular bound is $a > 0$, and the changes to frozen analytic frames cost at most C on a ball of radius σ. Test the analytic matrix on two unit vectors and apply the scalar Cauchy estimate. This gives derivative bound $k = CAC/(\sigma/2)$ on the half ball. On the strip

$$\eta = \min\{\sigma/2,\ (a/2)/(1 + k)\}$$

the displacement from its real anchor changes the matrix by at most $a/2$. The lower singular bound is therefore $a/2$, so its weighted inverse has norm at most $2/a$. This also proves invertibility

and analyticity of the inverse by the determinant formula. The constants are independent of the dimension of the block. In the site application the real lower form bound is proportional to h_l^2, and the two inverse weights cancel it; it is that weighted matrix to which the assertion applies.

The remaining mass bookkeeping can now be stated exactly. The normalized inverse film has bounded ordinary operator coefficients because its factor μ^2 cancels the two inverse-weight losses. The regulator film has both operator and trace-ideal price μ^2 times a fixed price, by its definition as $\mu^2(d_\mu^3)^3$. Converting it back to the weighted form costs at most μ^{-1} on each side, hence cancels this μ^2. This accounts for the mass-independent regulator toll in the link determinant comparison. It does not remove the explicit μ^2 in its ordinary frame.

Finally the actual matrices behind the last paragraph of this chapter are as follows. Let B be the fixed inclusion of the link block and let P be the complementary site projection; let S_i read the initial-site slot of direction i. Write C for the regulator film, R for the differentiated row, and H for the local second variation of the scalar correction. With the rear operation denoting the transpose at reversed momenta, including the reversed external legs, the relevant products are

$$\text{House} = R^{\text{rear}} H R, \qquad \text{Orbit} = \sum_{i=0}^{3} S_i^{\text{rear}} C S_i, \qquad \text{Motive} = P^{\text{rear}} P. \tag{MA11}$$

These are exact products of films, so their first and second coefficients are their ordinary ordered product derivatives. The motive is constant and its finite site rank is the reason its trace allowance is $4|\operatorname{Loc}(G,0)|$. The orbit allowance is the regulator trace allowance multiplied on each side by the Fourier-frame bound. The house allowance follows by first assigning the finite local trace allowance to H, then multiplying by the two row bounds. The two reflection cross products are bounded in the same order. The total toll is the positive sum

$$\text{payLedger} + \text{reflectLedger} + 4(4\,\text{orbitLedger}) + \text{popLedger}. \tag{MA12}$$

Each factor in this expression has thus been paid either by a local finite rank, the alias trace sum, or a weighted operator norm; no trace norm of a general extensive Hessian is assumed.

8.9 Constructing the common asymptotic calendar

Every finite stage contains only finitely many requests of its admitted degree and derivative order. Their maximum is finite. A countable family is met by letting the stage order increase, so that each fixed request is eventually present. This is weaker than a uniform estimate at all orders at one stage, and is exactly what the diagonal construction needs.

We give the numerical schedule construction explicitly, since both finite shooting and uniform running use it. For each stage $j \ge j_0 \ge 4$, all pre-block parameters have already been chosen: its derivative order, support exponent, analytic width, path bound, column bound and drift allowance $d(j)$. Let

$$A(j) = \log\Bigl(1 + \text{paths}(j) + \text{width}(j)^{-1} + \text{columns}(j) + \text{relaxation}(j) + \text{supportExp}(j)\Bigr).$$

Choose an odd integer block length $\mathscr{L}_j \ge 3$, at least every previously required lower bound, such that

$$\frac{d(j)}{\log \mathscr{L}_j} \le \min(s/4, j^{-2}), \qquad \frac{A(j)}{\log \mathscr{L}_j} \le j^{-2}. \tag{CAL1}$$

This is possible because the numerators are fixed before this choice and $1/\log n \to 0$. To make the block lengths nondecreasing, take for each j the maximum of the finitely many lower bounds

through j, then an odd integer larger than that maximum. These choices prove the drift and logarithmic-cost inequalities without making the drift allowance depend on the selected terminal cutoff.

After fixing the block lengths choose the positive protection exponents, decreasing with j, small enough for all post-block requests and at most $1/j$. This can be done by taking successive finite minima of positive admissible upper bounds. Next choose the stage's coupling-entry and scale-entry thresholds; every analytic estimate proved with these parameters has one eventual threshold, so a finite maximum handles all its requests. Later increases of these thresholds preserve their inequalities.

Choose the initial coupling $\bar\beta > 1$ above the first-stage entry threshold, and sufficiently large that

$$x^{-1/2}\exp((\log\log x)^2) \le 1 \qquad (x \ge \bar\beta). \tag{CAL2}$$

The ratio $(\log\log x)^2/\log x \to 0$ proves this. For fixed j, the required lattice-size bound at logarithmic scale u has the explicit form

$$\begin{aligned} \mathcal{V}_j(u) &= S_j(1 + c(u) + c(u + \log \mathscr{L}_j))^{q_j} 100(c(u)+1) \\ &\quad + 100((j+3)\mathscr{L}_j + 1), \\ c(u) &= \left\lceil \log^2(\bar\beta + su) \right\rceil. \end{aligned} \tag{CAL3}$$

Here S_j and q_j are the fixed shell-count and growth requests from the local estimates. The last two terms include the small-field and patch buffers. For each fixed stage, $c(u)$ and its fixed translate are $O(u)$; hence $\mathcal{V}_j(u) = O((1+u)^{q_j} + u + 1) = o(e^u)$. Choose an integer entrance bound B_j such that $m \ge B_j$ and $u \ge m\log 3$ imply simultaneously

$$\begin{gathered} \bar\beta + su \ge \mathrm{betaEntry}(j), \qquad u \ge \mathrm{scaleEntry}(j), \qquad u > 0, \\ \log \mathscr{L}_j \le u/j, \qquad \mathcal{V}_j(u) \le e^u. \end{gathered} \tag{CAL4}$$

Each condition holds eventually as $u \to \infty$; $m\log 3 \to \infty$ gives one common integer B_j.

There are also stagewise comparison-window requests $w_j(u)$ satisfying $w_j(u)/u \to 0$, and finite readiness thresholds r_j. Choose integers C_j so that for every $k \ge C_j$ and $k\log 3 \le u \le k\log \mathscr{L}_j$,

$$r_j \le u, \qquad w_j(u) \le k/2. \tag{CAL5}$$

Indeed eventually $w_j(u)/u \le (2\log \mathscr{L}_j)^{-1}$; multiplication by the upper bound on u proves the second inequality, and the lower bound on u supplies readiness. Integer rounding only strengthens the resulting bound by $\lfloor k/2 \rfloor$.

We now build the stage clock, rather than presupposing one meeting all these requests. Use an unshifted stage number $h \ge 0$, corresponding to $j = j_0 + h$, and write B_h, C_h for the chosen integer requests after this shift. Define plateau starting indices recursively by

$$a_0 = 0, \quad k_h = C_h + 2a_h + 1, \quad a_{h+1} = a_h + B_{h+1} + 2k_h + 1. \tag{CAL6}$$

Set $E(i) = h$ when $a_h \le i < a_{h+1}$. The sequence a_h is strictly increasing, so E is defined everywhere, nondecreasing, increases by at most one in a step, and tends to infinity. Also $E(0) = E(1) = 0$. Its recursion gives two useful properties:

$$E(i+1) > 0 \implies B_{E(i+1)} \le i, \tag{CAL7}$$

and, for every h, $k_h \ge C_h$ and

$$E(i) = h \quad \text{whenever} \quad k_h - \lfloor k_h/2 \rfloor \le i \le k_h + \lfloor k_h/2 \rfloor. \tag{CAL8}$$

For (CAL7), the first index of plateau $h > 0$ is $a_h = a_{h-1} + B_h + 2k_{h-1} + 1 \geq B_h + 1$. For (CAL8), $k_h = C_h + 2a_h + 1$ places its lower endpoint above a_h, while $a_{h+1} > 2k_h$ places its upper endpoint strictly before a_{h+1}. Thus every stage receives a sufficiently long central comparison window, however large its requests.

Define

$$J(i) = j_0 + E(i), \qquad L_i = \mathscr{L}_{J(i+1)}, \qquad \tau_0 = 0, \quad \tau_{i+1} = \tau_i + \log L_i. \tag{CAL9}$$

Then $i \log 3 \leq \tau_i \leq i \log \mathscr{L}_{J(i)}$, by monotonicity of the blocks. The first inequality proves $\tau_i \to \infty$. For every noninitial plateau (CAL7) lets us apply (CAL4) at $u = \tau_i$. In particular

$$0 \leq \frac{\log L_i}{\tau_i} \leq J(i+1)^{-1} \longrightarrow 0. \tag{CAL10}$$

This is the gentle-mesh condition. Equations (CAL1) and (CAL9) give (FCA4) and (RUN3). The stage order, power and support exponent stay nondecreasing, the protection exponent decreases, and (CAL2) gives the small-field threshold at every scale. These properties all belong to the same calendar.

For volume, let $d_i = \prod_{r<i} L_r = e^{\tau_i}$. On every noninitial plateau (CAL4) gives $\mathcal{V}_{J(i+1)}(\tau_i) \leq d_i$, so any positive physical period pays its lattice-size requirements. Only finitely many steps belong to the initial plateau, since $E \to \infty$. Choose an integer M larger than their finitely many size requirements and the required minimum period, and set the physical periods to

$$P_p = 4(M + p + 1), \qquad \text{lattice side at } (i,p) = P_p d_i. \tag{CAL11}$$

These periods are strictly increasing and divisible by four. They pay the initial exceptions directly and the later requirements by the preceding denominator bound. They also provide arbitrarily large periods divisible by any prescribed positive integer c: take $p = c(M + z + 1) - (M + 1)$, which is at least z, and then $c \mid P_p$. The selected complete construction may subsequently retain a period subsequence; it keeps all the size inequalities.

Finally use (CAL5) at k_h, τ_{k_h}, which lies between the two bounds preceding (CAL10). The requested window has width at most $k_h/2$, and (CAL8) keeps that entire window on stage $j_0 + h$. Thus the comparison windows, entry delays, local volume margins, gentle mesh and asymptotically small drift are realized together. No selection of a new calendar for a particular history, cutoff, period or observable is required.

8.10 Countable local demands on the same calendar

Why countably many demands can be met At any one stage only finitely many derivative, degree and support requirements have been activated. Their maximum determines a finite threshold at that stage. The increasing stage sequence eventually activates every fixed finite request, while later delays preserve the choices already made. This is a diagonal schedule of finite demands, not one estimate uniform in every degree and derivative order at once. The order of choices below is the mechanism that makes these quantifiers compatible.

The calendar construction must also accommodate estimates whose numerical values use the chosen block lengths. We describe this ordering explicitly. First choose the stage profiles, kernel, primary analytic price list, and the secondary price list for the flat covariance and exterior constructions. The three local constructions then supply finite integer lower bounds $f_{\mathrm{flat}}(j), f_{\mathrm{curv}}(j), f_{\mathrm{dual}}(j)$ for block lengths at profile j. Two further integer lower bounds $f_{\mathrm{stat}}(j), f_{\mathrm{comp}}(j)$ ensure the local stationary chart and its comparison construction are available. Raise the original length floor to their pointwise maximum:

$$f(j) = \max\{f_{\mathrm{original}}(j), f_{\mathrm{flat}}(j), f_{\mathrm{curv}}(j), f_{\mathrm{dual}}(j), f_{\mathrm{stat}}(j), f_{\mathrm{comp}}(j)\}. \tag{CAL12}$$

The choices of these supplier functions use the original stage profiles and price lists. Raising the floor changes neither that profile nor either list. Apply (CAL1)–(CAL11) with this floor to select the block lengths once. They now determine the local patches at all later steps.

For each fixed requested finite horizon E and profile j, the local construction gives a positive width $\rho_{E,j} \le 1$. Its reciprocal is a countable numerical demand. More generally, let $F : \mathbb{N}^2 \to \mathbb{R}$ be any already selected demand array. Define the finite envelope

$$H_F(j) = 1 + \sum_{n=0}^{j} \sum_{i=0}^{j} |F(n,i)|.$$

Increase only the coupling-entry delay at profile j to at least $\exp(\exp H_F(j))$. At every actual step m, with outgoing profile $j = J(m+1)$, the admissible clock then obeys

$$H_F(J(m+1)) \le \log\log \widehat{\beta}_m. \tag{CAL13}$$

For each fixed n, the divergence and monotonicity of J give a first step after which $n \le J(m+1)$. For all such steps and every $i \le J(m+1)$, the corresponding summand in the finite envelope yields

$$|F(n,i)| \le 1 + \log\log \widehat{\beta}_m. \tag{CAL14}$$

This buys every fixed demand eventually. It makes no assertion that infinitely many horizons have been paid at one finite step. Apply it first to $F(E,j) = \rho_{E,j}^{-1}$.

After the block lengths have been fixed, the local stationary and comparison estimates also have finite coupling bills depending on the resulting patches. Let $B_1(j)$ and $B_2(j)$ be their numerical bounds, and put $B(j) = \max(B_1(j), B_2(j))$. Add this function to the entry delay by taking another maximum. This gives, including at step zero,

$$B(J(m+1)) \le \widehat{\beta}_m. \tag{CAL15}$$

If these bills also have a countable horizon parameter, apply the finite-envelope construction to that array as well, in addition to the all-step bound for its distinguished horizon. Every such enlargement is made with the lengths held fixed. The requirement is therefore a numerical delay condition, rather than a circular demand to choose a block length larger than a function of a subsequently reselected block length. Intersecting finitely many delay conditions is achieved by successive maxima. Each persists under all later compatible enlargements, so the earlier treasury, period, local geometry and observable conditions survive.

The additional local calendars use exactly the same finite-max construction. To spell out their quantitative admission, enumerate their finite requests by natural-number codes. Request q specifies an integer precondition $p_q(j)$ and a finite real cost $c_q(j,i;L_j,E_i)$ for a legal adjacent pair $i \le j \le i{+}1$, where E_i is the fixed stage profile. Before selecting lengths, include

$$p_*(j) = \max_{\mathrm{code}(q) \le j} p_q(j) \tag{CAL16}$$

in their floor; an empty maximum is zero. After those lengths are fixed, there are only finitely many costs with $\mathrm{code}(q) \le j$ and $i \le j$. Increase the entry delay until all are at most

$$z(\widehat{\beta}) = 1 + \log(2 + \log \widehat{\beta}). \tag{CAL17}$$

This is possible because $z(v) \to \infty$. For example a sufficiently large double exponential of their finite maximum suffices. The legal flight in each cost is the one formed from the already selected

L_j, patch, kernel and stage; it is not a cost obtained from an independently chosen history. At an actual step m, substitute $j = J(m+1)$, $i = J(m)$ and $L_j = L_m$. Monotonicity of J implies that a request whose code is at most $J(n)$ has both its floor and its numerical cost paid at every $m \geq n$. In particular every fixed request is paid on its admitted tail. All the post-length maxima are held fixed when increasing the delays again, so the statement persists under every subsequent compatible increase.

For the stationary comparison, the finite-list assertion has a concrete index bound. A request records an input degree d, an extended degree $D > d$, and a jet order k, and is admitted only when $2(D+k+1)$ is at most the incoming smoothing grade. At outgoing grade j, there are consequently only finitely many such triples and adjacent predecessors. The maximum of their local length thresholds is the integer floor used in (CAL12). On the fixed patch, take a second finite maximum, at least one, of their polynomial estimate constants, the local stationary constants, and all reciprocal widths $\rho_{E,i}^{-1}$ with $E, i \leq j$. Call it R_j. Its post-length request is e^{R_j}; payment by $\widehat{\beta}_m$ is exactly $R_j \leq \log \widehat{\beta}_m$. Every auxiliary derivative index in a fixed triple ranges over finite intervals, so it introduces no supremum over arbitrarily high orders at one stage. This proves the finiteness of the native stationary and comparison prices used above, at their actual chosen lengths.

This includes the finite jet, local coordinate and path readouts together, with one selected system of their dictionaries and bases. The inverse-map request has zero precondition. Its actual cost is the inverse estimate of the specified local patch map at its selected analytic width and axis forest. It can therefore be paid from step zero: at each j, maximize its finitely many adjacent predecessor costs and use (CAL17). One does not need to assume a growth law for this inverse price. For a reporting window based at level k, the affine running reference coefficient gives

$$z(\widehat{\beta}_m) \leq z(\widehat{\beta}_k) + s|t_m - t_k| \leq (1+s)\{2 + Z_k + |t_m - t_k|\}, \tag{CAL18}$$

where Z_k includes $z(\widehat{\beta}_k)$. To obtain the first inequality, differentiate z on $v > 1$: $z'(v) = 1/[v(2 + \log v)] \leq 1$, and use $|\widehat{\beta}_m - \widehat{\beta}_k| = s|t_m - t_k|$. Thus the paid inverse cost has the same polynomial reporting-age control required by the local comparison estimates, on the actual windows of the actual history.

The common reference-period registry is fixed at the same pre-length stage. At outgoing grade j its finite list consists of horizons $e \leq j$ and the allowed adjacent incoming grades. Their integer size requirements are included in (CAL16). With the resulting lengths fixed, its uniform finite-patch, chart and normalization prices give a real entry gate for that same list, which is included among the later delay maxima. This retains one registry for every eventual horizon and every reference torus used in the periodic-copy argument; a later numerical enlargement does not reselect it. The physical periods themselves are still chosen as in (CAL11).

We explain the identification with the actual run, since it is needed for the order of quantifiers in the theorem. The run's stage profile is the original profile, its kernel is the original kernel, and its factor at profile j is the length just selected. Hence (CAL12) at $J(m+1)$ is the actual step's local length inequality. Its stationary horizon satisfies the original primary price list. Projecting the two maxima in (CAL15) gives each local patch payment, and therefore its stationary admissibility and the delay required for the flat covariance limit. The preselected flat supplier may now be applied with entrance index zero to this very horizon; the curvature and dual suppliers apply to it by the other two floors. Their local results concern the same levels, patches and retained horizon as the run. No new family is chosen when these numerical payments are consumed. The finite volume and token comparisons are established afterwards on this fixed run, and only then is a requested extraction order refined.

8.11 Paying the actual compulsory-emission forecast

The compulsory-emission estimate must retain its quadratic monitor exponent until it has been compared with the output forecast. Here is that comparison. Fix the block, the positive local emission parameter μ, and the finite output order. Put $y = \beta_F(x, L)$, $p = M(y)$, with $M(y) = \exp\{(\log\log y)^2\}$. The integrated emission bound obtained from the packet and connected-cluster estimates is

$$2E(x)\exp\{-\mu p^2/4\}, \qquad E(x) \longrightarrow 0. \tag{EF1}$$

This is the bound before weakening its exponent in the dictionary estimate. The requested strong forecast has form

$$C\exp\{-at^2\}, \quad C \geq 1, \qquad t = \frac{y^{-1/2}p}{\beta^{-1/2}}, \qquad \beta \in [y/2, 2y]. \tag{EF2}$$

In particular $0 \leq t \leq 2p$. After μ has been selected at this block and chart width, choose the positive forecast rate $a \leq \mu/32$, and, if necessary, decrease its requested rate to at most $\mu/64$ so that its strict positivity margin is retained. This is a choice of the numerical forecast, made before choosing its entry delay. It gives

$$at^2 \leq 4ap^2 \leq \mu p^2/8 \leq \mu p^2/4. \tag{EF3}$$

For any prescribed $c > 0$, choose the entry delay so that $2E(x) \leq c$. Equations (EF1)–(EF3) then put the actual emission bound below c times (EF2), uniformly for every coupling in the band. The source uses the weaker sufficient condition $a \leq \mu/16$; the chosen $\mu/32$ leaves room for subsequent attachments. These choices are made separately for each of the finitely many legal adjacent stage pairs and combined by their maximum. Thus the forced forecast used by the step, and later by local grafting, is purchased with its actual quadratic decay and on its actual chart. A previously weakened subquadratic estimate is not needed for this payment.

8.12 Uniform reserve at every step, including startup

The numerical purchases also leave a squared reserve for errors that will later be divided by local normalization factors. We derive that reserve, since eventual smallness at each fixed stage alone would not supply a bound at every step of the chosen history. Put

$$M(x) = e^{(\log\log x)^2}, \qquad \kappa(x) = e^{-M(x)^{1/5}}, \qquad x > e. \tag{CAL19}$$

Thus $\kappa(\widehat{\beta}_m)$ is exactly the forced allowance at step m. Fix a positive stage rate q. Eventually $qM(x) \geq 2$. Since $M(y) \geq M(x) \geq 1$ for $y \geq x > e$,

$$qM(y)^2 \geq qM(x)^2 \geq 2M(x) \geq 2M(x)^{1/5}, \qquad e^{-qM(y)^2} \leq \kappa(x)^2. \tag{CAL20}$$

The stage rate in this formula is the positive rate of the actual already chosen post-length estimates. Choose its threshold separately for each stage and include these thresholds among the entry-delay maxima. At step m, use the outgoing stage $J(m+1)$, $x = \widehat{\beta}_m$, and $y = \widehat{\beta}_{m+1} \geq x$. This gives the squared-reserve bound simultaneously for every m, with the same post-length data. The same delay can enforce $\log\widehat{\beta}_m \geq 1$.

Startup has a corresponding reserve without requiring a preceding block. In the startup formulas write $s = \log x$, so their large parameter is $M_s = e^{(\log s)^2} = M(x)$ and their small remainder exponent is $\gamma_s = e^{-M_s/2}$. The startup estimates established above provide its local expansion with

packet-energy charge $\sigma = M_s$, extracted suppression charge $\kappa_0 = M_s$, and remainder charge γ_s, at coupling $x/2$. The connected support budget for that remainder is

$$e^{-M_s}\left[e^{e^{\gamma_s}-1}-1\right] \le e^{-M_s}(e-1) \le e^{-2M_s^{1/5}} = \kappa(x)^2 \tag{CAL21}$$

for sufficiently large s. The first inequality uses $e^{\gamma_s} - 1 \le 1$, and the second follows from $M_s - 2M_s^{1/5} \to \infty$. In particular it pays the full connected budget, including its exponential conversion from the local remainder bound.

For each fixed stage the startup threshold is finite. Include it in the same entry-delay maximum. At level zero the calendar identifies the initial stage with its first entry stage, so this threshold is paid there too; at a successor level it follows from the previous entry payment and the positive coupling increment. Consequently the chosen calendar gives (CAL21) at every possible finest cutoff. For any actual coupling in $[x/2, 2x]$, the startup expansion is still valid: its lower-coupling hypothesis follows from the left endpoint, and its vacuum price is bounded using the right endpoint. The exact startup density identity then turns (CAL21) into the required forced-activity bound. These choices enlarge delays before the final calendar and history are selected, and preserve all earlier purchases. They add no assumption to the theorem.

8.13 Inverse curvature coordinates and the startup fourth coefficient

The fourth-degree correction used in startup is determined by the Wilson plaquette price in curvature coordinates. It is separate from the earlier inequality bounding the difference between chord squared and curvature squared. Write $c(U) \in \mathfrak{g}$ for the projected curvature, and define

$$k(a) = c(\exp a). \tag{SQ1}$$

This real analytic map has $k(0) = 0$ and $Dk(0) = I$. Indeed the exponential has derivative $a \mapsto a$ at zero, its inverse group factor has derivative $-a$, and taking the projected skew part of their difference divided by two returns a. The finite-dimensional analytic inverse theorem gives open neighborhoods of zero and an analytic inverse u with

$$u(k(a)) = a, \qquad k(u(z)) = z, \qquad u(0) = 0, \qquad Du(0) = I.$$

Set $R(z) = \exp(u(z))$. On a group neighborhood of the identity and a color neighborhood of zero,

$$c(R(z)) = z, \qquad R(c(U)) = U. \tag{SQ2}$$

For the second identity use the group logarithm chart constructed earlier: write $U = \exp a$, shrink so that a is in the inverse neighborhood, and substitute $c(U) = k(a)$. Thus no global group logarithm or connectedness assumption is involved.

Define the real analytic scalar function

$$f(z) = \mathrm{price}(R(z)). \tag{SQ3}$$

Since $R(0) = 1$, the price is zero there. Its first derivative is zero because the Wilson price is a squared chord. For $R(tz)$, its matrix derivative at zero is the matrix of z, so the physical color normalization gives

$$f(tz) = \tfrac{1}{2}t^2|z|^2 + o(t^2).$$

The second derivative is symmetric; polarization of this identity gives

$$D^2 f(0)[z, w] = \langle z, w\rangle. \tag{SQ4}$$

Both inversion and conjugation preserve the price, while $c(U^{-1}) = -c(U)$ and $c(hUh^{-1}) = \mathrm{Ad}_h\, c(U)$. Apply (SQ2) to the nearby group elements $R(z)^{-1}$ and $hR(z)h^{-1}$. Uniqueness of the inverse curvature coordinates gives

$$f(-z) = f(z), \qquad f(\mathrm{Ad}_h\, z) = f(z) \tag{SQ5}$$

as germs at zero, for every fixed $h \in G$. Differentiating these germ identities proves that every odd Taylor coefficient vanishes and every Taylor coefficient is adjoint invariant. The positive physical inner product is unchanged by the adjoint action.

Use Taylor coefficients, including the factorial, to set

$$Q_4(z) = \frac{1}{4!} D^4 f(0)[z, z, z, z]. \tag{SQ6}$$

This is the fourth-degree invariant polynomial called Q_4 in (FST1). By (SQ4)–(SQ5), the Taylor polynomial through order four is exactly $\frac{1}{2}|z|^2 + Q_4(z)$. Consequently

$$f(z) - \tfrac{1}{2}|z|^2 - Q_4(z) \quad \text{has all derivatives of orders } 0, \dots, 4 \text{ zero at } 0. \tag{SQ7}$$

For a smooth family of actual group elements $U(v)$ with $U(0) = 1$, (SQ2) identifies its Wilson residual

$$\mathcal{R}(U) = \mathrm{price}(U) - \tfrac{1}{2}|c(U)|^2 - Q_4(c(U)) \tag{SQ8}$$

with the composition of the remainder in (SQ7) and $c(U(v))$, near zero. The chain rule shows that all derivatives below order five vanish: each term contains a derivative below order five of the outer remainder. This proves the exact flat-jet assertion underlying the startup estimate.

We next explain the uniform derivative bound, including arbitrary plaquette links and threshold equality. Let $U = (U_0, U_1, U_2, U_3) \in G^4$ and $P(U) = U_0 U_1 U_2^{-1} U_3^{-1}$. For a list of four-link directions $A_1, \dots, A_q$, let $J_q(A; U)$ be the ordered left derivative of $\mathcal{R}(P(U))$, in the convention (AB3). It is jointly real analytic in the direction variables and the four matrix entries on their actual group locus. This follows by induction: differentiating the analytic expression along the analytic left increments preserves analyticity. It is homogeneous of degree one in every direction; in particular

$$J_q(tA; U) = t^q J_q(A; U). \tag{SQ9}$$

Keep U_1, U_2, U_3 arbitrary and replace the first link by

$$\widetilde{U}_0(z) = \exp(z) U_3 U_2 U_1^{-1}. \tag{SQ10}$$

The resulting plaquette is exactly $\exp z$. At $z = 0$, the residual has the vanishing property (SQ7). Each ordered link derivative lowers this vanishing order by at most one: introduce one extra real variable for its left increment, apply the vanishing composition to the joint variables, and then differentiate in that variable at zero. Induction gives vanishing through order strictly below $\max(5 - q, 0)$ for $J_q(A; \widetilde{U}(z))$.

For directions in a fixed ball $|A| \le r$, the external parameters consisting of these directions and the four group entries lie in a compact finite-dimensional set. On that set times $|z| \le 1$, the required derivatives of the analytic expression are bounded. Taylor's estimate, obtained by integrating successive derivatives along the segment from zero to z, therefore gives one $C_{q,r} > 0$ such that

$$|J_q(A; U)| \le C_{q,r} |z|^{\max(5-q,0)} \quad \text{if } |A| \le r,\ |z| \le 1,\ P(U) = \exp z. \tag{SQ11}$$

This argument bounds actual left derivatives; it is not a differentiation of a healthy-set mask. The three unaltered links need not be small.

The local logarithm estimate gives constants $c_0, b_0 > 0$ with $|\log V| \le c_0 \operatorname{chord}(V)$ for $\operatorname{chord}(V) \le b_0$. Choose $\delta \le \min(b_0, c_0^{-1}, 1)$. If $\operatorname{chord}(P(U)) \le \theta \le \delta$, and $|A| \le r\theta$, apply (SQ11) to A/θ, then (SQ9). It gives

$$|J_q(A;U)| \le C_{q,r} c_0^{\max(5-q,0)} \theta^{q+\max(5-q,0)} \le C'_{q,r} \theta^5. \tag{SQ12}$$

For $q \ge 5$ the last inequality uses $\theta^q \le \theta^5$; for $q \le 5$ its exponent is exactly five. All inequalities allow equality in the plaquette threshold. At a fixed derivative horizon take a maximum over its finitely many q's.

At a plaquette anchored at o, an allowed direction (AB4) restricted to its four links is at most $r\theta$, with r depending only on the fixed power p and the bounded plaquette distances from o. Pullback to these four links therefore converts (SQ12) into the derivative envelope used in (FST8). A plaquette wrapping around the torus still has these same four stored group entries and bounded periodic distances, so the estimate is unchanged. Summing over the fixed finite number of signed plaquette types and multiplying by β gives the anchored bound $C\beta\theta^5$. The constants are chosen before coupling, cutoff and torus period. This completes the startup Taylor and derivative input from *Extra/* WILSON$_S$ *TARTUP/Curvature* and Residual.

8.14 The actual finest-scale Wilson representative

The Wilson startup is an exact identity for the whole density, although its activity dictionary is nonzero. The regular and forced entries cancel the added principal path terms configuration by configuration. For the allowance estimates, the forced suppression needs only $\beta \ge x/2$. The regular term $C'\beta\theta^5$ also needs $\beta \le 2x$. Thus both allowances are asserted together on the shooting band $x/2 \le \beta \le 2x$; the exact density identity itself does not acquire this restriction.

The startup used in finite shooting has nonzero activities; its total density is nevertheless exactly the bare Wilson density. We describe this cancellation. Fix a finest-level clip c with $c_{\text{direct}} \le 1$, threshold $\theta > 0$, and support width $B \ge 2$. Put $b = 1 - c_{\text{direct}} \ge 0$. For a plaquette holonomy u, let

$$\mathcal{R}(u) = \operatorname{price}(u) - \tfrac{1}{2}|\operatorname{curvature}(u)|^2 - Q_4(\operatorname{curvature}(u)), \tag{FST1}$$

where Q_4 is the fixed fourth-degree invariant curvature correction. The finest path has row $\sqrt{b}$ times the elementary plaquette curvature, zero loop and cubic tensors, and quartic tensor bQ_4 on the same plaquette. Thus its quadratic-plus-quartic part is the corresponding local Taylor polynomial. It uses the entire plaquette footprint, including a wrapping plaquette.

For each plaquette P, write u_P for its holonomy, $r_P = \operatorname{chord}(u_P)$, and $k_P = b|\operatorname{curvature}(u_P)|^2/2$. Define its unpaid nonnegative part by

$$d_P(U) = \begin{cases} \operatorname{price}(u_P) - c_{\text{linear}}\theta r_P, & r_P > \theta, \\ [k_P \text{ if the footprint is healthy, else } b\operatorname{price}(u_P)] & \\ \quad - [k_P \text{ if its anchor is kept, else } 0], & r_P \le \theta. \end{cases} \tag{FST2}$$

The kept predicate is the exponentially relaxed plaquette test of the clip. On a kept anchor its own plaquette satisfies $r_P \le \theta$. The identities $\operatorname{price}(u) = \operatorname{chord}(u)^2/2$, $|\operatorname{curvature}(u)| \le \operatorname{chord}(u)$, and $c_{\text{linear}} \le 1/4$ show $0 \le d_P \le 2$. If $r_P > \theta$, they give the stronger $d_P \ge \theta^2/4$.

Let $\mathcal{A}_{\text{top}}(U)$ denote the actual principal path action, including its clipped direct part and its masked quartic term. Checking separately the healthy case, the unhealthy case with $r_P > \theta$, and the unhealthy case with $r_P \le \theta$, gives the pointwise identity

$$\operatorname{price}(u_P) = \operatorname{baseTerm}_P(U) + \operatorname{mask}_{\theta,P}[bQ_4(\operatorname{curvature}(u_P))] + b\operatorname{mask}_{\theta,P}[\mathcal{R}(u_P)] + d_P(U). \tag{FST3}$$

The first two terms here are precisely the plaquette terms in $\mathcal{A}_{\mathrm{top}}$. In the healthy case substitute (FST1) and cancel. In the large-plaquette case both masks and the kept quadratic term vanish, and the clipped direct term is $c_{\mathrm{linear}}\theta r_P$, canceled by (FST2). In the remaining case both masks vanish and the two alternatives for the kept quadratic term cancel those in (FST2). These are exhaustive cases, including threshold equality.

Average (FST3) over all signed coordinate frames with weight $\omega = |\,\mathrm{Signed}\,|^{-1}$, and sum over plaquettes. Index these oriented terms by i, with footprint S_i, anchor x_i, residual $\mathcal{R}_i$ and unpaid part d_i. Signed-frame invariance of the Wilson energy gives

$$\beta S(U) = \beta\mathcal{A}_{\mathrm{top}}(U) + \beta\omega b \sum_i \mathrm{mask}_{\theta,S_i}\,\mathcal{R}_i(U) + \beta\omega \sum_i d_i(U). \tag{FST4}$$

Take the middle sum as the regular activity action.

To represent the final sum by forced activities, partition the actual flag set into its connected components for the relation that two flags' radius-$2B$ neighborhoods overlap. If D is a nonempty component, let $R(D)$ be the union of those neighborhoods and set

$$q_D(U) = \omega \sum_{i:\, x_i \in \mathrm{halo}_B(D)} d_i(U). \tag{FST5}$$

The extended regions of distinct components are disjoint. Every nonzero d_i lies within distance B of some flag, and hence belongs to exactly one component's sum. The assertion about distance follows from the relaxed test: outside that distance the clip condition $\theta\exp((B-2)/c_{\mathrm{relaxation}}) \geq 2$ and the universal chord bound 2 force all its tests to be kept and healthy. Therefore $\sum_D q_D = \omega\sum_i d_i$.

Assign the forced entry $e^{-\beta q_D(U)}$ to the support-inventory pair $(R(D), D)$, eligible precisely when $R(D) \cap \mathrm{flags}_\theta(U) = D$, and assign zero to all other active entries. In the gas sum, a nonzero admissible collection must have disjoint supports and inventories partitioning the flags. Its inventories are consequently exactly the connected components just described. Conversely that collection is admissible. The gas sum thus has one nonzero term and equals

$$\mathrm{gas}(U) = \exp\left(-\beta\omega \sum_i d_i(U)\right). \tag{FST6}$$

Combining (FST4) with (FST6) proves the exact startup density $e^{-\beta S(U)}$, with scalar normalization zero.

These explicit entries also satisfy the needed bounds. The normalized derivatives of $\mathcal{R}$ on a healthy plaquette, through each fixed required order, are at most $C\theta^5$. This follows by Taylor expansion of the smooth four-link product and (FST1): its curvature correction cancels the jets below degree five; the permitted derivative directions have size $O(\theta)$, so differentiating does not lose the total factor θ^5. The constants are chosen for a fixed order before the coupling or period. There are only the fixed number of signed plaquette types at an anchor, and their cover cost is bounded. Consequently the anchored regular bound is at most $C'\beta\theta^5$.

For a forced component D, every flag is incident to a plaquette above threshold and each plaquette touches at most four sites. Using $d_P \geq \theta^2/4$ and keeping one signed frame gives

$$q_D(U) \geq |D|\,\omega\theta^2/16. \tag{FST7}$$

This exponential suppression pays the connected-component support count and its support weights. To see its size in the actual schedule, set $x = \widehat{\beta}_N$, $z = \log x$, $M(z) = e^{(\log z)^2}$, and $\theta = e^{-z/2}M(z)$. For the forced estimate, the lower bound $\beta \geq x/2$ is sufficient. On the actual shooting band

$x/2 \le \beta \le 2x$, (FST7) has strength proportional to $M(z)^2|D|$. The neighborhood count $(4\lceil z^2\rceil+1)^4$ and each fixed stage support cost are eventually at most $M(z)/4$. The component expansion may therefore use local allowance $e^{-M(z)/2}$; its summed connected-activity bound is at most $e^{-M(z)+1}$, which is eventually at most $e^{-M(z)^{1/5}}$, the required forced allowance. The regular bound satisfies

$$C'\beta\theta^5 \le 2C'e^{-3z/2}e^{5(\log z)^2} \le e^{-11z/10}, \tag{FST8}$$

the required regular allowance. The same finite thresholds also ensure the clip condition preceding (FST6), since $B = \lceil z^2\rceil$, and $\theta \to 0$. These thresholds are uniform in N, p once the stage is fixed, and have already been included in its entry delay.

The dictionaries of supports and flag predicates are independent of β. Regular entries are linear in β, and forced entries are exponentials of bounded finite-volume functions times fixed eligibility indicators. Their weighted sums obey the preceding majorants. Thus the startup family is continuous in precisely the activity budgets used for finite shooting. Countably many unused addresses can be appended with value zero, without changing density or bounds; they furnish the marker slots required later. This constructs (FCA2) with its actual representatives.

8.15 Finite coupling arcs and calibrated finite histories

Finite shooting selects a coupling on a continuous finite arc with Wilson startup. The estimates control that arc in the specified band, and the common calibration condition fixes the retained endpoint. Retirement removes only terms whose exact density identity and purchased error have already been recorded. The limiting history is unique within this fixed recipe and calendar; uniqueness across arbitrary choices of recipe is not part of the construction.

We construct the finite histories before using their compactness or uniqueness. Fix the common calendar, a terminal depth N, and a period p. Put

$$\tau_i = \log(a_i^{-1}), \quad s = \mathrm{slopeConstant}(G) > 0, \quad \widehat{\beta}_i = \bar{\beta} + s\tau_i, \quad \mathcal{B}_i = [\widehat{\beta}_i/2,\, 2\widehat{\beta}_i]. \tag{FCA1}$$

Here $\tau_0 = 0$, $\bar{\beta} > 0$, and $\tau_{i+1} - \tau_i = \log L_i > 0$. The path tensors at fixed N, i are fixed independently of p and of the bare-coupling shooting parameter. A represented state consists of those tensors, its prescribed clip and budget, its actual coupling β, scalar normalization v, and regular and forced activities. Admissibility means the local regularity, symmetry and absolute activity bounds proved in the nonlinear step, together with the retained startup and transition conditions. It includes $\beta \in \mathcal{B}_i$.

Continuity of a family of such states has a specific meaning. Its β and v are continuous; its regular and forced address dictionaries are constant; and differences of its activity values tend to zero in the absolute weighted derivative budgets. For regular terms derivatives are tested only on the common healthy domain, through the prescribed order and in the prescribed directions. For forced terms the bound is on every configuration. On a fixed period a global address sum is legitimate for this continuity assertion. It is not substituted for the period-independent anchored bounds in admissibility.

The nonlinear construction provides a forward family at every finite step on this one calendar. At the finest level its states $F_{N,p}(b)$, $b \in \mathcal{B}_N$, have

$$\beta(F_{N,p}(b)) = b, \qquad v(F_{N,p}(b)) = 0, \qquad F_{N,p}(b).\,\mathrm{eval} = \mathrm{bare}_b\,. \tag{FCA2}$$

They have the fixed startup representatives and insertion slots. For a continuous family $Q(b)$ of admissible level-$i+1$ states on a compact parameter interval, the same primitive integration and

reset construction produces continuous coarse states $T(b)$, with the same retained choices, such that

$$T(b).\,\mathrm{eval} = \mathcal{P}_i^{\beta(Q(b))}(Q(b).\,\mathrm{eval}), \qquad |\beta(Q(b)) - \beta(T(b)) - s\log L_i| \le d_{i+1}. \tag{FCA3}$$

The number d_{i+1} is the fixed drift allowance of this stage. The output satisfies all level-i local bounds before imposing its coupling band. This distinction permits the shooting argument to detect crossings of that band.

These forward families are obtained from the actual primitive formulas, rather than a choice of unrelated admissible outputs. At a step, keep one primitive construction, its ordinal and tree choice, and its fixed support dictionaries. Its output is the moment state produced by that primitive. The local integral comparison and the absolute activity estimates above prove continuity of the output in the stated budgets. The reset identity gives the density equality in (FCA3); its coupling bound gives the second inequality. For a continuous input curve these same formulas therefore define a continuous output curve. For two occurrences of the identical input state they give the identical output state: equality is equality of the represented moment states, including the activities. Choosing this rule once at each N, i, p, and the startup rule (FCA2) once at N, p, gives a single deterministic forward recipe. Its underlying path, clip, normalizers, source bounds and calendar remain those of the nonlinear construction. Restricting a parameter interval preserves every one of these properties.

To make the normalization in this construction explicit, write $H_b(U) = \mathcal{P}_i^{\beta(Q(b))}(Q(b).\,\mathrm{eval})(U)$. The positive push kernel and the retained positive vacuum give $H_b(1) > 0$. If $\alpha_{N,i}$ is the already committed path correction, the output coupling and scalar are

$$\beta(T(b)) = \beta(Q(b)) + \alpha_{N,i} + \Delta(b), \qquad v(T(b)) = \frac{\log H_b(1)}{|\,\mathrm{Site}_i\,|}. \tag{FCA9}$$

The scalar reset produces $|\Delta(b)| \le r_i$ and $\theta_i^2|\Delta(b)| \le e_i$, where r_i is its fixed correction radius and e_i its regular allowance. In the exact normalized splitting, the output trunk density times its attached-log exponential is $H_b(U)/H_b(1)$. Installing (FCA9) multiplies that density by $\exp(|\,\mathrm{Site}_i\,|v) = H_b(1)$, proving (FCA3) at every configuration, including flag boundaries. The site count, rather than the bond count, is essential in this cancellation.

The correction in this construction is the source Hessian projection:

$$\Delta(b) = \mathrm{Averaging.\,delta}(X.\,\mathrm{read}(a(b))). \tag{FCA10}$$

Here X is the actual fixed source listing and $a(b)$ its parameterized source. The affine-curvature Hessian projection, rebates, faults and low- and high-path loan identities in the coupling-reset exponent section prove the exact all-configuration representation with this value of Δ. In particular the correction is not defined by solving a two-probe contrast equation. The projection's diagonal bound gives $\theta_i^2|\Delta(b)| \le H(b)$, where $H(b)$ is the paid regular source bound. The schedule pays $H(b) \le f_i/4$ and the reset fees, giving the stated correction radius and the unit bound used in (DR7). The regular and forced bounds after the projection and the two dictionary fusions are paid independently by those same fees.

To see the normalization without an extra characterization of the correction, let $\tau(b)$ be the primitive scalar before reset and let $\ell_j(b,1)$ be its attached logarithms at the identity. The exact exponent calculation gives the installed scalar

$$v(b) = \tau(b) - |\,\mathrm{Site}_i\,|^{-1}\sum_j \ell_j(b,1). \tag{FCA11}$$

The resulting state has regular entries zero at the identity, no identity flags, and kinetic and loop actions zero there. Its density at the identity is therefore $e^{|\operatorname{Site}_i|v(b)}$. The already proved all-configuration density identity makes this equal $H_b(1)$, proving the logarithmic formula in (FCA9). Thus (FCA9) and (FCA11) are the same scalar normalization. The interval estimate for this same projection value is obtained by the probe argument below. We must still show that this projected correction lies in the purchased correction interval. The source uses a nonflat probe for this purpose, on the very same trial family. Its activity terms are independent of the correction dial z: they have already been computed with the Hessian projection value Δ. Only the scalar coupling $\beta + \alpha + z$ varies in that trial. Thus there is no second reset algorithm.

Here is the probe construction, uniform in the even torus side. Fix the clip's positive direct coefficient a and its common cubic/quartic path bound $C \geq 0$. A nonzero Lie-algebra direction gives a group curve through the identity with nonzero curvature derivative. Its chord is therefore positive somewhere; continuity and the intermediate value theorem give $\varepsilon > 0$ such that every chord in $[0, \varepsilon]$ is attained. Put

$$t = \min\{\varepsilon, 1/4, a/[512(1+C)]\}, \qquad \eta = at^2/4 > 0. \tag{RSP1}$$

Choose $u \in G$ with chord $t\theta$. Set all links to the identity except direction-zero links whose direction-one coordinate is even; set those links to u. This is periodic on every even side. Every $(0,1)$ plaquette has holonomy u or u^{-1}, and every plaquette has chord at most $4t\theta \leq \theta$. Thus the field V has no flags. Its direct plaquette charge in every signed-frame sum is at least $na(t\theta)^2/2$, where n is the number of sites. The cubic and quartic terms have absolute sum per site at most $C(4t\theta)^3 + C(4t\theta)^4 \leq a(t\theta)^2/4$, by (RSP1) and $\theta \leq 1$. Averaging over the signed frames leaves

$$K(V) \geq n\eta\theta^2 > 0, \tag{RSP2}$$

where K is the actual kinetic action. No smallness of the loop tensor enters this lower bound.

Write $H(U)$ for the actual transferred density at a fixed input parameter, and set $B = \beta + \alpha$. On any flag-free configuration its early representation is

$$H(U) = \exp\{n\tau - [\beta\ \mathrm{main}(U) + \alpha\ \mathrm{second}(U) + \mathrm{loop}(U)] - E(U)\}. \tag{RSP3}$$

The source's anchored regular budget is at most $F/4$, with $F = \mathrm{funds}_i$. Summing its anchors bounds $|E(U)| \leq nF/4$ at every flag-free configuration, including 1 and V. Formula (RSP3) proves positivity at both probes and gives $|\log H(1)/n - \tau| \leq F/4$. Subtracting its two logarithms gives contrast error at most $nF/2$ relative to $\beta\,\mathrm{main} + \alpha\,\mathrm{second} + \mathrm{loop}$.

On flag-free fields the main and loop actions equal their unmasked path expressions. The higher-order kinetic correction obeys $|K(V) - \mathrm{second}(V)| \leq 2nC\theta^3$. The stage's cubic payment is $|\alpha|, 2C\theta^3 \leq F/2$. Therefore, with $L_0(U)$ the zero-coupling loop background,

$$|BK(V) + L_0(V) - [\log H(1) - \log H(V)]| \leq nF. \tag{RSP4}$$

This is the actual contrast estimate, with its α and loop terms retained.

Let e be the final regular allowance, $r \in (0,1]$ the correction radius. The reset price is chosen before history selection so that

$$2(F+e) \leq \eta\theta^2 r. \tag{RSP5}$$

This is paid by the coupling-entry delay: $F = ce$, $e = g^{11/5}$, $\theta = gM$, and $g \leq 1$, whereas the monitor $M \to \infty$; it suffices to impose $2(c+1) \leq \eta r M^2$. This same condition is imposed for both permitted adjacent-stage choices.

For the fixed trial activities, let A_V be their regular action at V. Define

$$R(z) = (B + z)K(V) + L_0(V) + A_V - [\log H(1) - \log H(V)]. \tag{RSP6}$$

Their budget gives $|A_V| \le ne$. Equations (RSP2), (RSP4), and (RSP5) imply $|R(0)| \le n(F + e) \le K(V)r/2$. Since $R(z) - R(z') = (z - z')K(V)$, this affine function has a unique root in $[-r, r]$. Its root varies continuously with the input: the contrast and coupling are continuous, the fixed-trial activity series are continuous by their common deviation bounds, and $K(V) > 0$ is fixed.

The projected state already satisfies the full density identity, so evaluating it at $1, V$ gives $R(\Delta) = 0$. More explicitly the exact fixed-trial identity is $R(z) = (z - \Delta)K(V)$. Thus the unique interval root is exactly the previously defined Hessian projection Δ. This proves its radius bound and transfers the interval service's continuous curve to the actual represented state. The independent projection estimate $\theta^2|\Delta| \le H$ retains the sharper paid bound. The all-configuration density equality is still supplied by the exponent and gas identities, not inferred from equality at two probes. In the source the corresponding shape comparison is also explicit: for any configuration U, the trial shape is its value at Δ times $e^{(\Delta - z)K(U)}$; the inequality $|e^x - 1| \le e^{|x|}|x|$ bounds its discrepancy by a finite configuration-dependent constant times $|R(z)|$, using $K(V) > 0$. Hence the interval and projection formulations are fully compatible on this same trial family. For a compact input curve, the connected-expansion tails can be made small at one common finite retirement index. Their absolute tail bounds and the finitely many local bounds permit that choice throughout the curve. In the deterministic recipe the index is chosen as a function of the incoming address layout. A curve's layout is constant, so that index is constant on the curve. The identities for the primitive expressions show that changing the incidental local guide or the point at which the curve was presented does not change the resulting state: all such expressions have the same tensors, canonical matching coordinates, source list and scalar reset. This supplies the continuous forward rule used above without silently choosing new dictionaries at different parameter values.

The source-level continuity and admissibility statement can now be written without suppressing its series hypotheses. Fix a closed parameter interval I and an incoming curve $Q(t)$ in its admissible band. Use one point $t_0 \in I$ as reference. The incoming path equals the prescribed path at every point, and the regular and forced layouts equal those at t_0. The matching maps therefore use identity label correspondence on these dictionaries. Local physical guides may be selected separately at each point; the matched primitive expression is independent of that selection. Transport all entries to the reference layout. This fixes the unretired source dictionary, its support windows and the finite debris predicate used for retirement.

For its source $a(t)$, the actual continuity assertions are the following. Each regular entry, with all permitted derivatives on its fixed healthy support, and each forced entry on all configurations, converges as $t \to t_0$ within I. Their absolute support-weighted envelopes have uniformly small tails near t_0. These two facts imply the deviation budget tends to zero: given $\varepsilon > 0$, choose a finite set of addresses leaving total tail less than $\varepsilon/2$; then choose a neighborhood on which the finitely many remaining differences cost less than $\varepsilon/2$. This uses no continuity of flag indicators as functions of the gauge field: their support and threshold predicates are held fixed in the compared evaluation.

On a compact input interval one retirement index can be used throughout. The uniform tail estimate for the unretired source and compactness give an eventual regular tail bound $g > 0$, while the local finite-volume comparison gives an eventual bound for all the retained local beacon entries. Intersect these eventual conditions, choose one integer above their thresholds, and retain it at every parameter value. The resulting source obeys the explicit regular estimate

$$\mathrm{Budget}_{P_i}(a(t)) \le (1/16 + 2/16)E_i + \mathrm{turnFee}(P_i)g + g, \tag{FCC1}$$

where $E_i = \mathrm{forecastE}_i$. Each beacon on its fixed neighbor support has the derivative bound gleam_i. The source graft combines the retired regular entries, beacon entries and the fixed local profile. Its forced ingredient estimate is paid by its original forced allowance: the primitive forced part costs at most $\mathrm{forecastK}_i(\beta)/8$, while the added local profile has size at most the prescribed $\beta_F^{-1/8}$ reserve after its cover weight. The local gluing estimate sums these contributions into $\mathrm{forecastK}_i(\beta)$. The turn fee times the local beacon bound is included in the regular price. The stage delays make the right side of (FCC1), with that additional beacon price, at most E_i. This proves the regular and forced local estimates for the grafted source simultaneously at every parameter value.

The uniformity over all presentations of one layout requires a separate argument. A layout is the regular and forced address lists, including supports, without coefficient functions. If no admissible state has a layout, its assertion is empty. Otherwise choose one such state U. Form the canonical curve U_β on the entire closed incoming coupling band by retaining its path, cutoff and address lists, setting every regular and forced coefficient function and the vacuum scalar to zero, and assigning coupling β. Keep the address symmetry permutations. Zero functions satisfy every derivative and hit bound, so this curve remains admissible and is continuous.

Choose local guides along this canonical curve and match their expressions to its reference layout. At a fixed coupling the weighted connected series has a summable majorant. For a target tail bound $g > 0$, take a finite set of terms outside which its sum is below $g/2$. Each of these finitely many terms eventually leaves the debris set: a term is either kept or eventually appears in the finite retirement roster. Hence one depth removes them all. Continuity in the deviation budget bounds the change of the entire weighted series by $g/2$ in a neighborhood of that coupling. The retired part is a subfamily of the same majorant, so the same neighborhood works for every larger retirement depth. Finitely many such neighborhoods cover the closed coupling band. Their maximum depth supplies a common tail bound g. Intersect it with the eventual finite beacon bounds, taking $g = \mathrm{gold} > 0$.

Now let $Q(t)$ be any other admissible curve with this layout and any allowed guide choices. At a point t, compare its guide with the canonical guide at $\beta = Q(t).\beta$. The paths, lists, cutoffs, budgets and couplings agree. The matching map is the identity on the lists, and its induced relabeling preserves the finite retirement roster and debris predicate. For a debris address the primitive term does not read the incoming coefficient functions: it is an omitted foreground cluster formed from the common geometric and Gaussian factors, or an unretired scalar path term. Matching preserves its connected coefficient, integrand and scalar coupling. Thus the two debris functions and beacon contractions agree:

$$1_{\mathrm{debris}(d,i)} W_i(Q(t)) = 1_{\mathrm{debris}(d,i)} W_i(U_\beta).$$

This is equality of local functions, so it preserves every derivative envelope tested in the budget. Transfer the canonical tail and beacon bounds through it. Adding the principal unretired contributions gives (FCC1) for the arbitrary presentation. The threshold therefore depends on the layout, stage, depth, period and fixed prices, but not on the curve, its coefficients or its guide choices.

Apply this construction both with standard allowances $1/16, 1/16$ and with the allowances needed for the contraction. Intersect the two eventual sets with $d \geq N$, and choose one integer $d(\mathcal{L})$ for each layout. Along every admissible curve $\mathcal{L}$ is constant, hence so is $d(\mathcal{L})$. The matching identities identify the deterministic rule with the fixed-index continuous formula on that curve. No continuity of this integer-valued selection as the layout changes is needed. For the deterministic operation, the retirement index is selected once from the incoming layout, at least the terminal depth N, to satisfy these bounds for every compatible presentation of that layout. The universal retirement estimate supplies such an index. Because a curve has one layout, this selected index is constant on the curve. The primitive changes of local guide leave the matched source unchanged,

including its finite debris predicate and beacon contraction. Thus the deterministic moment state agrees with the continuous curve construction just described. Restricting the interval in the shooting argument changes neither index nor dictionary.

The Hessian projection is continuous in the resulting deviation budget. For two source values a, a', its bound is

$$\theta_i^2 |\Delta(a) - \Delta(a')| \le \mathrm{deviation}_{P_i}(a, a'). \tag{FCC2}$$

The threshold is a fixed positive number at this step. The rebate, fault, freshness and loan difference identities are linear in the source difference and in $\Delta(a) - \Delta(a')$, as proved in the reset chapter. Their summable quote envelopes therefore pass continuity and uniform tails to the reset entries. The finite-support dictionary fusions preserve these properties by their absolutely convergent product bounds. In particular both the final coefficient families and their actual fixed address layouts form a curve in the stated sense.

Finally the local primitive identity identifies the transferred density with its early-value expression on every configuration. The reset's exact exponent and gas identities then give its installed state with that same transferred density; its regular and forced closure estimates give all the coarse admissibility conditions. The scalar vacuum term is continuous by the same source-series estimate. Equivalently, evaluate the positive identity at the flat configuration to obtain its logarithmic normalization (FCA9). Together with continuity of the incoming coupling and (FCC2), this proves continuity of the output coupling as well. These are precisely the continuity, density, drift and band-free output conditions used in (FCA3); the subsequent interval restriction is what imposes the coarse coupling band. The calendar makes the following drift condition hold at every step:

$$d_{i+1} \le \frac{s}{4} \log L_i. \tag{FCA4}$$

Define narrower target intervals

$$\ell_i = \frac{3}{4}\bar{\beta} + \frac{5}{8} s\tau_i, \qquad u_i = \frac{5}{4}\bar{\beta} + \frac{3}{2} s\tau_i. \tag{FCA5}$$

Since $\tau_i \ge 0$, they obey $\ell_i \le u_i$ and $[\ell_i, u_i] \subset \mathcal{B}_i$. If the fine coupling in (FCA3) equals ℓ_{i+1}, then

$$\beta(T) \le \ell_{i+1} - s \log L_i + d_{i+1} \le \ell_i - \tfrac{1}{8} s \log L_i < \ell_i.$$

If it equals u_{i+1}, then

$$\beta(T) \ge u_{i+1} - s \log L_i - d_{i+1} \ge u_i + \tfrac{1}{4} s \log L_i > u_i. \tag{FCA6}$$

These strict crossings require no monotonicity or derivative bound on the coarse coupling as a function of the bare coupling.

We use the following elementary interval restriction. Suppose f is continuous on $[a, b]$, $f(a) \le \ell \le u \le f(b)$. Let b' be the first point at which $f = u$, and let a' be the last point in $[a, b']$ at which $f = \ell$. These points exist by the intermediate value theorem and compactness of the two level sets. Then $a' \le b'$ and $f([a', b']) \subset [\ell, u]$. A value below ℓ would force a later crossing of ℓ before b'; a value above u would force an earlier crossing of u. Both contradict their definitions. This argument also covers $\ell = u$.

Starting at N, restrict the bare parameter to $[\ell_N, u_N]$ in (FCA2). Descend by finite induction. If an interval already carries trajectories through levels $i+1, \dots, N$ and its endpoint level-$i+1$ couplings are ℓ_{i+1}, u_{i+1}, apply (FCA3) to their level-$i+1$ curve. Equations (FCA6) and the preceding interval restriction give a closed subinterval on which its level-i coupling stays in $[\ell_i, u_i]$

and takes these exact endpoint values. Retain the already constructed higher states and insert this new coarse state. Every higher density identity, startup condition and continuity estimate survives the restriction.

At level zero we have a nonempty closed bare-parameter interval $[a_{N,p}, b_{N,p}]$ and complete forward trajectories

$$b \longmapsto (R^{b}_{N,p,i})_{0\le i\le N}. \tag{FCA7}$$

Their finest coupling is exactly b; each level is continuous in the stated sense; all states are admissible; and every adjacent pair satisfies (FCA3). Their bottom couplings take the endpoint values $3\bar{\beta}/4$ and $5\bar{\beta}/4$. Thus every bottom value in this interval is attained by a trajectory of the same recipe. In particular choose $b^{*}_{N,p}$ with bottom coupling $\bar{\beta}$, and set $R_{N,p,i} = R^{b^{*}_{N,p}}_{N,p,i}$. This constructs the calibrated finite histories for every N, p.

The same retained arc supplies the auxiliary trajectories at nearby prescribed bottom couplings used in the confluence argument. It supplies whole trajectories, not just scalar sequences satisfying the drift inequality. No uniqueness of an inverse shooting map has been assumed. Determinism asserts only that equal finest states give identical forward trajectories; the later contraction argument establishes the required uniqueness of their limiting calibrated data.

There is also a uniform interior margin. If the bottom value of any successful trajectory lies in $[3\bar{\beta}/4, 5\bar{\beta}/4]$, then

$$\frac{3}{4}\widehat{\beta}_i \le \beta(R_i) \le \frac{5}{4}\widehat{\beta}_i \qquad (0 \le i \le N). \tag{FCA8}$$

At $i = 0$ this is the given bottom interval. If it holds at i, solve (FCA3) for the fine coupling and use (FCA4): its increment lies between $(3s/4)\log L_i$ and $(5s/4)\log L_i$. Adding these bounds proves (FCA8) at $i+1$. This margin justifies the nearby-trajectory comparisons while retaining the original full bands.

Finally, assembling these states into the history preserves the common path at fixed N, i, its period independence, the prescribed budget and clip, the finest bare Wilson law, and all exact push identities. Induction in the push integral identity gives the same finest Wilson partition function at every coarser level. It is this calibrated history, rather than a separately selected sequence of couplings, that enters the limits below.

8.16 Exact density identity after retirement

The retirement used above is an exact decomposition before it is an estimate. Keep one anchored listing for the whole input curve. For its full regular source write $A_{\mathrm{full}}(V) = \sum_j a_j(V)$. Partition its indices by the chosen retirement predicate and write $a_j = a_j^{\mathrm{keep}} + a_j^{\mathrm{tail}}$, including their original healthy masks. The same predicate is used throughout the compatible curve. The local absolute envelopes make both subseries summable, hence

$$A_{\mathrm{full}} = A_{\mathrm{keep}} + A_{\mathrm{tail}}. \tag{RET1}$$

The forced gas is retained in the kept source exactly; it is not split into two multiplicative copies. If the density of the full source is

$$e^{v|\Lambda| - H(V) - A_{\mathrm{full}}(V)}\, \mathcal{G}(V),$$

then the identical integrated density after retirement is

$$e^{-A_{\mathrm{tail}}(V)} e^{v|\Lambda| - H(V) - A_{\mathrm{keep}}(V)}\, \mathcal{G}(V). \tag{RET2}$$

This follows by substituting (RET1) in the exponential; in particular the tail has a minus sign. Summability is needed here to identify the sums in (RET1), rather than merely cancelling finitely many formal letters. Replacing the old listing by the anchored common listing preserves each masked occurrence, so it preserves this equality on all configurations.

The full-source density identity follows from the connected expansion already established. Its empty-flag pool is strictly positive and equals the exponential of its absolutely convergent connected logarithm. The scalar normalization contains the positive Gaussian and coordinate prefactors, the incoming scalar volume term, and the reference-forest normalization. Taking their logarithm per coarse site defines v. The cubic, quartic, loop and local correction scores have the exact totals displayed above. Their signs cancel those terms in $H + A_{\mathrm{full}}$, leaving precisely the logarithm of the positive pool. Multiplying by the forced pool then recovers the original conditional integral, including its prefactors. Thus (RET2) holds before either the kept source or the tail has its sharper output budget.

The kept-source estimate and the tail estimate are purchased separately on this same listing. The former gives the desired small regular price together with the original forced forecast; the latter is the remainder used by the finite reference beacons. Reattaching that remainder restores (RET2). Consequently neither continuity nor projection replaces the exact density by a merely close one.

8.17 Finite reference beacons and weighted retirement continuity

The finite beacons used in (FCC1) have an explicit five-volume construction. For integer half-width w, use reference torus sides $2(w+j)$, $j = 0, \dots, 4$, and coefficients

$$\rho_j = \frac{(1, -4, 6, -4, 1)_j}{2^4\, 4!}. \tag{RB1}$$

Fold each coordinate of a reference torus by reflection into its interval of length $w + j$. To pull a configuration to this folded torus, a positively traversed original edge is read directly, a negatively traversed edge is inverted at its lower endpoint, and a collapsed edge is the identity. Translate this folded view to a chosen coarse anchor o. Denote the resulting group configuration by $S_{j,o}V$. This finite word map defines the views even at folded faces.

Fix the local read radius r, with $2r \le w$. A position in a folded coordinate has a type: within r of the left face, within r of the right face, or in the interior. A boundary type records the face and its distance $0, \dots, r-1$. Its number of preimages under folding is independent of j; an interior type has $2(w - 2r + 1) + 2j$ preimages. In four coordinates a full type therefore has a multiplicity which is a product of four constant or affine functions of j. The fourth difference (RB1) kills every such product having a constant factor, and gives one on the product of four interior factors: its fourth difference is $4!\, 2^4$. Hence for any function f of the four-coordinate type,

$$\sum_{j=0}^{4} \rho_j \sum_{x \in \mathrm{torus}_{2(w+j)}} f(\mathrm{type}_j(x)) = f(\text{all-interior type}). \tag{RB2}$$

This is a finite counting identity, including faces and edges of every codimension.

Let $f_j(x, V)$ be the compatible anchored local reference densities. Center them by $f_j^0(x, V) = f_j(x, V) - f_j(x, 1)$. Compatibility means that on healthy configurations two such readings agree whenever their radius-r cover inputs agree, and they obey the exact signed and translation laws. Define the beacon at o by

$$b_o(V) = \sum_{j=0}^{4} \rho_j \sum_{x} f_j^0(x, S_{j,o}V). \tag{RB3}$$

Choose its support radius $z \ge w + 4$. Introduce also the shifted expression which uses anchor $o - \mathrm{stage}_j(x)$ for the x-summand and masks that summand on its radius-z neighborhood. Summing this expression over o gives exactly the sum of the masked beacons, by the finite translation permutation of the torus. On a healthy radius-$2z$ neighborhood all these masks are one. Their folded local readings depend only on their four-coordinate type; compatibility and the signed laws identify equal types. Equation (RB2) reduces the shifted expression to the original centered local density at o. The all-interior type reads its unchanged cover neighborhood. Thus the difference between the assigned retained density, its identity value and this shifted expression vanishes on that healthy neighborhood. This is the local remainder used in the forced part of the graft. Its global identity follows by summing over o; there is no discarded face correction.

The derivative allowance of (RB3) is equally concrete. If the centered full reference sum on side $2(w+j)$ has left-derivative bound q_j on the small reference chart through order k, and the ambient budget has direction power p, then

$$\|b_o\|_{\mathrm{budget}} \le \sum_{j=0}^{4} |\rho_j|\, 2^k (1+z)^{pk}\, q_j. \tag{RB4}$$

On a healthy neighborhood the folded reference configuration can be gauge transformed into that small chart, by the local axial chart construction; its displacement leaves the fixed margin required by the reference estimate. Gauge invariance transfers that estimate back. Each pulled-back direction reads one original edge or its reversed adjoint transform. Its norm is bounded by $(1+z)^p\theta$; reversal preserves the color norm. Repeated left derivatives of the finite word substitution produce at most the 2^k factor in (RB4). Finite summation over the five centered reference sums proves the displayed allowance. These are the actual finite offers used for the beacon, with their common threshold and derivative order.

The q_j's in (RB4) are supplied by the finite reference construction. Put $k = B.\mathrm{order}$, $s = \mathrm{turnFee}(B)$, $h = \mathrm{gold}(G, L, x) > 0$, and $V_j = |\mathrm{Site}(2(w+j))|$. Its actual allowance is

$$q_j = k!\, h\, 3^k + 2V_j\{s(h+h)+h\}. \tag{RB4a}$$

The centered chart expression has bounds $\|\mathrm{D}^\ell f\| \le hp^{-\ell}$, $\ell \le k$, at every admitted point of the chart of radius $p = \mathrm{eaves}(G, L, x)$, with margin $2\,\mathrm{furl}(G, L, x)$. The logarithm of each varying link recovers its local coordinate. If J bounds that logarithm and C is the size of a permitted left variation, the scheduled inequalities are

$$(1 + \mathrm{royalty}(G))C \le 1, \qquad k!J(1 + \mathrm{royalty}(G))C \le p.$$

The chain rule bounds the coordinate map through order k by p, with slope one. Composition with the chart expression costs at most $k!h(1+(1+1))^k = k!h3^k$.

On the strict healthy chart the exact decomposition expresses this chart function as a constant minus the retained full reference sum minus the summable activity source. The source budget bounds the latter by $V_j\{s(h+h)+h\}$, through the same left-derivative order. Centering at the identity costs at most twice this number; derivatives of positive order kill the centering constant. The triangle inequality proves (RB4a). To include a healthy endpoint on the boundary of the strict chart, approach it along tP, $t \uparrow 1$. This path keeps the prescribed chart margin and is eventually strictly healthy. Regularity of the retained reference function gives continuity of each left derivative at the endpoint, and hence the same bound there. No threshold indicator is differentiated.

The reference estimates hold eventually in their approximation depth uniformly on the compact coupling band. The five reference volumes therefore have a common depth. The actual result after

folding is

$$\operatorname{gleam} = \sum_{j=0}^{4} |\rho_j|\, 2^k (1+z)^{B.\mathrm{power}\, k}\, q_j. \tag{RB4b}$$

The radii and reference cardinalities have rate zero, and h has rate $-23/20$; so does this allowance. The small chart required here is supplied on any healthy neighborhood by axial gauge: its link distances are bounded by a fixed multiple of $(2z+1)^5$ times the threshold. The interior extension and its inverse chart place it inside radius $p - 2\,\mathrm{furl}$. The radius has rate zero and the threshold rate $-1/2$, whereas p decays more slowly; the scheduling thus gives the strict margin eventually.

The reference offers are indexed by the coupling itself: at a curve parameter t, the selected offer is the one at $\beta(t)$. It is therefore the same offer for two presentations with the same coupling, independently of how that point was reached. The fixed-reference finite-volume estimates hold eventually in their approximation depth simultaneously on the compact coupling band; (RB4b) gives the number denoted by gleam_i in (FCC1). The compatibility needed above comes from the same local cylinder formulas in the different reference views.

Here is also the weighted-tail step behind the retirement argument. At one finite torus all supports have cover number bounded by its site count. Thus multiplying any summable envelope by its support exponential, even the early exponent $s+3$, costs only a finite constant depending on that torus. Changing its derivative anchor also costs a finite constant:

$$K = \big((1+D)^p\big)^k, \tag{RB5}$$

where D is the maximal torus separation; the direction triangle inequality gives this bound. Consequently the local source estimate supplies one summable envelope in the common anchored listing. If the debris indicator of each fixed address is eventually zero, dominated convergence for series makes its weighted debris sum tend to zero at that parameter. This eventual removal is literal: a nonkept address eventually belongs to the finite retirement roster, whereas a kept address never belongs to the debris.

Continuity in the deviation budget extends that small-tail bound to a parameter neighborhood. More explicitly, spend half a requested allowance on the tail at the center and half on the full source difference to a neighboring parameter, converting the latter to the common anchored weighted norm by the finite constant (RB5) and its support-weight counterpart. Restricting the difference to the debris subfamily cannot increase its positive majorant. Thus the neighborhood works for every sufficiently large depth. A finite subcover of the compact band gives a single depth. Combining it with the uniform beacon depth proves the simultaneous retirement used in (FCC1). The canonical fixed-layout curve argument already given then transfers this depth to every compatible presentation; the construction here supplies its weighted continuity and actual finite beacons.

Here are the reference and continuity inputs with their finite-volume dependence exposed. In this section alone n is a fixed reference side, $V = n^4$, B is the fixed budget, $k = B.\mathrm{order}$, $p = B.\mathrm{power}$, and $e = B.\mathrm{supportExponent}$. The five signed fourth-difference coefficients in (RB1) are denoted throughout by ρ_j; the λ_j formerly used for them in (RB4b) mean the same coefficients. They are unrelated to the parameter λ in the local fusion argument.

The representation constant used here is $\mathrm{royalty}(G) = \sqrt{m}$, where m is the dimension of the fixed faithful unitary representation. In particular it is at least one. It is distinct from the Fourier aliasing price, which depends on a strip width and an analytic norm.

The reference-chart radius and margin have the explicit values

$$p_0 = (x + s_G \log L)^{-7/16}, \qquad d_0 = (x + s_G \log L)^{-1}, \qquad h = (x + s_G \log L)^{-23/20}. \tag{RB6}$$

The logarithm of the normalized finite Gaussian integral is centered by subtracting its value at the identity. On the reference chart $\|P\| + 2d_0 < p_0$, with no flags, its derivatives through the purchased order satisfy $\|\mathrm{D}^\ell f(P)\| \le h p_0^{-\ell}$. The lower bound on its unnormalized integral is strictly positive. These statements are obtained on the actual reference forest and its Gaussian marginal, with the same stationary section and coupling, by the Gaussian differentiation and connected estimates (ASP2a)–(ASP2c), the local integral normalization (PRI2)–(PRI5), and the logarithm series (ASP1)–(ASP2). In particular the bound concerns the centered logarithm, not the integral itself. The reference side must exceed the prescribed patch and covariance read radius. The path width, path norm, loop allowance, correction parameter and derivative horizon are fixed first; the odd block length is then above their finite lower bounds; finally x is above their common delay. Bounds for the two adjacent loop allowances and the correction parameter put both sides in these same estimates.

The exact logarithmic density identity on the strict healthy chart has the form

$$f(P) = c - \sum_{y \in \mathrm{Site}(n)} F_y(\text{actual } P) - \sum_{\alpha} A_\alpha(\text{actual } P). \tag{RB7}$$

Here F_y is the retained finite ordinary reference density and A_α the emitted regular source. Their actual occurrence indices and supports are retained in the sum. The nonretained source consists of the omitted ordinary occurrences and omitted short stationary paths, not of a difference between Gaussian integrals on different tori. Its anchored early allowance is at most h once its approximation depth is large. The retained-source part has allowance at most $\mathrm{turnFee}(B)(h+h)$. Summing these two bounds gives $\mathrm{turnFee}(B)(h+h)+h$ per site. Summation over V sites and centering (RB7) gives the second term of (RB4a); the reference logarithm gives its first term after the logarithmic chart chain rule already proved there. This explains every summand in q_j.

We give the tail argument that supplies the omitted-source allowance used in (RB7). A regular address is a pair (o_α, X_α), its anchor and finite support. For nonnegative numbers a_α, its early bound u means

$$\sup_{o} \sum_{\alpha : o_\alpha = o} a_\alpha \exp\{(e+3)\,\mathrm{cov}_B(X_\alpha)\} \le u. \tag{RB8}$$

Sum over the V anchors. Since each exponential is at least one, $\sum_\alpha a_\alpha \le Vu$. This proves absolute summability of the derivative envelopes, including derivatives in any fixed finite list of directions: divide those directions by one common scalar until they satisfy the threshold bound, apply (RB8), then multiply by that scalar to the derivative order. The same scalar works for every address. Termwise differentiation on the strict healthy open set therefore follows from uniform absolute convergence of the derivative series through order k.

Changing every anchor to the one in the common listing multiplies its envelope by at most $K_{\mathrm{anc}} = ((1+D)^p)^k$, as in (RB5). On this fixed torus,

$$\mathrm{cov}_B(X) \le V, \qquad K_{\sup} = \exp\{(e+3)V\}$$

bounds every early support weight. Thus an unweighted summable envelope a_α gives the summable weighted envelope $g_\alpha = K_{\mathrm{anc}} a_\alpha e^{(e+3)\,\mathrm{cov}_B(X_\alpha)}$. For the debris predicate

$$P_m(\alpha) = \{\alpha \text{ is not kept and does not belong to the depth-}m \text{ finite roster}\}$$

one has $1_{P_m(\alpha)} \to 0$ at every fixed address: ordinary occurrences exhaust their finite shelves and short paths exhaust their length cutoffs. Dominated convergence for the nonnegative series yields $\sum_\alpha 1_{P_m(\alpha)} g_\alpha \to 0$. This proves a small early bound, rather than merely a small unweighted sum.

For clarity, source deviation has the following meaning. Two sources have identical regular and forced address layouts. There are nonnegative envelopes c_α, d_α such that on each regular healthy domain every left derivative through order k of their difference is bounded by c_α, and on every configuration every forced coefficient difference is bounded by d_α. Both total weighted sums

$$\sum_\alpha e^{e\,\mathrm{cov}_B(X_\alpha)} c_\alpha, \qquad \sum_\alpha e^{e\,\mathrm{cov}_B(Y_\alpha)} d_\alpha$$

are at most the deviation allowance δ. Regular directions use the anchor-dependent budget (AB1)–(AB4). Forced coefficients use their full configuration domain.

The dependence on the coupling and input activities is continuous in this deviation sense. On a finite set of occurrence addresses, the local integrands and their derivatives vary continuously: the Gaussian covariance and mean, finite Haar factors, cutoff profiles and stationary coordinates are the actual ones already constructed. Their local integral bounds (ASP2b)–(ASP2c) provide a common integrable majorant. The normalized denominator stays positive in a sufficiently small neighborhood. Differentiation and dominated convergence hence give continuity of each regular derivative supremum on its compact direction and healthy test sets; the full-domain forced bounds use the corresponding full-domain envelopes. For the infinite family, the signed connected expansion gives a common summable majorant on that neighborhood, by (ASP3)–(ASP5). Given $\delta > 0$, first choose a finite address set whose common tail is below $\delta/4$, then use finite continuity to make its sum below $\delta/2$. The two tails contribute at most $\delta/2$. This proves deviation continuity with the actual support weights.

More explicitly, the dependence of the stationary path-only terms on the coupling is affine: if $A_\alpha(\beta)$ denotes one such term, then

$$A_\alpha(\beta) - A_\alpha(\beta') = (\beta - \beta')\{A_\alpha(1) - A_\alpha(0)\}.$$

The two early envelopes at couplings zero and one give a summable Lipschitz envelope for this difference. The other connected terms are compared by the same marked interpolation before summation. If the input source discrepancy is t, their regular and forced absolute envelopes have sums bounded by tM and tQ for fixed finite M, Q, whenever $0 < t \le t_0$. Signed translations cost

$$Z = |\mathrm{Seat}(n)|\,\mathrm{Rim}(n)\,\mathrm{turnFee}(B)M.$$

Reindexing does not change a sum; multiplication by the symmetry normalization and the direction-change allowance gives exactly these factors. Converting a total unweighted envelope to a weighted one costs at most $A = e^{eV}$. Thus the outgoing source deviation is bounded by $t\{1 + AZ + AQ\}$. All these constants may depend on this reference torus, budget, fixed path bounds and derivative horizon, but are independent of the approximation depth.

Finally fix a parameter value and an early target $u > 0$. Pay $u/2$ for its debris tail using (RB8). Choose a parameter neighborhood with full source deviation below $u/(2K_{\mathrm{sup}})$; then its restricted debris differences have early allowance at most $u/2$. The triangle inequality gives the same tail bound u throughout that neighborhood for all sufficiently large depths. A finite subcover of the compact coupling band makes the depth uniform. Applying this to the five fixed reference tori proves the common h allowance in (RB7), then (RB4a) and (RB4b), and hence the beacon term in (FCC1). No torus-uniform version of K_{anc} or K_{sup} is used.

8.18 Derivation and transfer of the reference chart estimate

The reference chart comparison uses a stationary expansion on fixed reference tori. Constants at this point may depend on the chosen reference torus and finite horizon. The subsequent periodic-copy

and activity arguments identify where that dependence is admissible. A finite reference calibration should therefore not be read as a uniform inverse estimate at every level.

We first derive the protected-logarithm jet input. Fix the stage and derivative order before the delay. Put $y = \beta_F$, $u = y^{-7/16}$, $m = y^{-1}$, $\delta = \beta^{-1/2}$, and $D = d\,\mathrm{fineMonitor}(t)$. Work at centers with $\|v\| + 2m < u$; eventually $\delta < u$ and the stationary radius exceeds u. Every $z \in \overline{B}(v, m)$ then satisfies $\|z\| + m < u$, which is the margin needed for the Gaussian logarithm estimates throughout this ball. Let I be the actual protected integral, $V(\delta, \cdot)$ its truncated Gaussian approximation with fluctuation cutoff D, and $p = \Omega \exp(-\mathrm{coverLog}) > 0$. The Gaussian comparison on the whole closed ball $\overline{B}(v, m)$, using the vertical gap $g > 0$, is

$$0 \le I(z) - pV(\delta, z) \le pe^{-gD^2/8}, \qquad V(\delta, z) \ge a = e^{-1}/2. \tag{PL1}$$

The positive omitted tail is bounded by the quadratic Gaussian tail estimate; the local comparison includes the chart and Haar factors. Here Ω is field independent, whereas $p(z) = \Omega e^{-\mathrm{coverLog}(z)}$ is a positive function of the center. The inequalities and the following logarithm argument are pointwise in that center. The cutoffs equal one near zero fluctuation, where the density is strictly positive, so $I > 0$. Its Gaussian determinant is positive by the vertical gap. Writing $I = pV(1 + q)$ gives $0 \le q \le e^{-gD^2/8}/a$. Hence the logarithmic discrepancy

$$f = \log I - \log \Omega + \mathrm{coverLog} - \log V \quad \text{satisfies} \quad |f| \le a^{-1} e^{-gD^2/8}, \tag{PL2}$$

by $0 \le \log(1 + q) \le q$. This is only a value estimate.

We record how the actual integral comparison in (PL1) is obtained, so that its positive error and its domain are explicit. On the fixed forest fiber let $\mathcal{U}_z$ be the set represented by the fluctuation chart $U = \mathrm{draw}(V_z, \mu(z) + \delta i)$, $\|i\| \le D$, where μ is the stationary graph. The chart radius inequalities are $(1 + \mathrm{trailCost})u + \delta D <$ initial and the corresponding color-scaled displacement is below the group-chart radius. On this set all protecting flags vanish. The exact change of variables, including its Haar Jacobian and the vertical Gaussian determinant, gives

$$\int_{\mathcal{U}_z} \mathcal{I}_z(U)\, d\nu_F(U) = p(z)V(\delta, z), \qquad I(z) = p(z)V(\delta, z) + \int_{\mathcal{U}_z^c} \mathcal{I}_z(U)\, d\nu_F(U). \tag{PL1a}$$

Here $\mathcal{I}_z$ is the nonnegative summed fiber density used in the exact push. This is why the comparison has a one-sided error.

The outside-field energy bound is paid in the same chart. Denote by k the number of free bonds, by $w_g(\beta) > 0$ the single-bond chart normalization, by v_C the common finite size bound, and by K_0, T_0 respectively the local derivative and determinant expenses. With $R_D = gD^2/4$, its normalized logarithmic expense is

$$Y_D = -k \log w_g(\beta) - R_D + v_C\, \mathrm{loopsPrice} + K_0 + T_0. \tag{PL1b}$$

The fiber complement estimate is $\int_{\mathcal{U}_z^c} \mathcal{I}_z \le p(z) e^{Y_D}$. It uses the local link reconstruction to put every small-toll, small-wire configuration in the safety chart. Otherwise its fine or prediction penalty pays at least the threshold energy. The parameter inequality

$$\beta\, \mathrm{meter}(v_C)u + R_D \le \beta p_{\mathrm{penalty}} m_0^2/2$$

makes that alternative at least as expensive as the fluctuation tail; m_0 is a fixed positive safety threshold, not the shrinking radius m. The quadratic lower bound within the chart pays the other alternative. Both comparisons use the same actual field and forest fiber.

The delay makes $Y_D \le -R_D/2$, and also the stronger inequality $Y_D + E + 2 \le -R_D/2$. Here the extra expense used to locate a clean chart point is

$$E = \beta\{v_C C_{\mathrm{cl}}\|z\|^5 + v_C C'(\theta'^3 + \theta'^4)\} + K_0 + T_0 + v_C A_{\mathrm{coarse}}/2.$$

The two terms multiplied by β and the loop expense are eventually bounded by one: the field radius has rate $-7/16$, the threshold has rate $-1/2$, and the remaining size factors have rate zero. The nondecaying expense K_0 is affine in the fine monitor, $K_0 = a(t) + b(t)\mathrm{fineMonitor}(t)$, with a, b bounded by the already fixed polynomial size budgets. The normalization and determinant expenses have the same type of bound. The square $R_D = (gd^2/4)\mathrm{fineMonitor}(t)^2$ dominates this fixed finite list, after the stage parameters and positive d are fixed. This proves both inequalities. The stronger one, together with the positive lower bound for the total push, forces a clean point inside $\mathcal{U}_z$; the local chord and displacement estimates then keep all chart flags zero on the surrounding radius-m ball. Thus (PL1a) applies throughout that ball. The first inequality and (PL1b) give its error $e^{-R_D/2} = e^{-gD^2/8}$, as claimed. The truncated Gaussian lower bound and its jet estimates use the same size and radius choices. Their derivation is the local Gaussian expansion above; no full-field lower bound for an inverse density is used here.

The Gaussian approximation uses the actual functions constructed above. Keep the retained maximum-norm real coordinates x, and give the vertical real coordinate array its Euclidean norm, denoting it by h. The linear map $\mathcal{L}(x, h)$ that joins these arrays into the joint maximum-norm array has norm at most one. In this notation set

$$\phi(x, h) = F(\mathcal{L}(x, h)), \qquad \psi(x, h) = \mathrm{loop}(U(x, h)) - \sum_{e \text{ free}} \log J_G(\mathrm{tint}(h)_e). \tag{GL0}$$

Here F is the main-action-plus-prediction score (AS1), U is its actual fine configuration, loop is the retained quadratic-loop contribution, and J_G is the local Haar density of the chart in the earlier primitive integral. The minus sign makes $e^{-\psi}$ contain the positive Haar factor. Let $\rho(x)$ be the vertical Euclidean version of the unique graph (US7), and let $M(x)$ be its Hessian in these vertical coordinates. The equality of counting and Euclidean color squared norms proves

$$D_h\phi(x, \rho(x)) = 0, \quad D_h^2\phi(x, \rho(x))[z, z] = z^{\mathsf{T}} M(x) z, \quad M(x) \ge (g/2)I. \tag{GL0a}$$

Thus the inverse covariance bound used below is $M(x)^{-1} \le (2/g)I$; it is the retained half-gap of the actual stationary Hessian.

We explain the common derivative hypotheses for these functions. The joined points used below satisfy

$$\|\mathcal{L}(x, \rho(x) + tbz)\| \le R_{\mathrm{path}}\|x\| + |b|D, \qquad 0 \le t \le 1, \quad \|z\| \le D,$$

where R_{path} is the stationary Lipschitz price, enlarged to include the retained identity coordinate. The radius inequalities before (PL1a) put this whole set inside the purchased stationary domain and put every color logarithm inside half the Haar-chart radius. The graph jets are (US8). The score jets follow from (AS6) after the norm-at-most-one map $\mathcal{L}$; converting coordinate multilinear bounds to operator bounds costs at most the number of joint coordinates to the derivative order. The loop jets follow by summing the absolutely weighted path jets, multiplied by the fine-site count. A Haar summand is the fixed analytic $\log J_G$ composed with a color projection of norm at most $1 + \dim \mathfrak{g}$, so its order-k price is at most its purchased chart price times $(2 + \dim \mathfrak{g})^k$; summing costs the free-bond count.

The centered determinant jets required below are controlled on this same domain. Differentiate $D \log\det M = \mathrm{tr}(M^{-1}DM)$ and $D(M^{-1}) = -M^{-1}(DM)M^{-1}$. Repeating gives a finite sum of

traces of products of inverse factors and derivatives of M, indexed by ordered derivative partitions. Each trace is bounded by the vertical dimension times the product of the operator norms. The inverse factor has bound $2/g$, and derivatives of M are obtained from two score derivatives and the graph chain rule, all at the already purchased higher orders. This gives a finite polynomial price in those jet bounds, the coordinate counts and $2/g$. The value of the centered determinant logarithm follows by integrating its first derivative from zero. These operations give the common price C for the graph, classical score, quantum contribution and centered half-log-determinant used below. Its dependence on the auxiliary size bounds is polynomial at fixed derivative order; those sizes have the assigned polynomial logarithmic bounds. Consequently C has rate zero. This establishes the hypotheses of the local Gaussian calculation on the entire displaced set, using the actual score, chart and graph.

Define the Taylor-integral expression

$$A(b,x,z)=\int_0^1 (1-t)D_h^2\phi(x,\rho(x)+tbz)[z,z],dt. \tag{GL1}$$

Stationarity and twice integrating a line derivative give $\phi(x,\rho(x)+bz)-\phi(x,\rho(x))=b^2A(b,x,z)$, including $b=0$ by the displayed integral. In particular $A(0,x,z)=\frac{1}{2}z^{\mathsf{T}}M(x)z$. Put

$$\begin{aligned}G(b,x,z)=\tfrac{1}{2}\log\det M(x)-\tfrac{1}{2}\log\det M(0)+A(0,0,z)-A(b,x,z)\\ +\psi(x,\rho(x))-\psi(x,\rho(x)+bz).\end{aligned} \tag{GL2}$$

Then the same truncated approximation in (PL1) is

$$V(b,x)=\int_{\|z\|\le D} e^{G(b,x,z)}d\gamma_{M(0)^{-1}}(z). \tag{GL3}$$

At $b=0$, multiplying its fixed Gaussian density by the exponential in (GL2) gives exactly the density of $\gamma_{M(x)^{-1}}$. Hence

$$V(0,x)=\gamma_{M(x)^{-1}}(\overline{B}(0,D)). \tag{GL4}$$

This identity retains the moving determinant and covariance; replacing either by its value at zero would lose the derivative argument.

Here is a uniform estimate for the mass in (GL4). If a centered Gaussian covariance Q satisfies $Q\le cI$, its one-coordinate exponential moment gives $\mathbb{P}(|Z_i|\ge v)\le 2e^{-sv+cs^2/2}$ for $s\ge 0$: apply Markov's inequality to $e^{\pm sZ_i}$ and add. If the number of coordinates is at most $P\ge 1$, and $Pv\le D$, the complement of the radius-D ball is contained in the union of those coordinate events. Indeed, if every $|Z_i|<v$, then $\|Z\|\le\sqrt{P}v\le Pv\le D$. Thus its probability is at most $2Pe^{-sv+cs^2/2}$. Take $P=1+Q_0(t)$, where Q_0 is the fixed polynomial size budget, and write $H=\mathrm{fineMonitor}(t)$. Eventually $P\le\sqrt{H}$. For $D=dH$, choose $v=d\sqrt{H}$ and $s=v/c$. The error is at most

$$2P\exp(-d^2H/(2c)). \tag{GL5}$$

The monitor is chosen to dominate these polynomial size budgets and logarithmic coupling powers, so (GL5) is smaller than y^a for every fixed exponent a, after increasing the delay. This estimate is uniform for the moving covariances, since the inverse-Hessian bound gives one common c.

For completeness the derivative price in (GL3) is finite without any derivative of a moving integration set. At order s, put

$$\begin{aligned}&U_s=2(1+C+D),\quad A_s=D^2\mathcal{B}_s(U_s,C),\\ &F_s=2C+2A_s+2\mathcal{B}_s(U_s,C).\end{aligned}$$

Here C bounds the purchased section, score, quantum and centered determinant jets, with two extra score and section orders for (GL1). The affine displacement map in (GL1) has price U_s; evaluating its second score derivative on two copies of z gives the factor D^2. Integration against $1-t$, followed by the sum in (GL2), bounds all jets of G by F_s. Since $G(0,0,z)=0$, its first derivative also gives $|G(b,x,z)| \le F_s \max(|b|,\|x\|)$. The shrinking chart makes this at most one. Consequently $e^{-1} \le e^G \le e$ on the fixed cutoff ball. Its base Gaussian mass is eventually at least one half by (GL5), so $e^{-1}/2 \le V(b,x) \le e$. The exponential composition rule bounds its jets by $\mathcal{B}_s(F_s,E_s)$; the fixed Gaussian measure is a probability, so integration preserves this bound. The logarithm on the compact interval $[e^{-1}/2,e]$ has a finite fixed jet price L_s, giving the joint (b,x)-jet bound $B(t)=\mathcal{B}_s(\mathcal{B}_s(F_s,E_s),L_s)$. All these prices have rate zero.

Finally (GL2) is unchanged by simultaneously replacing (b,z) with $(-b,-z)$. Symmetry of the cutoff ball and base Gaussian makes $V(-b,x)=V(b,x)$, and therefore its logarithm is even in b. Each of its x-derivatives has zero first b-derivative at zero. Integrating its second b-derivative bounds the difference from its $b=0$ value by $B(t)b^2$, through order j, using the purchased joint order $j+2$. At $b=0$, (GL4)–(GL5) and $|\log v| \le 2(1-v)$ for $1/2 \le v \le 1$ give an arbitrarily small polynomial value error on the surrounding ball. Apply the interpolation (PL4) below with a rough joint bound $B(t) \le y$ and a sufficiently small chosen value exponent. This bounds the x-jets of $\log V(0,x)$ by y^{-2}. Combining the even difference and this zero-parameter estimate gives the asserted Gaussian logarithm jet price $B(t)\delta^2+y^{-2}$ at $b=\delta$, as well as its lower bound. This proves the Gaussian input used in (PL1) and (RBG2).

The separate rough derivative estimate for the discrepancy on that closed ball is, for a requested order j,

$$\sup_{z\in\overline{B}(v,m)} \|D^i f(z)\| \le y^{S_0} \quad (0 \le i \le j+1), \qquad S_0=(j+1)^2+2. \tag{PL3}$$

For this rough estimate one can avoid estimating an inverse integral on a whole chart. We give the relative derivative bound used in the construction. Write the exact push as an integral over a fixed forest fiber of a finite sum of nonnegative functions

$$I(x)=\int \sum_{\alpha} a_\alpha(U) e^{-\beta N_\alpha(x,U)}\, d\nu_F(U), \qquad a_\alpha(U) \ge 0. \tag{PL3a}$$

The coefficients include the fine density and all factors independent of the retained coordinate x. The only x-dependence here is the prediction penalty N_α, with the fiber configuration fixed. In particular this argument does not require a nearby point to retain a selected protecting flag. Let r bound the number of coarse bonds. The link product and trace estimates in (AS3) give the penalty jet price

$$A_N = p_{\text{penalty}}\, r\, \mathcal{P}(2^n o_n \,\text{colorScale}).$$

The predictor at fixed U is a constant group element of matrix norm colorScale; hence this price is uniform in U and the finite choice α.

For clarity write the composition polynomial as

$$\mathcal{B}_n(a,b)=n!b(2+a)^n.$$

It bounds the finite partition sum for composition when the inner and outer jets have respective prices a,b at scale one. Subtracting the value of $-\beta N_\alpha$ at the center costs a factor two in its jet bound. If E_n bounds the exponential jets at zero, each integrand in (PL3a) therefore has jets through order n bounded by its own value times

$$K_n=\mathcal{B}_n(2|\beta|A_N,E_n).$$

Integrating preserves this relative bound. Differentiation under the integral follows from the same estimate on a neighborhood together with the bounded fine-density envelope and the finite Haar fiber measure. Consequently $\|D^i I(x)\| \le K_n I(x)$ for $i \le n$. At each center where (PL1) gives $I(x) > 0$, normalize by this single positive value: $J_x(z) = I(z)/I(x)$. Its jets at $z = x$ are bounded by K_n, and $J_x(x) = 1$. If L_n bounds the logarithm jets at one, the chain rule gives jets of $\log I(z) - \log I(x) = \log J_x(z)$ bounded by $\mathcal{B}_n(K_n, L_n)$. The size of $I(x)^{-1}$ has canceled before estimating. This is the normalization that supplies rough logarithm derivatives uniformly over the surrounding ball.

The other terms of f have explicit prices as well. If $C(t)$ is the stationary chart-jet price, set

$$T_n(t) = \mathcal{B}_n(2 + 2C(t), C(t)), \qquad O_n(\beta, t) = \beta T_n(t) + T_n(t) + C(t).$$

Composition of the classical action, quantum action and Gaussian determinant on the stationary graph bounds the centered cover logarithm by $2O_n$. Let $w(t)$ be the rate-zero Gaussian-log coefficient and put $Q(t) = |w(t)| + y^{-2}$; since $\delta \le 1$, the Gaussian logarithm jet price is at most $Q(t)$. Combining the centered logarithms and then restoring the value bound (PL2), which is eventually at most one, gives the rough price

$$M(t) = \mathcal{B}_n(K_n(2y, r(t)), L_n) + 2O_n(2y, t) + 2Q(t) + 1, \qquad n = j + 1. \tag{PL3b}$$

The actual coupling lies below $2y$, so positivity and monotonicity of these polynomials justify that substitution. The bond-count bound $r(t)$, chart coefficient $C(t)$ and $w(t)$ have rate zero. Thus A_N has rate zero, K_n has rate n, and $\mathcal{B}_n(K_n, L_n)$ has rate n^2. The cover term has rate one and the Gaussian term rate zero. Their sum has rate at most $S_0 - 1 = n^2 + 1$, and hence is eventually at most y^{S_0}. This proves (PL3), including the zeroth derivative by (PL2). It uses actual integral derivatives and the surrounding ball, rather than differentiating the small value error.

Here is the interpolation and its order of choices. If $\|D^i h\| \le v_i$ on a ball of radius s and $\|D^{i+2} h\| \le C$, integration along unit directions gives

$$\|D^{i+1} h(0)\| \le 2v_i/s + Cs. \tag{PL4}$$

For multilinear derivatives one tests unit vectors before applying this line estimate. Suppose rough derivatives through order $j + 1$ are bounded by y^d on a ball of radius y^{-a}, $a > 0$. For any target exponent b, a sufficiently small value bound y^B on that ball implies every jet through order j at its center is bounded by y^b. At order zero take $B = b$. For the step from order i to $i + 1$, choose

$$c = a + \max(0, d - b) + 1, \qquad b' = b - 1 - c,$$

then use the order-i assertion with target b' and ball exponent c to determine B. Eventually $2y^{-c} \le y^{-a}$, so the smaller balls around points within y^{-c} of the center lie in the original ball. The induction gives $\|D^i h\| \le y^{b'}$ there. Formula (PL4) gives $2y^{b'+c} + y^{d-c}$, whose two exponents are strictly below b; their sum is eventually at most y^b. Lower jets already satisfy the stronger bound.

Apply this recursion with $a = 1, d = S_0, b = -2, j = n$. It fixes B first; then choose the tail accuracy and delay so that (PL2) is at most y^B. Thus the discrepancy jets are bounded by y^{-2}. Gaussian differentiation bounds the separate local Gaussian logarithm jets by $H(t)\delta^2 + y^{-2}$, with H of rate zero. Keeping that term in L_0 leaves the discrepancy Y in (RBG1). Since $u \le 1$,

$$\max_{i \le n} u^i \|D^i Y(v)\| \le y^{-2}. \tag{PL5}$$

This includes $v = 0$. The critical extension and actual chart composition use the already purchased higher jets and finite composition prices. Those prices are fixed before the tail and reference delays.

This supplies the input to the rate calculation below and hence (RB6); the retirement argument then uses the resulting bound with its existing delay order.

We supply the quantitative calculation behind (RB6). Put $y = x + s_G \log L$, $A = y^{-7/16}$, and $\epsilon = y^{-2}$. All derivative orders and stage parameters in this calculation are fixed before x tends to infinity. A function of rate r is eventually bounded in absolute value by $y^{r+\eta}$ for each $\eta > 0$. Products add rates, a finite sum has the largest of its rates, and multiplication by any fixed constant preserves a rate. Thus a rate strictly smaller than $-23/20$ gives the bound in (RB6), with coefficient one, after increasing the delay.

Write the actual reference logarithm, before centering, as

$$g = \log \Omega + L_0 + Y + \beta f_1 + W. \tag{RBG1}$$

Here Ω is the field-independent Gaussian normalization; Y is the logarithm of the normalized protected local integral after its Gaussian main part is removed; L_0 is the remaining Gaussian normalization term; f_1 is the difference between the prescribed coarse classical action and the fine classical action evaluated on the stationary graph; and W is the difference between the prescribed coarse quantum action and the thermal determinant and Haar contribution on that graph. These are functions on the same real primary-coordinate space. In particular (RBG1) is an equality of the actual densities before any estimate is taken.

For a function F, use the scaled derivative bound

$$\|F\|_{n,A;z} = \max_{0 \le q \le n} A^q \|D^q F(z)\|.$$

The protected-integral comparison and interpolation (PL1)–(PL5) give $\|Y\|_{n,A;z} \le \epsilon$, including at the origin. The derivative price of L_0 on the radial segment from zero to z is

$$d = X(x) g_{\mathrm{step}}(\beta)^2 + \epsilon, \qquad g_{\mathrm{step}}(\beta) = \beta^{-1/2}, \tag{RBG2}$$

with X of rate zero. The integral itself is positive: the Gaussian measure is positive, the protected integrand is nonnegative, and on a nonempty open neighborhood of the zero fluctuation all protecting cutoffs are one and the exponential density is strictly positive. The Gaussian has positive density there. Its normalizing Gaussian determinant is positive by the vertical Hessian gap. Hence its real logarithm is defined on the chart. The estimates for its derivatives follow by differentiating its absolutely dominated integral and using the logarithm expansion established in (ASP1)–(ASP2).

Let $M \ge 1$ bound the reference coordinate and site counts, let H be the stationary-response price, and let K_n be the sum of the purchased bare-action jet price and the stationary-composition price. More explicitly, the coordinate conversion multiplies an order-q multilinear price by M^q; the latter composition price is the fixed polynomial $\mathrm{clump}_n(1 + H + (1 + H))H$ appearing in the ordered-partition derivative bound. Thus K_n has rate zero. Classical matching removes the jets through order four, giving

$$\|\beta(f_1 - f_1(0))\|_{n,A;z} \le \beta K_n A^5. \tag{RBG3}$$

The higher jets of W, of orders three through n, have a price C_n of rate zero. This price is the bare jet bound multiplied by $M^n(1+|\alpha|+\Lambda)$, together with the same stationary-composition polynomial and H. Here α is the bounded kinetic correction and Λ is the path allowance. The finite-period thermal comparison bounds its second jet at zero by

$$\|D^2 W(0)\| \le A_0, \qquad A_0 = M^2\{M^4 m(\mathfrak{g}) M\}\, \Theta p\, e^{-\gamma c}. \tag{RBG4}$$

In this formula $m(\mathfrak{g}) = (1 + \dim \mathfrak{g})^2$ is the fixed color-coordinate conversion factor. The reference side is c, $\gamma, \Theta > 0$ are the purchased decay constants, and $p \ge 0$ is bounded by the adjacent loop allowance. Only the displayed exponential dependence on c is used below.

The first jet is smaller than a generic extensive-action bound would suggest. The signed symmetries give $Dg(0) = 0$, while classical matching gives $Df_1(0) = 0$. Differentiating (RBG1) therefore yields

$$DW(0) = -DL_0(0) - DY(0), \qquad \|DW(0)\| \le d + \epsilon. \tag{RBG5}$$

Taylor's formula along segments in the chart, with these first and second jets and the third-and-higher jet prices, gives

$$\|W - W(0)\|_{n,A;z} \le C_n A^3 + 3\{(\epsilon + d)A + A_0 A^2\}. \tag{RBG6}$$

For orders zero, one and two this follows by integrating the next derivative and using (RBG4)–(RBG5). For an order $q \ge 3$, multiply the unscaled bound $C_n A^{3-q}$ by A^q. Taking the maximum and enlarging the three low-order contributions gives the factor three. Centering Y costs at most 2ϵ; centering L_0 costs dA. Equations (RBG1), (RBG3), and (RBG6) now give

$$\|g - g(0)\|_{n,A;z} \le dA + 2\epsilon + \beta K_n A^5 + C_n A^3 + 3\{(\epsilon + d)A + A_0 A^2\}. \tag{RBG7}$$

Here $y/2 \le \beta \le 2y$. Consequently d has rate -1. The lower bound on the reference side makes $\Theta p e^{-\gamma c}$, and hence A_0, have rate -1. All remaining prices have rate zero. The respective rates on the right of (RBG7) are at most

$$-23/16, \quad -2, \quad -19/16, \quad -21/16, \quad -23/16, \quad -15/8.$$

Every one is strictly below $-37/32$, and $-37/32 < -23/20$. The defining rate inequalities, with a sufficiently small positive slack and then a common delay, therefore make the whole finite sum at most $y^{-23/20}$. This derives the quantitative gain rather than assuming that a small integral automatically has a small logarithmic jet.

It remains to explain why the estimate applies to the five reference tori, whose sides need not be physical periods. Choose a physical coarse side divisible by the given reference side; the divisibility requirement was included in the block calendar. Copy the reference forest periodically to this larger torus. The uniform stationary construction gives a critical section on each torus. On the larger torus the newly constructed section and the retained physical section have the same zero value and solve the same critical equation in the same uniqueness corridor. They agree near zero. Their tested classical jets and their thermal second jets therefore agree there. The periodic-copy comparison transfers classical matching and the finite-period Hessian error to the reference section. Its decay exponent can be replaced by the minimum of the two positive exponents, and its price by the sum of the two prices. This is sufficient for (RBG4). No equality of the Gaussian integrals on different tori is required.

On a fixed reference torus, differentiation along the logarithmic link chart costs the factorial and direction factors already displayed in (RB4a)–(RB4b). The omitted-source allowance comes from (RB8), uniformly in the compact coupling band. The same normalized offers depend on the coupling value, so substitution of the actual coupling curve preserves these estimates and their continuity. Summing the finite-site logarithmic and omitted-source bounds gives the reference offer bound used in (FCC1). This completes the passage from the actual stationary integral to the retirement beacons.

The uniformity required in this last assembly can be made explicit. For a fixed transition between stage profiles, the local argument is purchased at derivative order $n + 6$ for the stationary score and at the indicated higher orders for the bare-action and Haar jets. Its conclusion uses order $n \ge 2$. The thermal comparison is first known for the retained critical section. A newly selected uniform analytic stationary graph is an extension of this section: the two solve the same

vertical equation, and uniqueness near zero identifies them. Therefore their classical matching and finite-period second-jet comparison are the same tests, with the enlarged common price and the smaller common positive width. This establishes the hypotheses of (RBG3)–(RBG5) for the graph used in (PL1), rather than replacing that graph by a different critical solution.

For each fixed incoming stage only finitely many outgoing transition tests are required. Take the maximum of their coupling delays. Enlarging the calendar's delay preserves all the estimates, since they hold for every coupling above that delay. At an actual calendar step the identities for the block length and next coupling identify its local parameter with $y = \beta_F(t, L)$. The two adjacent loop caps are bounded by the same local loop allowance, and the kinetic correction satisfies

$$|\alpha| \le \text{drift} + |s_G \log L| + 1.$$

The assumed stage bounds on paths and loops therefore place both sides of this actual step in the local comparison just proved. This accounts for the order of choices and for the fact that the estimates hold for all allowed coupling values in the compact band.

Finally put $s = y^{-23/20}$, let q be the derivative order of the current budget, and let T be its fixed conversion price for left derivatives. After retiring the source tail to budget s, the centered retained source on any one reference torus of side c_i has left derivatives bounded by

$$F_i = q!\, 3^q s + 2c_i^4\{T(s+s)+s\}. \tag{RBG8}$$

The first term is the finite composition bound from the logarithmic coordinates to the prescribed left derivatives. The other terms pay the two source differences and the finite sum over sites; their factors are positive, so restriction to the omitted subfamily does not enlarge them. The retirement depth is common over the compact coupling band by (RB5)–(RB8), and common over the finite reference tori by taking its maximum. Each reference side is chosen above the read radius, and its finite joint-coordinate, fiber and site counts have the same preassigned size bound. No constant in this step is chosen from the later macroscopic volume subsequence.

If the fixed interpolation stencil has signed weights a_i, its support reach is e, and the budget support power is p, transporting these local estimates to one insertion gives

$$\sum_i |a_i|\, 2^q (1+e)^{pq} F_i. \tag{RBG9}$$

Indeed the directional substitution costs 2^q, the support enlargement costs $(1+e)^{pq}$, and the finite sum is bounded with absolute weights. Smooth dependence of each reference offer on its coupling allows substitution of the actual continuous coupling curve. Formula (RBG9) is the bound for the reference offer in (FCC1). It is on the same ordinary source and common finite stencil as the retirement construction, and supplies the needed passage from the local protected integral to its assembled source bound.

8.19 The local activity norm and analytic comparison estimates

The norms in finite calibration are the paid activity norms, not the uniform norm of the resulting density. Fix the current budget P and a common address layout. Let $b_P, k_P > 0$ be its regular and forced normalization allowances. For an activity pair $A = (E, K)$, define

$$\|A\|_P = \inf\{t \ge 0 : E \text{ has anchored regular derivative budget } tb_P,\ K \text{ has forced hit budget } tk_P\}. \tag{LCN1}$$

The regular budget tests every left derivative through P's order, in its permitted directions, on the healthy domain of that address, with its exponential support weight. The forced hit budget takes the

absolute bound on every configuration and sums over addresses hitting each site. Both include the fixed locality and invariance conditions. The infimum is attained: take the least supremum envelope for each regular derivative family and each forced coefficient, and pass every finite nonnegative partial-sum inequality to the infimum. Hence $\|A\|_P \le t$ is equivalent to the two displayed budgets with that exact t. Triangle inequalities for derivatives and coefficients give subadditivity; scaling gives $\|cA\|_P = |c|\|A\|_P$.

Fix, for each signed translation, one permutation implementing that symmetry on the layout. The activity pairs obeying all these fixed permutations form a linear subspace $\mathcal{H}_P$. Its seminorm is (LCN1). Use its separated completion $\widehat{\mathcal{H}}_P$; the map from the original presentation is an isometry after identifying seminorm-zero differences. This retains the off-mask regular germs required by the derivative budget, rather than identifying them merely by their evaluated values outside the healthy domain. The scalar-plus-activity space has norm $\|(\lambda, A)\| = \max(|\lambda|, \|A\|_P)$. The actual retained activities have norm at most $1/8$ in these coordinates.

We explain how the local primitive estimates give an analytic estimate in this norm. For a finitely supported signed measure $w = \sum_j c_j \delta_{t_j}$ on $[-1, 1]$, say it has moment price r at radius $p \ge 1$ if

$$\left| \sum_j c_j (t_j/p)^n \right| \le r \qquad \text{for every } n \ge 0. \tag{LCN2}$$

A family $f(t)$ has comparison price C at radius p if every such measure satisfies $\| \sum_j c_j f(t_j)\| \le rC$. This is the precise signed-combination bound used here. If $f(t) = \sum_n (t/p)^n a_n$ with $\sum_n \|a_n\| \le C$, the bound follows by interchanging a finite sum and an absolutely convergent series. Conversely, define the moment vector $((t/p)^n)_n$ in ℓ^∞. The signed-combination bound defines a bounded linear map on their span and then its closure. Coordinate vectors belong to that closure: subtract the already recovered lower coordinates from the moment vector at a small nonzero t, divide by $(t/p)^n$, and let $t \to 0$. Its image supplies the coefficients. Restricting from radius p to $p' < p$ sums the geometric coefficient majorant and gives an absolutely convergent series. This is the recovery step used to turn signed-combination bounds into analytic line estimates.

On an affine activity line $A + tB$, regular occurrences and forced coefficients are affine in t on the same configuration: their masks and inventories are fixed. The finite calibrated coupling along the line is $\beta(t) = \Omega[\Omega^{-1/32}(\lambda + tu)]^2$, with $|u| \le 1$. The comparison center has $\Omega^{-1/32}\lambda \in [33/40, 47/40]$, and the climate bounds ensure the whole line stays in the allowed coupling band. Assume $\|A + tB\|_P \le 1$ for $|t| \le 1$. The primitive's signed regular and forced estimates, followed by the Hessian projection and the fixed dictionary fusions, give prices

$$\text{output activity: } d \text{ at radius } q/2, \qquad \text{scalar correction: } f_i \text{ at radius } q, \quad f_i = \text{chores}_i\,/\theta_i^2. \tag{LCN3}$$

The regular signed bound is tested in the anchored derivative budget. The forced signed bound has one nonnegative summable hit-envelope common to all line parameters and configurations. For any weights satisfying (LCN2), these two envelopes give budgets $r d\, b_{P'}$, $r d\, k_{P'}$ on the actual output layout, so (LCN1) gives the activity price in (LCN3). The scalar price follows by applying the linear Hessian reading to the same signed regular source estimate and dividing its paid bound by θ_i^2. This explains both the norm and the threshold denominator in (LCN3). The primitive locality, derivative and signed-sum bounds used in this step are the marked integral and source projection estimates of the preceding chapters; no assertion about analytic dependence of a chosen forest or local guide is required.

We make explicit the gain that pays (LCN3), since an individual bound on each line value would not imply the signed-combination bound. Use the 301 distinct interpolation nodes $r_j = j/301$,

$0 \le j \le 300$, in $[-1,1]$, and their Lagrange polynomials L_j. If w has moment price η at radius q, let $c_j = \sum_t w_t L_j(t)$ and $w_0 = \sum_j c_j \delta_{r_j}$. The coefficient bounds of these finitely many polynomials give $\sum_j |c_j| \le C_q \eta$ for one fixed $C_q \ge 0$. Every polynomial of degree at most 300 is interpolated exactly, so $w - w_0$ annihilates its moments through that degree. If $V \ge q$, its remaining moment prices at radius V are bounded by $\eta(1 + C_q)(q/V)^{301}$: factor $(q/V)^n$ from the degree-n moment and use $n \ge 301$. Thus a family with signed price B at radius V and individual value bound b has signed price at radius q at most

$$C_q b + (1 + C_q)(q/V)^{301} B. \tag{LCN8}$$

This is valid for the regular derivative budget as well as a scalar norm, because finite sums, absolute scalar multiplication and envelope sums preserve that budget.

We derive the preliminary large-radius signed price used in (LCN8). Fix the stage, block ratio and address layout. Write $y = \beta_F(x, L)$, $\phi = y^{1/256}$, $B = B(x)$, and let $A > 0$ be the fixed turning price. We first explain the parameter estimate on each vertex. A function polynomial in the line parameter t, of degree at most two, is reconstructed exactly from its values at $-1, 0, 1$:

$$f(t) = \tfrac{1}{2} t(t-1) f(-1) + (1 - t^2) f(0) + \tfrac{1}{2} t(t+1) f(1). \tag{LCN11}$$

The sum of the absolute coefficients of these three Lagrange polynomials is 4. If each nodal function has all required spatial derivatives bounded by $v d^m$ at derivative order m, (LCN2), applied to each parameter power, bounds its signed interpolation by $4p^2 v d^m$ at radius $p \ge 1$. The same identity holds after spatial differentiation: its parameter coefficients are fixed scalars. Thus this is a signed derivative bound, with no factor $\sum_t |w_t|$.

Logarithmic vertices enter as $e^{f(t)} - 1$, rather than $f(t)$. If their nodal derivative budget is at most $v \le 1/2$, first recover the budget of f using $f = \log(1 + (e^f - 1))$. The scalar value bound is $|f| \le 2v$; all derivatives of $\log(1 + z)$ through the fixed order are bounded on $[-1/2, 1/2]$. Its finite chain-rule sums therefore give a bound $b_j v$ with the same spatial derivative scale, for a constant b_j depending only on that order. Apply (LCN11), and then the absolutely convergent exponential series in the signed moment norm. Multiplication in that norm multiplies prices, so a log price $s \le 1$ gives $e^s - 1 \le (e - 1)s$; the differentiated exponential is controlled by the finite chain-rule sums at the same scale. Consequently there are constants $A_j \ge 0, B_j \ge 1$ such that the signed price of this vertex is at most $A_j p^2 v$, provided $B_j p^2 v \le 1$. These are the estimates FullVolume/Ray.Badge.starpath and White.faded, applied at the three displayed nodes.

They apply to the actual vertices because the line dependence is placed in the couplings and occurrence amplitudes, not in a newly chosen forest. With the spatial configuration held fixed, an old vacuum logarithm is constant or is multiplied by the square of the affine coupling-amplitude line, hence is quadratic in t. A regular occurrence has the prescribed degree-at-most-two amplitude representation on this line; multiplication by its fixed local enforcement factor preserves that degree. The forced occurrence has the same representation. Warning vertices, including their common steering and course variables, have the retained parameter-independent derivative envelope and therefore satisfy the constant-function signed estimate. In the differentiated formula the background, steering and Gaussian-coordinate variables are varied together as specified in the common chart. The spatial chain-rule prices are included in the nodal bounds. These are the four cases in FullVolume/Sector/Grant.mural, lured, loglift and alerts. In particular polynomial interpolation is applied to the log before exponentiation, not to the exponential.

The preliminary signed price now follows by the connected expansion with these vertex prices. Three nodal log envelopes sum to at most $3q$, where the old target q has rate at most $-7/16$. At $p = \phi$, their enlarged price has rate at most

$$-7/16 + 2/256 = -55/128 < -2/5.$$

Thus both $3q \le 1/2$ and $B_j\phi^2(3q) \le 1$ hold eventually, and the reweighted regular and log prices fit the target of rate $-2/5$. This is the explicit target enlargement in FullVolume/Sector/Reign.shows. The finite-order product rules multiply these prices at each connected tree. The support exponent exceeds the final exponent by more than three; split this excess into three positive reserves to pay the tree-degree factors, collar enlargement and log attachments. The positive tree sums have ratios at most $1/2$, hence cost at most 2 each. All remaining fixed-order spatial and collar prices have rate zero. Multiplying them by the enlarged forecast of rate $-2/5$ is eventually at most $y^{-3/8}$. Restricting the moment radius to $\phi/2$ improves the bound. The forced entries retain a common nonnegative hit envelope by the same positive majorants. This establishes the preliminary signed estimate for every finite combination satisfying (LCN2); it does not infer it from individual activity bounds.

Here are the product and summation steps in that signed estimate. A radius-R series certificate means an actual representation on $[-1, 1]$,

$$f(t) = \sum_{n\ge 0} (t/R)^n b_n, \qquad \sum_n \|b_n\| \le A, \qquad R \ge 1. \tag{ASP1}$$

For a finite signed combination $w = (w_t)$, put $\Lambda_w f = \sum_t w_t f(t)$. If $|\sum_t w_t (t/R)^n| \le \eta$ for every n, absolute convergence gives $\|\Lambda_w f\| \le \eta A$. This estimate uses cancellation of the combination; replacing w_t by $|w_t|$ would lose it. Products preserve (ASP1) with the product allowance, by the absolutely summable Cauchy product. Summable families preserve it by the double-series theorem. Exponentiation gives allowance $e^A - 1$ for $e^f - 1$, by summing the positive scalar majorant $\sum_{q\ge 1} A^q/q!$.

For the simultaneous spatial derivative estimates, fix a base coarse field x. Suppose each diagonal derivative through order j, in a unit spatial direction v, has (ASP1) allowance Ab^q at order q. If another factor has allowance Cd^q, the derivative of their product is

$$\sum_{r=0}^{q} \binom{q}{r} D_v^r f \, D_v^{q-r} g.$$

The product rule for the series certificates bounds its allowance by

$$\sum_{r=0}^{q} \binom{q}{r} Ab^r Cd^{q-r} = AC(b+d)^q. \tag{ASP2}$$

Induction gives, for a finite product of vertices, allowance $\prod A_i$ and slope $\sum b_i$. Thus the signed cancellation is retained before multiplying the vertices. After applying Λ_w, polarization bounds the full symmetric q-linear derivative by $\mathrm{polarCost}(q)\eta Ab^q$. Taking $\sum_{q\le j} \mathrm{polarCost}(q)$ as one common constant controls all orders through j. Finite linearity commutes Λ_w with these derivatives.

Here the polarization constant used by the source is explicitly

$$P_q = 2^q (q+1)^q, \qquad P_{\le j} = \sum_{q=0}^{j} P_q. \tag{ASP2a}$$

For a symmetric q-linear form A and vectors $\|v_i\| \le 1$, expand

$$\sum_{S\subseteq\{1,\dots,q\}} (-1)^{q-|S|} A\left(\sum_{i\in S} v_i, \dots, \sum_{i\in S} v_i\right) = q!\, A(v_1, \dots, v_q).$$

Every term whose list omits an index cancels by the alternating subset sum; a surviving list contains every index exactly once, giving $q!$ identical terms. If the unit diagonal bound is C, each subset term is bounded by $(q+1)^q C$, by homogeneity. There are 2^q terms and $q! \geq 1$. This proves the stated constant, including $q = 0$, and normalization of the nonzero vectors proves the operator norm bound. Applying it after the signed parameter combination gives the allowance $P_{\leq j}\eta A$ at all orders through j. It uses no bound depending on the dimension of the spatial chart.

The Gaussian integration prices entering these allowances also have a direct finite-dimensional proof. Let γ_Σ be the centered Gaussian actually used for a packet, with $\Sigma \geq 0$ and $|\Sigma_{ii}| \leq B$. The one-coordinate marginal gives

$$\int e^{tz_i}\, d\gamma_\Sigma = e^{t^2\Sigma_{ii}/2} \leq e^{Bt^2/2}.$$

Thus $(1+|z_i|)^q \leq e^{qz_i} + e^{-qz_i}$ implies $\int (1+|z_i|)^q\, d\gamma_\Sigma \leq 2e^{Bq^2/2}$. Independence of the coordinates is unnecessary. For a local edge set I, let J be its colored coordinate set, $a = |I|$, $b = |J|$, and

$$F_I(z) = 1 + a + b + \sum_{i\in J} |z_i|, \qquad Y = 1 + a + 2b.$$

Choose a coordinate attaining the maximum when $J \neq \varnothing$; then $F_I(z) \leq Y(1+|z_i|)$. The empty set case follows directly. Consequently

$$\int F_I(z)^q\, d\gamma_\Sigma \leq Y^q\{1 + 2be^{Bq^2/2}\}. \tag{ASP2b}$$

This proves integrability as well as the bound, by domination by a finite sum of the two exponential moments.

A tilt changes only the marginal on J. Its covariance Σ_J is positive definite. With $c \geq 1$, its exact density multiplier is

$$M_c(z) = c^b \exp\left(\frac{1-c^2}{2} z_J^T \Sigma_J^{-1} z_J\right), \qquad 0 \leq M_c(z) \leq c^b, \qquad F_I(cz) \leq cF_I(z). \tag{ASP2c}$$

This is obtained by the substitution $y_J = cz_J$ in the Gaussian density; the determinant contributes c^b. The quadratic form is nonnegative, proving the bound. The remaining scalp factor and this multiplier combine exactly into the original scalp times c^b times the exponential of the tilt's quadratic cost. Applying the product and exponential estimates to that equality gives derivative price charge$(j)(1+\delta)^{2b}$, with spatial slope bounded by $L'2^K F_I(z)^K$. Formula (ASP2b) pays its moments.

Finally a warning event E with probability at most r^{2m} and a polynomial factor with second moment at most A^2 contribute at most Ar^m, by Cauchy–Schwarz. Haar variables have total mass one and do not change this bound. All polynomial moment prices are included by taking q twice the required derivative degree. If the colored coordinate count is bounded by

$$2b \leq 2\,\mathrm{cov}_B(Y)\,\mathrm{ballCost}(B)\,\mathrm{density}(L)\,\dim \mathfrak{g}$$

and the tilt margin satisfies $2\delta\,\mathrm{ballCost}(B)\mathrm{density}(L)\dim\mathfrak{g} \leq \alpha$, then $\log(1+\delta) \leq \delta$ gives $(1+\delta)^{2b} \leq e^{\alpha \mathrm{cov}_B(Y)}$. This is precisely the reserved cover weight; the Gaussian determinant factor is paid before the final summation.

The interpolation step preceding this argument supplies those certificates at the real nodes. Explicitly, at $k+1$ distinct nodes t_i, let l_i be the Lagrange basis and $B_k = \sum_i \sum_{r=0}^k |[t^r]l_i|$. A

degree-k family with nodal derivative allowance Ab^q has (ASP1) allowance $B_k R^k Ab^q$: expand $l_i(Rz)$ and bound $R^r \le R^k$. For optional exponential vertices the nodal hypothesis is on $e^{f(t_i)} - 1$. When its allowance is at most $1/2$, $\log(1+y)$ has bounded derivatives on $[-1/2, 1/2]$ and $|\log(1+y)| \le 2|y|$. The finite chain rule therefore bounds the nodal logarithm through the same spatial order. Interpolate this logarithm, and only then exponentiate by (ASP1). The condition $B_j R^2(3q) \le 1$, with its fixed derivative constants included in B_j, pays this last exponential step. This explains why the source can use polynomial dependence of the log without asserting polynomial dependence of the exponential.

The remaining tree summation is positive. If vertex prices have hit allowance γ at support exponent $s + \delta$, $\delta > 0$, the rooted forest estimate of the connected expansion gives for the terms with m nonroot vertices the bound

$$e^{\delta \operatorname{cov}_B(Y)} \left(\frac{4 \operatorname{ballCost}(B) \gamma}{\delta} \right)^m. \tag{ASP3}$$

This estimate is proved by summing the leaf label at its first overlap with its parent: a radius-B cover of the parent has at most its cover number times $\operatorname{ballCost}(B)$ possible hit sites. The hit bound pays that sum; the reserve $e^{\delta \operatorname{cov}_B}$ pays the parent cover factors and the enumeration of parent assignments. The rooted occurrence enumeration and factorial weights give the factor 4^m already included in (ASP3), as in the absolute tree estimate used earlier. A designated marked root instead pays its own weighted allowance, multiplied by the same geometric factor. For ratio at most $1/2$, summing $m \ge 0$ costs at most two. Polynomial factors from the spatial direction growth are paid by the collar reserve:

$$\operatorname{collar}_{j,p} = (\operatorname{shift}(x))^{pj} \sup_{r \ge 0} r^{pj} e^{-\delta r}. \tag{ASP4}$$

The supremum is finite (with the usual value one when $pj = 0$); it is just the elementary absorption of a fixed polynomial by a positive exponential. The shift is bounded by a fixed multiple of the local support width plus one.

The packet's additional factor $(1+\varepsilon)^{n_A}$, $0 \le \varepsilon \le 1$, is also paid on the whole packet. The number n_A of relevant off-tree color coordinates satisfies $n_A \le C \operatorname{shares}(A) \operatorname{Delay}(x)^5$, by the link census and the $10B$ visit bound. Since $1 + \varepsilon \le e$, and any fixed positive derivative constant is at most its exponential, these factors are absorbed by increasing the coefficient of $\operatorname{Delay}(x)^5 \operatorname{shares}(A)$ in the packet income. They do not change its gain exponent or replace a whole-packet cost by a single-site hit.

In the final scheduling set $\alpha = (Q_{\exp} - e)/4 > 0$, where $e < Q_{\exp}$ is the requested output exponent. Then $e + 3\alpha \le Q_{\exp}$. Spend the three α's separately in the tree bound, collar bound and attached-log bound. The enlarged root price has rate $-2/5$; the degree ratio in (ASP3) has that rate times a width factor of rate zero, so it is eventually at most $1/2$. The fixed derivative and collar factors also have rate zero. Their product with the root price is eventually below $\widehat{\beta}^{-3/8}$, since $-2/5 < -3/8$. The parameter radius is $R = \widehat{\beta}^{1/256}$, eventually at least two; restrict to $R/2$ for the final bound. These are the simultaneous choices giving the common nonnegative hit envelopes for the signed forced entries and the series certificates for the attached logs.

Retirement selects a fixed subset of this same nonnegative majorant. Adding the boundary correction and turning operation gives the common signed price

$$w(x) = A\{y^{-3/8} + 4y^{-7/16}\}.$$

Every reference torus and comparison ball in the graft has size polynomial in $B(x)$ at a fixed stage. Their site counts $N_i(x)$ and $N_*(x)$, and their derivative factors, have rate zero. Summing the

two compared anchored expressions over a reference torus gives signed price $p_i(x) = 2N_i(x)w(x)$. Summing the unrestricted value bound over the comparison ball gives price $N_*(x)w(x)$. Both retain rate $-3/8$.

Let z be the graft radius, let c_* bound the cover number of its radius-z support, and let λ_i be the fixed finite interpolation weights. If the derivative order and spatial derivative power of the budget are j and p, put

$$T_z = 2^j(1+z)^{pj}, \qquad O(x) = \sum_i |\lambda_i| T_z p_i(x).$$

For a signed measure with moment price η at radius $\phi/2$, the finite interpolation derivative formula gives beacon envelope $\eta O(x)$ at each site. There is one anchored beacon per site. Its support weight therefore gives anchored budget at most $\eta O(x) e^{(\text{supportExponent}+3)c_*}$. The retired part has budget $\eta w(x)$. Since the regular graft is linear in these parts, the complete regular signed price is

$$w(x) + AO(x)e^{(\text{supportExponent}+3)c_*}. \tag{LCN10}$$

The cover bound c_* is fixed before x, and all radius factors have rate zero. Thus (LCN10) has rate $-3/8$, and is eventually at most $y^{-1/4}$. This is the rough price at radius $V = \phi/2$ required in (LCN8). It applies to the deterministic source because the retirement index depends only on the layout and the matching identities identify the actual source with this graft.

The forced half uses separate all-configuration envelopes. Its primitive entries retain their exponential flag suppression. The attached local insight entries have a signed scalar price obtained from the comparison-ball and reference-torus value prices, multiplied by their ball count and support weight. Denote the resulting primitive and attached hit prices by a and b. The delay purchases $\kappa \geq 0$, $b \leq 1$, and

$$\gamma = e^{\kappa}a + 2b, \qquad (2B+1)^4\gamma \leq 1, \qquad e^{-\kappa}(e^{\gamma}-1) \leq \text{forecastK}\,.$$

The local contribution b has rate $-3/8$ and is eventually below the comparison ceiling $y^{-1/8}$. The forced forecast uses the primitive exponential suppression as well; it does not follow from comparing two power rates. To fuse the two families, inflate each old forced atom by e^{κ}, bound each attached exponential minus one by twice its logarithm envelope, and sum the connected components using the displayed support reserve. Each component contains an old forced atom, so undoing the inflation saves at least $e^{-\kappa}$. The resulting hit price is at most $e^{-\kappa}(e^{\gamma}-1)$, as proved in the attached-log fusion argument. This gives the forced forecast at radius $\phi/4$, with one support-exponent unit reserved. Shrinking later to $q/2$ improves the moment inequalities and supplies the common forced envelope in (LCN3).

The sharper individual price in (LCN9) comes from universal retirement. Its remaining beacon envelope is a finite sum, with fixed coefficients, of

$$\text{gold}(x), \qquad N_i(x)\,\text{gold}(x), \qquad 2^j(1+z)^{pj}N_i(x)\,\text{gold}(x),$$

where $\text{gold}(x) = y^{-23/20}$. Every multiplier has rate zero. The fixed support exponential and turning price preserve rate $-23/20$, proving the claimed rate of ε_j. Thus finite interpolation combines a large-radius signed estimate and a sharper real individual estimate. For the primitive source, the individual bound b is its regular price (FCC1) plus the beacon contribution. With $x = \widehat{\beta}_i$, it has form

$$b = 3t\,C_j x^{-11/10} + \varepsilon_j(x), \qquad \varepsilon_j \text{ has rate } -23/20, \tag{LCN9}$$

where $t > 0$ is the freely chosen small primitive allowance. The three principal payments are the coefficient $t + 2t$. The other terms are the turning fee times the retired tail, the tail itself, and the

beacon price times its fixed support exponential. The local estimates give them the stated $-23/20$ rate at fixed block and stage. The larger signed radius is $V = \mathrm{fire}(x)/2$, where $\mathrm{fire}(x) = \beta_F(x)^{1/256}$; the rough signed price has rate $-1/4$. Consequently the last term of (LCN8) has rate at most $-301/256 - 1/4 < -23/20$. Choose t so that $6C_q t \le v$, and then the coupling-entry delay so that the two remaining terms cost at most $(v/2)C_j x^{-11/10}$. Equations (LCN8)–(LCN9) give signed regular source price at most $vC_j x^{-11/10} = \mathrm{chores}_i$. This is the numerator of the scalar price f_i.

The Hessian reading of this source yields the scalar correction price f_i by linearity and its bound $\theta_i^2|\Delta| \le \mathrm{Budget}$. Applying the regular reset price to that same signed source, including the committed coefficient bound $|\alpha_{N,i}| \le d_{i+1} + s\log L_i$, gives regular output price $db_{P'}$ once the small allowance and delay have been chosen. The forced reset entries use their common all-configuration envelopes; the forced and attached-log fusion estimates retain one unit of support-weight reserve and yield hit price $dk_{P'}$ at radius $q/2$. The expenses and the forced cutoff are the paid local reset inequalities, not a consequence of the scalar reading alone. Combining these two prices through (LCN1) proves the activity assertion in (LCN3). Thus the analytic recovery uses a signed estimate obtained by a finite interpolation cancellation, with the exact exponent 301, rather than assuming analyticity of arbitrarily selected primitive witnesses. In fact the output values are independent of such incidental choices. To compare two representatives with identical coupling and activity data, use the same constant affine input line but select the first representative at $t = 0$ and the second at $t \neq 0$. The bounds (LCN3) still hold. For $q \ge 4$, the signed measure $\delta_t - \delta_0$ has moment price at most $|t|$ at either radius. Therefore its two output differences are bounded by $d|t|$ and $f_i|t|$; letting $t \to 0$ proves equality. This supplies the single-valued activity and scalar maps on the common normed presentation.

Choose $q \ge 4$ so large that

$$\frac{2e}{(1/4)(q/4)} \le 1, \qquad 2d \le \gamma = 1/128. \tag{LCN4}$$

At any center (λ, A) with $\Omega^{-1/32}\lambda \in [17/20, 23/20]$ and $\|A\|_P \le 1/8$, every neighboring center within distance $1/8$, followed by any direction of norm at most $1/4$, stays in the enlarged scalar window and activity ball used in (LCN3). Recover the line series at radius $q/4$. Polarization makes its degree-n coefficient a symmetric multilinear map: the finite difference formula divides by $n!$, has 2^n summands, and each sum of n unit vectors has norm at most n. Since $n^n/n! \le e^n$, the resulting multilinear coefficient is bounded by

$$\left(\frac{2e}{(1/4)(q/4)}\right)^n C, \qquad C = 2d \text{ or } 2f_i. \tag{LCN5}$$

The formulas agree on overlaps because they recover the same scalar line coefficients. They extend through the isometric dense inclusion into the completion, preserving their operator norms. By (LCN4), the activity coefficients are bounded by γ, and the scalar coefficients by $2f_i$, at every order.

For increments z, z' of norm at most $1/4$, telescope each multilinear diagonal and sum:

$$\sum_{n\ge 1} n(1/4)^{n-1}\|z - z'\|C = \frac{16}{9}C\|z - z'\| \le 2C\|z - z'\|. \tag{LCN6}$$

Thus the local Lipschitz constants are 2γ for activities and $4f_i$ for the scalar correction. To get the estimates for any two centers in the whole stated window and activity ball, split into distances at most $1/4$, where (LCN6) applies, and distances greater than $1/4$. In the latter case (LCN3) applied

to singleton measures bounds the two values separately by d and f_i; their differences are at most $2d \le \gamma$ and $2f_i$. Both cases imply

$$\|\Phi_i(y) - \Phi_i(x)\| \le 8\gamma\|y - x\|, \qquad |\Delta_i(y) - \Delta_i(x)| \le 8f_i\|y - x\|. \tag{LCN7}$$

With $s_i = 2f_i$, this is exactly the scalar constant $4s_i$ and activity constant 8γ used in (FCB). All prices are for the fixed step of the common calendar and are uniform in the actual cutoff and period when its paid stage and support margins hold. The local analytic constants are chosen before the final history is selected; the completion introduces no new choice of primitive data. The common root envelope in this argument uses the energy gain proved before the expansion, in its actual order of quantifiers. For every sufficiently large odd block side L, all sufficiently large fixed window widths W admit a number $\kappa > 0$ and a coupling threshold such that

$$\kappa\, t_{\mathrm{sf}}(\beta, L)^2\, n_{\mathrm{seed}} \le E_{\mathrm{assessed}} \tag{ASP5}$$

for every admitted tube, sector, tree and packet store above that threshold. Here n_{seed} is its number of charged seeds, t_{sf} its small-field threshold and E_{assessed} the sum of the actual fine, coarse and mismatch departure costs. The assignment of those costs to packet cores, the paid annuli and the three clipping ranges was proved in the energy chapter. Formula (ASP5) is a per-seed energy statement; the expansion additionally uses the support census, normalization expenses and reserves (ASP3)–(ASP4).

For adjacent scheduled profiles use the same angular covers for both copies of (ASP5). They provide one positive coefficient for each of the two requests; replace both by their smaller value and increase the common coupling threshold. The geometric block factor is fixed before these choices. The resulting calendar delays are chosen above both thresholds, and its actual coupling update identifies the next small-field threshold with the one in (ASP5). Every admitted step therefore has the claimed bound with the same charged seeds and the same angular read sets. The pointwise domination obtained this way gives one nonnegative envelope for all the signed entries, including when a parameter path changes which sector is active. Combining it with the tree and collar sums proves the all-root envelope used above.

8.20 Stability of finite calibration and its comparison weights

The retained forward recipe also gives quantitative stability of the complete finite trajectory under changes of its bottom coupling. This is the finite statement behind the auxiliary trajectories in the uniqueness argument. Write $x_i = \widehat{\beta}_i$ and define

$$w_i = x_i^{-1/32}, \quad \mathfrak{b}_i = x_i w_i = x_i^{31/32}, \quad \lambda_i = \frac{\sqrt{\beta_i/x_i}}{w_i}, \quad \beta_i = x_i(w_i\lambda_i)^2. \tag{FCB1}$$

The positive square root is used. By (FCA8), $w_i\lambda_i \in [17/20, 23/20]$; this wider interval contains $[\sqrt{3/4}, \sqrt{5/4}]$. The activity coordinate X_i is its image in the common completed activity chart. At a fixed finite horizon all these coordinates live in one finite product of their scalar and activity spaces.

Let $\gamma = 1/128$. The actual local analytic comparison, including the initial state at N, bounds differences of activity outputs by 8γ times the maximum input-coordinate difference. Its scalar correction output has Lipschitz bound $4s_i$ times that maximum. The weights are explicitly

$$s_i = 2\,\frac{\mathrm{chores}_i}{\theta_i^2}, \qquad \mathrm{chores}_i = v\, C_{\mathrm{stage}(i)} x_i^{-11/10}, \tag{FCB2}$$

where $v > 0$ is the fixed comparison cost and $C_j > 0$ the stage's local analytic price. The common small-field threshold is $\theta_i = x_i^{-1/2} \exp((\log\log x_i)^2)$. Thus $s_i \le 2vC_{\mathrm{stage}(i)} x_i^{-1/10}$. These comparison bounds apply to the same retained primitive operation at both points, as ensured by its fixed layout and choices.

The calendar pays the cumulative weight inequality

$$\sum_{j<i} s_j \le \gamma \mathfrak{b}_i, \qquad \mathfrak{b}_0 \le \mathfrak{b}_i. \tag{FCB3}$$

Here is its derivation from delays. At a fixed stage transition the positive reference increment is $a = s \log L > 0$. The mean value theorem, for $x \ge a$, gives

$$(x+a)^{31/32} - x^{31/32} \ge a \, \frac{31}{32} \, 2^{-1/32} x^{-1/32}.$$

Since $x^{-1/10}/x^{-1/32} \to 0$, eventually the upper bound in (FCB2) is at most γ times this increment. Include this one eventual threshold in the coupling-entry delay for each of the two allowed adjacent-stage possibilities. It follows at every realized step that $s_i \le \gamma(\mathfrak{b}_{i+1} - \mathfrak{b}_i)$. Summing gives $\sum_{j<i} s_j \le \gamma(\mathfrak{b}_i - \mathfrak{b}_0) \le \gamma \mathfrak{b}_i$. Positivity and monotonicity of x_i give $\mathfrak{b}_0 \le \mathfrak{b}_i$. Thus (FCB3) is a consequence of the same schedule, not a supplementary assumption on a chosen pair of trajectories.

Compare two successful trajectories of the deterministic recipe, with bottom-coupling difference δ. Let

$$U = \max_{0 \le i \le N} \max\{|\lambda_i' - \lambda_i|, \|X_i' - X_i\|\}.$$

The local comparisons give $\|X_i' - X_i\| \le 8\gamma U$, including the finest-state comparison, and correction differences at most $4 s_i U$. Subtracting the finite coupling recursions cancels their common path coefficients:

$$\beta_i' - \beta_i = \delta - \sum_{j<i} (\Delta_j' - \Delta_j). \tag{FCB4}$$

On the other hand (FCB1) factors their difference exactly as

$$\beta_i' - \beta_i = \mathfrak{b}_i (\lambda_i' - \lambda_i)(w_i \lambda_i' + w_i \lambda_i).$$

The last factor is at least $17/10 > 1$. Equations (FCB3)–(FCB4) imply

$$|\lambda_i' - \lambda_i| \le \frac{|\delta|}{\mathfrak{b}_i} + 4\gamma U \le \frac{|\delta|}{\mathfrak{b}_0} + 4\gamma U.$$

Taking the maximum over all levels and both coordinates proves

$$U \le \frac{|\delta|}{\mathfrak{b}_0} + 8\gamma U, \qquad U \le 2|\delta|/\mathfrak{b}_0. \tag{FCB5}$$

Only finite products and finite sums occur at this stage.

If $\delta = 0$, (FCB5) gives identical finest normalized couplings, hence identical finest actual couplings. Their startup states are equal by their fixed startup rule. Descending determinism of the primitive recipe then gives equality of all represented states. Thus the successful finite trajectory with a prescribed bottom value is unique within this retained recipe, although monotonicity of a scalar coarse readout was unnecessary for existence.

The arcs prove existence for every bottom value in the nearby interval. Combining them with (FCB5) gives continuous dependence of its normalized scalar and activity coordinates on that value,

on the same finite calendar. For example, choose a bottom perturbation $|\delta| \le (\ell/4)\mathfrak{b}_0$, with the fixed small comparison radius $\ell > 0$, and small enough to stay in $[3\bar{\beta}/4, 5\bar{\beta}/4]$. Then $U \le \ell/2$, placing the entire perturbed trajectory inside the common analytic comparison neighborhood. This constructs the nearby finite trajectories used in passing local comparison estimates to banked limits. Equations (FCB1)–(FCB4) also identify explicitly the positive scalar weight and summable correction weights in that argument.

8.21 Uniform running of the actual coupling

The drift estimate also proves the uniform running bound needed in the public theorem. Define

$$D_0 = 0, \qquad D_{i+1} = D_i + d_{i+1}. \tag{RUN1}$$

For the calibrated finite history, (FCA3), $\beta_{N,p,0} = \bar{\beta}$, and the exact recursion for $\widehat{\beta}_i$ imply

$$|\beta_{N,p,i} - \widehat{\beta}_i| \le D_i \qquad (i \le N). \tag{RUN2}$$

Indeed subtract the two recursions and apply the triangle inequality at each step. The right side depends only on the calendar, so this bound is uniform in both N and p.

The stronger asymptotic drift condition of that calendar is

$$\frac{d_{i+1}}{\log L_i} \le J(i+1)^{-2}, \qquad J(i) \longrightarrow \infty. \tag{RUN3}$$

Here $J(i)$ is the stage index, not a cutoff or period. Given $\varepsilon > 0$, choose j so that $J(i+1)^{-2} \le \varepsilon/2$ for every $i \ge j$. Summing (RUN3) gives, for $m \ge j$,

$$D_m \le D_j + \frac{\varepsilon}{2}(\tau_m - \tau_j).$$

Since $\tau_m \to \infty$, eventually $D_j \le (\varepsilon/2)\tau_m$. Together with (RUN2) this proves

$$\forall \varepsilon > 0\ \exists m_0\ \forall m \ge m_0\ \forall N \ge m\ \forall p, \qquad |\beta_{N,p,m} - \widehat{\beta}_m| \le \varepsilon\tau_m. \tag{RUN4}$$

This is uniform control of the actual finite-cutoff coupling, including $m = N$. It is stronger than merely writing down the reference line $\widehat{\beta}_m = \bar{\beta} + s\tau_m$.

Dividing by $\tau_m > 0$, and then absorbing the fixed intercept $\bar{\beta}/\tau_m$, yields

$$\sup_{N \ge m,\ p} \left| \frac{\beta_{N,p,m}}{\tau_m} - s \right| \longrightarrow 0. \tag{RUN5}$$

Explicitly apply (RUN4) with error $\varepsilon/4$, and take $\tau_m \ge 4(|\bar{\beta}| + 1)/\varepsilon$; the triangle inequality gives an error smaller than ε. Taking $m = N$ proves the bare coupling slope in the theorem, uniformly in the physical period.

Set $b_0 = 11C_A/(48\pi^2)$, so $2b_0 = s$. For arbitrary cutoff and period sequences $N_i \ge i$ eventually and p_i, let $g_i = \beta_{N_i,p_i,i}^{-1/2}$. Positivity of the coupling and (RUN5), followed by continuity of the square root and reciprocal at $s > 0$, give

$$g_i \sqrt{\log(a_i^{-1})} \longrightarrow (2b_0)^{-1/2}. \tag{RUN6}$$

The equality used here is $\beta^{-1/2}\sqrt{\tau} = (\beta/\tau)^{-1/2}$ for $\beta, \tau > 0$. This regulator coupling is defined from the inverse action coefficient. Its numerical leading logarithm does not by itself identify it with a coupling defined by a physical position-space observable. That separate identification, and its normalization, are proved in the trace and Ward calculation.

8.22 Comparing omitted interactions at a common normalization

The fine-to-intermediate comparison uses one fixed denominator throughout. If N is the finest level and p is the period, define, for a polynomial-valued integrable insertion P at level m,

$$\mathcal{I}_m(P;t) = \frac{\int [P(U)]_t \, d\mu_m(U)}{Z_{N,p}},$$

where t is a polynomial mode, μ_m is the product Haar prior, and $Z_{N,p} > 0$ is the original finest-level Wilson partition function. In particular the denominator is not replaced by a new partition function when m changes. The root estimate, for insertion weights $w_v > 0$, is

$$|\mathcal{I}_k(P;t)| \le B_k(t) := \frac{\exp\{|\Lambda_k|\,\tau_k + |V|(x+y+2)\}}{\mathrm{toll}(w,t)}.$$

Here τ_k is the vacuum treasury, (x,y) bound the two insertion envelopes, and the positive toll is the coefficient-extraction weight. This bound follows by applying the small-field integral bound to the represented density, bounding its vacuum bill by the treasury, and using the exact identity between the integrated represented density and $Z_{N,p}$.

At an intermediate level $j > k$, separate the two omitted interaction polynomials P_0, P_1 from the retained insertion data. Multiply both omitted parts by a complex parameter z, restore the resulting density, and call it (P(z)). Then (P(1)) is the full insertion and (P(0)) is the insertion with those parts omitted. The paid analytic estimate makes

$$Y(z) = \mathcal{I}_j(P(z);t)$$

holomorphic for $|z| < S_j$, where $S_j = \mathrm{brisk}(B_j) > 2$. For every such z, transport (P(z)) through the actual retained Mark steps from j to k. The retention estimate preserves the served-insertion bounds with root weights w_k; integration of the transported polynomial equals integration of the original polynomial. Coefficient extraction commutes with this integral because it is a continuous linear map on the finite-dimensional polynomial space. The denominator remains $Z_{N,p}$, so the root estimate gives $|Y(z)| \le B_k(t)$ uniformly throughout this disk.

For any holomorphic Y bounded by B on $|z| < S$, $S > 2$, every disk of radius (S/2) centered on the segment ([0,1]) lies inside that disk. Cauchy's derivative estimate gives $|Y'(z)| \le 2B/S$ on the segment. Integrating along the segment proves

$$|Y(1) - Y(0)| \le \frac{2B}{S}.$$

Consequently the omitted interactions change the normalized coefficient by at most $2B_k(t)/S_j$. The root cost depends on k, not on the intermediate volume at j. The analytic paid bound, the retained transport estimate, and the root small-field bound are the substantive inputs to this comparison and must be established below.

The holomorphy under the integral in this comparison can be justified directly. The polynomial algebra has finitely many modes; its truncated exponential is a finite sum of powers. For a fixed configuration the restored polynomial is a finite sum of products of this truncated exponential and the local field factors. Its masks and flags depend on the configuration, not on z. Replacing the omitted terms by zP_0, zP_1 therefore makes each mode an entire function of z. All masks are measurable: a prescribed flag pattern is a measurable event, and restoration is a finite sum of products on these events.

The weighted polynomial norm is equivalent to the ordinary finite-dimensional coefficient norm when all insertion weights are positive. The served-insertion estimate for the scaled omitted terms,

together with the envelope bounds and nonnegative coupling, gives a single finite constant Y such that

$$\|P(z,U)\| \le Y \qquad (|z| < S_j,\ U \in \mathcal{C}_j).$$

Since the Haar prior is a probability measure, this is an integrable majorant. It proves integrability at every z. Dominated integration of Cauchy's formula on small circles, or equivalently Morera's theorem with Fubini, proves holomorphy of each integrated coefficient. Thus no differentiability of the configuration masks is needed. The served-insertion envelope proved in the marked-integration chapter gives this uniform majorant.

The retained transport statement has a descending induction with no limiting step. At the top j, scaling the omitted parts spends two units of support exponent. Suppose the transported polynomial has the required served bound at (m+1). Its first-degree coefficient is the prescribed insertion manifest there, by the relay identity. Integration through the Mark kernel gives the prescribed manifest at m, by the insertion feed identity. The parent map sends the chosen pins at (m+1) to those at m. The matching assumptions supply their occupancy bound and the allowed reweighting of the insertion weights; the wake data supply the local finite transport tales and their estimates. The one-step duty estimate now gives the served bound at m with the full support exponent at that level. This proves the assertion down to k. Integrability at each stage follows from integrability of the Mark push. This induction explains why the two-unit expenditure occurs only at the top; it does not accumulate once per step. The finite inventory and one-step duty estimates in the marked-integration chapter establish the inputs to this induction.

8.23 Comparing the varying Gaussian integrals

A change of protected variables also changes the Gaussian law. If $z' = z'(z)$ has Jacobian $J_V(z)$, the general density ratio is

$$J_V(z)\,\frac{\rho_{C_V}(z'(z))}{\rho_{C_0}(z)}.$$

For positive covariances on the same d-dimensional real space this is

$$J_V(z)\left(\frac{\det C_0}{\det C_V}\right)^{1/2}\exp\!\left\{-\frac{1}{2}z'^{\mathsf{T}}C_V^{-1}z' + \frac{1}{2}z^{\mathsf{T}}C_0^{-1}z\right\}.$$

The fixed-covariance expression follows only when $C_V = C_0$. The determinant and inverse derivatives in the protected-variable estimates below are precisely what control the genuinely varying case.

The fluctuation space at a fixed step is one finite-dimensional Euclidean space of dimension D. A prepared datum has Gaussian density $\rho_C(z)$, mean displacement t_X in physical coordinates, and scale $g = \beta^{-1/2} > 0$. Use the common physical variable $w = t_X + gz$. Then

$$z = g^{-1}(w - t_X), \qquad c_X(w) = g^{-D}\rho_C(g^{-1}(w - t_X)),$$

and for every integrable vector-valued function A,

$$\int A(z)\,d\nu_C(z) = \int c_X(w)A(g^{-1}(w - t_X))\,dw.$$

This is ordinary translation invariance of Lebesgue measure followed by scalar change of variables, including its determinant g^{-D}. It also proves integrability of the transformed integrand.

The common local covariance bounds give $\rho_C(z) \le C_0 e^{-\tau\|z\|^2}$ with $C_0, \tau > 0$ independent of the history index and gauge-slice variable. The protected mean gives $\|t_X\| \le B_0$: sum its finitely many squared coordinate bounds and take the square root. The coupling window gives fixed $0 < u \le g^{-1} \le v$. For $r = g^{-1}$ these imply

$$1 + \|r(w - t_X)\| \le (1 + v + vB_0)(1 + \|w\|),$$

$$e^{-\tau\|r(w-t_X)\|^2} \le e^{\tau u^2 B_0^2} e^{-\tau u^2\|w\|^2/2}.$$

The second inequality uses $\|w\|^2 \le 2\|w - t_X\|^2 + 2B_0^2$. For any fixed set of tested bonds its face-scale function is at least one and at most $K_0(1 + \|z\|)$, since each coordinate is bounded by the Euclidean norm. Consequently every transformed integrand whose original norm is bounded by a fixed power of that face scale has a common majorant

$$C_1 e^{-\tau u^2\|w\|^2/2}(1 + \|w\|)^k.$$

This is Lebesgue integrable: in finite dimension Gaussian decay dominates any polynomial. Its constant is independent of the gauge-slice variable.

The integral convergence argument is applied twice. At each gauge-slice point, the common majorant and pointwise convergence of the difference of the two transformed integrands give convergence of their Lebesgue integrals by dominated convergence. Fubini identifies these with their original Gaussian integrals for almost every gauge-slice point, simultaneously for the countably many history indices. Their norms are bounded by the integrals of the two majorants. The outer gauge-slice measure is a probability measure, so this constant is integrable. A second dominated-convergence application proves convergence of the full Gaussian–gauge integrals.

For ordinary primitive derivatives, differentiation under the integral identifies the derivative of the local activity with the integral of its differentiated integrand. Its paid bound is

$$C\big(1 + (p\,\mathrm{facescale}^K + 1)\big)^j \le C(p+2)^j\,\mathrm{facescale}^{Kj},$$

because the face scale is at least one. Thus it has the required polynomial majorant. Marked primitive derivatives have the same form of estimate from their common marked budget. The undifferentiated marked amplitude has a uniform constant bound, which is the case $k = 0$.

Pointwise comparison of the transformed integrands is taken with one common amended draft. At each fixed gauge-slice point, the draft data have only finitely many discrete possibilities; an ultrafilter selects one possibility eventually. The same viability, threshold and choice data are then used in both histories. Their density factors and local vertex products compare in these common physical coordinates; products retain both real and imaginary marked ingredients. The preceding two-stage integral argument converts this comparison into convergence of ordinary derivatives, marked derivatives and marked amplitudes. The finite reserve products and boundary-choice averages in the actual primitive definitions preserve it. Their healthy-support local identities identify these formulas with the actual leaflet and marked-grove germs near the tested base configuration. This is the Gaussian integral passage used in the single-stream comparison, with all discrete masks retained inside the common draft.

8.24 Common drafts and Gaussian density factors

The finite choice used in the preceding integral argument is precise. A draft records a gauge-slice variable, a store and a boundary choice, together with the fixed selector, scale and width. After fixing the gauge-slice integration variable, the paid conditions fix the latter three parameters. Only

the store and boundary choice remain variable, and they range over a finite set at a fixed lattice step. For any ultrafilter one pair is therefore constant eventually: otherwise the complements of all its finitely many fibers would have an empty intersection in the ultrafilter. Replacing the gauge-slice coordinate by the fixed integration point produces one common draft. Viability transfers to it on that eventual set. No continuous gauge-slice variable is claimed to become constant along the original histories.

In this common draft, the atmospheric physical fluctuation w is the same in the two histories. The restart map is

$$z_\beta(w) = \sqrt{\beta}\,(w - t_X).$$

The protected-mean comparison and coupling comparison imply comparison of this map. More generally its corrected version along an allowed output-field family replaces w by the common smooth steered field before subtracting the protected mean. The exact identity $t_X + \beta^{-1/2} z_\beta(w) = w$ ensures that the original real integrand is recovered; the analogous steered identity is used for its derivative integrand.

The Gaussian density formula is

$$\rho_C(z) = (\sqrt{2\pi})^{-D} \exp\{-\tfrac{1}{2}\log\det C - \tfrac{1}{2} z^{\mathsf{T}} C^{-1} z\}.$$

The covariance comparison, its uniform positive floor, and the finite-dimensional inverse and determinant formulas compare every term in this expression. The floor also ensures that any subsequential density limit at a fixed finite z is strictly positive: the logarithmic exponent has a finite limit and its exponential is positive. Thus division by the corresponding marginal density is legitimate in the local reweighting formula.

The same argument applies to every principal covariance restriction on a fixed finite set of tested bond/color coordinates. Its positive floor is inherited from the full covariance, and its inverse and log determinant depend smoothly on that restricted matrix. The local correction factor is exactly its finite Jacobian factor times the density of the corrected local fluctuation divided by the density of the uncorrected local fluctuation. The Jacobian is a finite product of the smooth link steering factors. Each steering factor is evaluated near the identity in the common draft; the chosen local chart supplies its smoothness through the required derivative order. Consequently this entire correction factor compares through that order. This establishes the density and local-change-of-variable factors used in the ordinary and marked integral comparisons; comparison of their remaining vertex factors is treated separately.

8.25 The protected change of variables and its determinant

The steering factors in the preceding ratio have the following construction. Fix a small positive radius d, independent of the later scale delay, on which the logarithm is a left inverse of the exponential through radius $2d$. Shrink a matrix neighborhood of 1 so that $\log(\exp a\, M)$ is analytic for $\|a\| \le 2d$, and so that for group elements M it exponentiates back to $\exp a\, M$ when $\|a\| \le d$. These choices follow by continuity on a closed ball strictly inside the original chart. Let $\chi_d(a)$ be the smooth inner-share cutoff, one on $\|a\| \le d$, zero for $\|a\| \ge 2d$. Define, with this multiplication order,

$$p_M(a) = \chi_d(a)\{\log(\exp a\, M) - a\}, \qquad \Phi_M(a) = a + p_M(a), \qquad j_M(a) = \det(1 + D_a p_M(a)). \tag{PV1}$$

Outside the cutoff support the first expression is identically zero in a neighborhood, so the map is smooth there regardless of the exterior values of the logarithm. On the support, smoothness follows from the chosen analytic chart. At $M = 1$, p_M and its derivative vanish, $\Phi_1 = \mathrm{id}$, and $j_1 = 1$.

Compactness of $\{\|a\| \le 2d\}$, joint continuity, and the exterior zero identity supply one neighborhood of $M = 1$ on which, uniformly for every a,

$$\|p_M(a)\| < d/2, \qquad \|D_a p_M(a)\| < 1/2, \qquad j_M(a) > 0. \tag{PV2}$$

For the determinant assertion, use its continuity and its value one on the compact ball; outside the ball it is exactly one. The derivative bound makes p_M Lipschitz with constant $1/2$. Thus Φ_M is injective, since equality of two images implies $\|a - b\| \le \frac{1}{2}\|a - b\|$, and surjective, since for each target b the contraction $a \mapsto b - p_M(a)$ has a fixed point. Its derivative is invertible by the Neumann series. Consequently it is a global smooth change of real coordinates with positive Jacobian j_M.

For an original coarse background V_0 and a nearby background V, the correction at a fine link is

$$M_e = (\text{replica } V_0)_e (\text{replica } V)_e^{-1}, \qquad u'_e = \Phi_{M_e}(u_e). \tag{PV3}$$

The order matters: for $\|u_e\| \le d$,

$$\exp u'_e \, (\text{replica } V)_e = \exp u_e \, (\text{replica } V_0)_e. \tag{PV4}$$

On inactive links the fine slice input remains unchanged. On active links the identity is needed only where their protecting cutoffs are nonzero. Those cutoffs imply $\|u'_e\| \le 8a\, t_{\text{field}}$. Choose the scale delay so that $8a\, t_{\text{field}} \le d/2$. Then (PV2) implies $\|u_e\| \le \|u'_e\| + \|p_{M_e}(u_e)\| \le d$, so (PV4) applies there. This proves equality of the old and new reconstructed configurations on every protected read. A cutoff product vanishing at the read makes the corresponding term zero. A nonzero product forces equality on its whole actual footprint. This argument proves the equality for old activities, atom words, row predictions, sector profiles and their complements individually, including the hard decisions inside them. It does not differentiate a hard decision.

Writing the physical fluctuation as $u = \mu_{V_0} + gz$, $g = \beta^{-1/2} > 0$, the corresponding Gaussian coordinate change is

$$z' = g^{-1}\{\Phi_M(\mu_{V_0} + gz) - \mu_V\}. \tag{PV5}$$

It obeys $\mu_V + gz' = \Phi_M(u)$ exactly, and becomes the identity at $V = V_0$. Restrict to the finite set I of bond/color coordinates required by the ground. Each link block of (PV5) has derivative $g^{-1} D\Phi_{M_e} g = D\Phi_{M_e}$; the scalar dilation factors cancel. The coordinate permutation between link/color arrays and Euclidean vectors contributes no absolute determinant. Therefore

$$\det Dz' = \prod_{e \in I_{\text{links}}} j_{M_e}(u_e) > 0. \tag{PV6}$$

This is a product over tested links, with one color-space determinant per link, rather than one determinant per color component.

Let C_0, C_V be the principal covariance matrices on I, with Gaussian densities ρ_0, ρ_V. Finite-dimensional change of variables gives for an integrable local function f

$$\int f(w)\rho_V(w)\, dw = \int f(z'(z)) \underbrace{\left(\prod_e j_{M_e}(u_e)\right) \frac{\rho_V(z'(z))}{\rho_0(z)}}_{L_V(z)} \rho_0(z)\, dz. \tag{PV7}$$

The denominator is strictly positive at every finite z. Both integrals can equivalently be written against the original full Gaussian, since the functions depend only on the indicated marginal and the unused coordinates integrate to one. This proves the local likelihood formula, including its direction and normalization.

Along a smooth coarse field family based at V_0, its finitely many matrices M_e lie in the common neighborhood from (PV2). Differentiation of (PV1) on the compact cutoff ball gives uniform finite-order bounds on the displacement and determinant; exterior derivatives vanish or equal those of the constant one. Composition with the coarse family gives a displacement price of the same small order as that family's motion and a determinant price of order zero. Equations (PV5)–(PV7) then supply the Gaussian score derivatives. The covariance floor and the finite inverse/determinant formulas control the score coefficients. The integrable Gaussian envelopes established above justify differentiation under the integral. Thus a marked entry can remain evaluated on the fixed original fine input: the derivatives fall on its protecting cutoff, on the smooth ordinary vertices, and on L_V. The protected equality (PV4) identifies this formula with the original integral on a neighborhood of the base point. This is the mechanism used in the marked primitive comparisons and their product estimates.

8.26 Finite-window covariance comparison

For clarity, covariance comparison here is obtained from its actual construction. The form matrix consists of the selected physical row operators evaluated on coordinate impulses. Selection is determined by whether the row base is outside the store's fixed six-radius interior. In an unselected row both histories have zero form entries. In a selected row the physical row comparison gives entrywise comparison through the required derivatives, with the fixed impulse held the same. The plate is the form Gram matrix plus the fixed dormant diagonal. Finite matrix multiplication and addition therefore compare its entries.

For each covariance column, its window is the fixed finite set prescribed by the reference frame. The principal plate restriction inherits the uniform positive floor. Inversion on this set of positive definite finite matrices is smooth, with derivatives controlled by the same inverse bound; hence the two restricted inverse matrices compare through the required order. Entries outside the window are identically zero in both histories. Placing the inverse column back in the full coordinates and symmetrizing gives the actual reference covariance, so it too compares. This retains the column-window construction and does not replace it by the full plate inverse.

The protected mean is obtained from these same window columns and its local bill vector. Their comparison gives the mean comparison used by restart. Together with the coupling-window comparison it supplies the affine physical-variable comparison. The scalar density proof above consequently rests on finite inverse matrices with a common lower eigenvalue bound. Its remaining physical-row and bill-vector inputs are precisely the stationary and local nonlinear comparisons, rather than an independent covariance convergence assumption.

8.27 Extracting the common activity comparisons

At each fixed level and physical period, the activity data lie in one common normed presentation space. Its dictionary, supports and threshold are fixed, and its norm controls anchored regular derivative mass and forced absolute mass. Its completion is used only for limits. The actual retained finite histories remain elements of the original space.

We spell out how a Cauchy sequence in this norm gives the evaluated comparisons needed above. For either the regular-occurrence map or the forced-coefficient map, the activity norm implies a nonnegative absolute budget of total at most $C\|a\|$, uniformly over configurations. At a fixed period one may take C to be the number of sites times the corresponding anchored budget constant. For regular terms sum the anchored masses over their unique anchors; for forced terms choose a site from each nonempty forced inventory contained in its support. The local mass dominates the

unweighted bound. This yields the stated global budget.

Choose its canonical least envelope

$$q(a,i) = \sup_U |F(a,i,U)|.$$

The existing absolute budget makes this supremum finite, nonnegative and summable with total at most $C\|a\|$. Linearity of F gives $q(a,i) \le q(a-b,i) + q(b,i)$. For a Cauchy sequence a_n, these envelopes have uniformly small tails. Indeed choose N with $C\|a_n - a_N\| \le \epsilon/2$ for $n \ge N$, and a finite set controlling the tails of the finitely many $a_0, \dots, a_N$ by $\epsilon/2$. On its complement the displayed subadditivity bounds every later tail by ϵ; the earlier tails are already controlled.

If a_n, b_n are Cauchy and $\|a_n - b_n\| \to 0$, then

$$|F(a_n,i,U) - F(b_n,i,U)| \le q(a_n - b_n, i) \le C\|a_n - b_n\|$$

uniformly in U for each label. The two envelope families also have the uniform tail property just proved. Thus regular occurrences with their actual fixed healthy masks, and forced entries with their actual inventories, have exactly the evaluated comparison and tail bounds used in the local construction. Masked regular evaluation is linear because the threshold and each labelled support are fixed in this presentation.

For any two escaping subsequences of a single Cauchy activity sequence, their norm difference tends to zero by the Cauchy criterion. The coupling sequence has the same property and converges in $\mathbb{R}$. The retained history's actual activities are eventually these normed presentation elements at every fixed level. Hence all incoming regular, incoming forced and outgoing forced comparisons required by the two-history transfer argument follow from the common activity Cauchy sequence. Completion convergence suffices: the completion embedding is isometric, so convergence there makes the original presentation sequence Cauchy.

The scalar and activity convergence themselves are obtained by the retained-history construction as follows. At each level its normalized scalar coordinate is bounded in a fixed finite interval, by the positive width and the uniform normalized-coupling bound. Compactness of the countable product of these intervals gives one diagonal escaping subsequence on which all scalar coordinates converge. The fixed-scalar construction then gives convergent activity data and the associated intermediate quantities at all levels along that subsequence. To promote this to the full sequence, repeat the same extraction on any escaping subsequence. The comparison of two such banked limits with the same limiting terminal coupling identifies their scalar limits. The subsequence criterion for convergence therefore gives full scalar convergence, after which the fixed-scalar construction gives the full activity convergence. The section on uniqueness of banked history limits proves the comparison identifying these limits. Compactness is used only to obtain the candidate limits.

For the winning trajectories the bottom coupling at level zero is the same prescribed center $\bar\beta$, so its convergence hypothesis holds. The finite-horizon alignment identifies their paths with the original stationary paths and their normed activity presentations with the actual retained activities. This provides the Cauchy sequences used throughout the fixed-level insertion proof, on the original common choices.

8.28 Finite-depth approximation of the activity trajectory

At a fixed finite horizon, regard the normalized scalar coordinates as an externally specified admissible sequence. The activity update maps the next-level activity to the current one and keeps its norm at most $1/8$. On this ball its paid comparison is

$$\|\Phi_i(\lambda,x) - \Phi_i(\lambda',x')\| \le c_0 \max\{|\lambda-\lambda'|, \|x-x'\|\}, \qquad c_0 = 8/128 = 1/16.$$

The accompanying scalar correction satisfies the corresponding bound with coefficient $4s_i$. These are the finite-history estimates used in the banked-limit comparison.

For any intermediate depth b, construct a comparison activity trajectory by setting its activity at depth b to zero and applying the same updates backwards below b; set unused later levels to zero. Finite decreasing induction constructs this trajectory. The paid ball-invariance estimate keeps every value admissible. Compare it with any full trajectory having the same scalar script. At depth b the distance is at most one (in fact at most $1/4$); each backward step multiplies that distance by at most c_0. Thus at $m \le b$ the discrepancy is at most c_0^{b-m}.

Now compare two sequences of histories whose scalar scripts have the same limit at every fixed level. For the trajectories truncated at the common fixed depth b, the top activity is identically zero in both histories. Descending induction, using the fixed-input one-step comparison at each of the finitely many levels, gives a common limit of their truncated activities. Given $\epsilon > 0$, choose $b - m$ so $c_0^{b-m} < \epsilon/3$. For sufficiently late histories both horizons exceed b and their truncated activities differ by less than $\epsilon/3$. The two truncation errors prove that the original level-m activities differ by less than ϵ.

Applying this argument to every pair of escaping subsequences proves that the activities have a common limit in their complete activity space. The scalar correction limit follows by the one-step comparison using the already convergent next-level activity. Finite initial segments discarded to ensure all required horizons are present do not change these limits. This supplies the fixed-scalar convergence used in the diagonal extraction and banked uniqueness argument. Its fixed-input one-step comparison is the local Gaussian/cluster/reset comparison developed above; the ball-invariance and contraction constants are the paid one-step estimates on that same update.

The physical scalar identity also survives this passage. At finite horizon the correction equals the difference of adjacent physical couplings minus the committed path coefficient. The latter has its stationary limit. Taking the three scalar limits therefore gives the displayed banked flow identity, with no infinite telescoping sum.

8.29 From labelwise comparison to reset budgets

At a fixed physical period there are finitely many site supports. The budget weight is linear in the nonnegative scalar bill, so there is a finite constant $P_* > 0$ with

$$p.\mathrm{mass}(D,t) \le P_* t \quad (t \ge 0)$$

for every support D. This constant may depend on the fixed period; no uniform volume assertion is used in this convergence argument.

Suppose a labelled difference has a nonnegative uniform-tail majorant $f_{n,i}$, and each fixed label eventually admits every positive tolerance. Given a desired total weighted budget ϵ, put $t = \epsilon/P_*$. Choose a finite label set E with every majorant tail at most $t/2$, and eventually bound each label in E by $t/[2(|E|+1)]$. Use these small bounds on E and the original majorant elsewhere. The resulting nonnegative array has total at most t, hence weighted budget at most ϵ. This applies both to the paid regular difference bounds and to the all-configuration forced difference bounds. It proves the quantitative difference-budget statements used to recover the activity norm.

The same construction controls the projected coupling increment. Each source separately has a common finite paid total, obtained by multiplying its uniform-tail bound by P_*. The diagonal projection estimate gives

$$t_0^2|\Delta_n| \le C, \qquad t_0 = p.\mathrm{threshold} > 0.$$

For the difference of two sources, choose its paid deviation budget below $t_0^2\epsilon/2$. Linearity of projection and its diagonal drift estimate then give $|\Delta_n - \Delta_n'| \le \epsilon/2$ eventually. Thus the two increment

sequences are uniformly bounded and their difference tends to zero. This proves the increment comparison used in the exact coupling update, rather than assuming a separate scalar continuity principle.

The reset dictionaries themselves are on common supports. The canonical stationary path presentations retain their fixed slot and atom supports, so the low homes and high footprints do not change with their numerical coefficients. The reset listing is therefore identical for both path families. Its regular and forced comparisons can be transferred through the final dictionary consumption on this same listing. The final vacuum and coupling scalars are irrelevant to that activity comparison. These observations isolate the remaining task in reset convergence: compare the actual residual, freshness, loan and attached-log coefficient functions with uniform tail budgets.

8.30 Comparison of the reset coefficient families

The regular reset family is the disjoint union of the source rebates, low freshness terms and high freshness terms. For a rebate, the difference identity is the rebate of the source difference. Its paid bound is multiplied by the fixed reset-loss constant, so both the uniform tail bound and each label's arbitrarily small comparison bound pass through. The fault difference has the same linear identity and the same all-field reset-loss bound, yielding the forced comparison.

For a fixed low label, the actual slots and support are independent of the stationary numerical coefficients. Only its mark m varies. Choose once a reference stationary coefficient array for which this label's mark $m_* \neq 0$ (all path coefficients equal to one suffice for this single-label algebraic test, and the direct mark is positive). Its freshness is then the fixed reference freshness multiplied by m/m_*. The coefficientwise path comparison makes these scalar ratios compare, so the fixed-label paid freshness difference tends to zero. Its uniform majorant is the original weighted low-mark budget times its fixed low-cost polynomial and the common threshold factor. Bounded-length low labels form a finite set, since each pair of path labels has total length bounded. The spare positive exponential width therefore makes these majorants have uniformly small tails.

High freshness is minus the sum of the two bounded compared scalar parameters times the actual cubic or quartic atom. The atom coefficient comparison gives fixed-label convergence; the weighted atom budget with spare width gives uniform tails. Finite signed-seat replication and the fixed orientation multiplier preserve all these bounds. The projected source increment is a bounded compared scalar by the previous argument, so it may be used as one of those two parameters.

For the low loans, the same reference-mark factorization gives exactly

$$\operatorname{loan}_R(\alpha, \Delta) = \operatorname{loan}_{R_*}\big((m/m_*)\alpha, (m/m_*)\Delta\big).$$

This follows from the literal loan definition: its live pair term is proportional to the mark, its proxy is independent of the numerical path coefficients, and its home and kept tests are common. In the direct case the mark is fixed. Subtracting the two reference loans is the reference loan at the differences of their two scalar arguments. Its all-field shape bound tends to zero as those arguments tend to zero. Thus the loan differences converge uniformly over configurations, even on the boundary of a failed health test.

The common absolute majorant for these loans is

$$e^{l/\rho}|m| \left(2t(B_\alpha + B_\Delta) + \tfrac{1}{2} B_\Delta \operatorname{pairBill} t^2\right).$$

The width hypothesis leaves a positive exponential margin after $e^{l/\rho}$ is spent. The weighted low-pair budget and the finiteness of bounded-length labels therefore give uniformly small tails. Combining faults and loans proves the attached-log comparison. Combining rebates and the two freshness

families proves the regular reset comparison, all on the same canonical listing and with its original countable enumeration.

The final consumption first turns every attached logarithm l_i into $e^{l_i} - 1$. If $|l_i| \leq v_i \leq B$, the mean value estimate gives $|e^{l_i} - 1| \leq e^B v_i$. A common total absolute budget supplies such a B, so the increment family has uniformly small summable tails. At each fixed label exponential continuity, applied to the bounded compared values under arbitrary varying configurations, gives convergence of the increment difference, and the varying-configuration selection argument makes this uniform in configurations.

The resulting forced family is indexed by finite subsets of the union of the old hard labels and these exponential-increment labels. Its coefficient is their finite product when the canonical support/eligibility test holds, and zero otherwise. The same test is used on both sides. Finite products of bounded compared values compare, and multiplication by this common indicator preserves the comparison even if the indicator varies with the history index. The finite-subset product budget preserves uniform tails. Canonical injective reindexing places these coefficients in the drawer, with zero outside its image; finite sums over the preimage establish preservation of both totals and tails. The regular family is likewise placed by its canonical support embedding. Hence the completed output arrays have the quantitative weighted difference budgets used in the activity-norm closure. This closes reset convergence without differentiating a hard indicator.

8.31 Closing the one-step convergence in the activity norm

Suppose two incoming activity presentations are Cauchy with norm difference tending to zero, and their couplings have a common limit. The canonical-envelope argument above gives the ordinary regular-occurrence and forced comparisons. The common stationary path and kernel limits then supply the microscopic comparisons for the two physical steps. Their actual output source dictionaries are expressed on one common listing by the local construction. The path reset is performed with those same dictionaries and the same committed coefficients.

The coupling after this step is exactly

$$\beta_{\mathrm{out}} = \beta_{\mathrm{in}} + \alpha_{\mathrm{path}} + \Delta_{\mathrm{source}}.$$

The projection comparison gives a vanishing difference of the two source increments Δ_{source}, and the stationary path coefficients have the same limit. Thus the two output couplings differ by a quantity tending to zero. The regular and forced reset comparisons give the analogous convergence of the output activity presentations.

To see the last implication in the actual norm, its budget characterization has two positive normalizing constants, b_{reg} and b_{force}. For a target $\epsilon > 0$, choose the regular difference budget below $(\epsilon/2)b_{\mathrm{reg}}$ and the forced difference budget below $(\epsilon/2)b_{\mathrm{force}}$, simultaneously at every anchor. The regular and forced comparison tail lemmas supply these bounds eventually. The common support and format data permit their use for the actual difference presentation. The norm characterization then bounds its norm by $\epsilon/2 < \epsilon$. Consequently the output norm difference tends to zero. This step requires the quantitative tail comparisons after reset, not just pointwise convergence of individual output coefficients.

The incoming completion data used by the finite-history update satisfy exactly these hypotheses. Their common completion limit makes their original presentation sequences Cauchy and their difference tend to zero by the isometric embedding. The continuous physical-coupling map sends their common normalized-scalar limit to a common physical limit. The chosen physical step identifies its actual output activities with the update's original presentation element; no new output

representative is selected for the comparison. The displayed one-step coupling identity also identifies the scalar correction as output coupling minus input coupling minus the committed path coefficient. This proves the fixed-input update convergence invoked in the finite-depth argument, once the regular and forced reset tail comparisons are established.

8.32 Uniqueness of banked history limits

The banked contraction compares histories at the same retained normalization. Its tail error tends to zero after the common finite-depth approximation is fixed. The double-tail Cauchy argument proves convergence of the whole relevant inner family, not merely convergence of an extracted subsequence. This is the convergence input used for the eventual full fixed-volume mesh limit.

Here is the comparison that makes the diagonal extraction unique. Consider two banked limits with normalized scalar coordinates f_i, f_i', completed activity data X_i, X_i', and local scalar correction limits g_i, g_i'. Their bottom coupling limits at level zero agree. Write

$$d_i = \max\{|f_i - f_i'|,\ \|X_i - X_i'\|\}, \qquad \gamma_0 = 1/128.$$

The one-step comparison, applied to the two limits by approximation with actual finite trajectories, gives

$$\|X_i - X_i'\| \le 8\gamma_0 d_{i+1}, \qquad |g_i - g_i'| \le 4 s_i d_{i+1}.$$

The scalar flow identity and equality at the base give, in physical coupling coordinates,

$$\Omega_i - \Omega_i' = -\sum_{j<i} (g_j - g_j').$$

The normalized-coordinate comparison has a positive weight b_i with

$$b_i |f_i - f_i'| \le |\Omega_i - \Omega_i'|, \qquad \sum_{j<i} s_j \le \gamma_0 b_i.$$

All these inequalities use finite sums at each i.

The original small activity and scalar bounds first give $d_i \le 1/4$ uniformly: for the scalar part apply the finite scalar-flow identity and the coarse individual correction bounds; for the activity part add the two individual norm bounds. Let $D = \sup_i d_i$, which is therefore finite and nonnegative. The one-step activity inequality yields $\|X_i - X_i'\| \le 8\gamma_0 D$. The scalar identity and correction inequality yield

$$b_i |f_i - f_i'| \le 4D \sum_{j<i} s_j \le 4\gamma_0 D b_i.$$

Since $b_i > 0$, the scalar difference is at most $4\gamma_0 D$. Hence $d_i \le 8\gamma_0 D$ for every i, and taking the supremum gives $D \le 8\gamma_0 D = D/16$. Thus $D = 0$, proving equality of every scalar and activity limit. The one-step inequalities arise by taking limits in the paid finite-trajectory comparison; the nearby finite trajectories are constructed by (FCA7) and (FCB5), with positive comparison weights and cumulative bounds (FCB1)–(FCB3); the fixed-scalar comparison identifies their activity limits. This is the precise contraction used to upgrade the subsequence construction above.

8.33 Extracting stationary path coefficients and their higher jets

The path book retains stationary coefficients, their higher jets and the true supports of the corresponding rows. Absolute convergence at a smaller analytic width permits termwise differentiation. Insertion transfer then consumes those same coefficient comparisons, so that changing the path presentation does not amount to changing the represented local insertion.

What the path extraction does and does not select The already established contraction fixes the banked activity limits. The path coefficients and their higher jets are additional presentation records and may require compact extraction. Their selected limits are used to represent those same activities and insertions. The extraction does not change the Wilson family or turn its already unique banked limit into a choice between different physical values. Conversely, uniqueness of the banked reader alone does not assert convergence of every unreduced path presentation.

There is a separate extraction behind the stationary paths just used. It concerns numerical path coefficients, not completed activity presentations. Combine all repeated occurrences of the same transported-curvature slot in each row by adding their coefficients. Do the same for each color component of each ordered pattern of two, three or four slots in the loop, cubic and quartic atoms. Denote this countable coefficient array by $a_{N,m}(\kappa)$, where N is the upper horizon and m the fixed level. The length $\ell(\kappa)$ is the slot length or the sum of the pattern's slot lengths. There are finitely many keys of bounded length: a bounded-length walk has a bounded number of steps from a finite alphabet, and there are only finitely many plaquette and color types.

The path bounds give, for each row or color component separately,

$$\sum_{\kappa} |a_{N,m}(\kappa)| e^{\sigma_m \ell(\kappa)} \le C_m \quad \text{eventually in } N. \tag{SH1}$$

Here C_m is the row allowance, the loop allowance for two-slot atoms, or the common higher-atom allowance for three- and four-slot atoms, as appropriate. An individual coefficient is bounded by C_m, because the exponential weight is at least one. Enlarge these bounds to include the finitely many exceptional N's. The product of the resulting compact intervals, over the countable set of pairs (m, κ), is compact and metrizable. Every escaping horizon sequence consequently has one subsequence on which all these coefficients converge simultaneously.

Write its limit $a_m(\kappa)$. For each finite set of keys, pass the inequality (SH1) to the limit. Taking the supremum over finite sets gives (SH1) for a_m at the original width σ_m. More is true at every strictly smaller width $\tau < \sigma_m$:

$$\sum_{\kappa} |a_{N,m}(\kappa) - a_m(\kappa)| e^{\tau \ell(\kappa)} \longrightarrow 0. \tag{SH2}$$

Indeed the part with length at most M is a finite sum tending to zero, whereas the remaining part is at most

$$2C_m e^{-(\sigma_m - \tau)M}. \tag{SH3}$$

Choose M first and then the horizon. This order proves (SH2) and permits pairing the coefficients with every function bounded by $De^{\tau \ell}$. It also preserves the zero and first loop moments, since each displacement coordinate is bounded by path length and a polynomial in length is bounded by an exponential of any positive width. The row symbol at zero is such a pairing with a bounded indicator of the plaquette type. Its limit preserves the finite metric identity. Color invariance passes by the finite sum over color coordinates. Reconstructing the rows and atoms from this common coefficient array therefore retains all the original path hypotheses and allowances.

The convergence includes derivatives of the unmasked physical functions at every order. To justify this assertion, complexify matrix entries around a physical configuration and use inverse links as independent conjugate-entry coordinates. The extension of a slot of length ℓ is a holomorphic polynomial bounded on a sufficiently small common ball by $Ce^{\tau \ell}$. A product of q slots has the same bound with C^q and the sum of their lengths. Equations (SH2)–(SH3) give uniform convergence of the row and atom series on that ball. Cauchy's integral formula on a smaller ball gives convergence of each derivative, with its fixed radius loss and factorial. Restricting to real smooth parameter

curves then gives convergence of all real Fréchet jets. This argument applies to the actual group words; no finite cutoff of the limiting path is substituted for them.

For completeness, the stationary-graph derivative argument used to pass through the nonlinear minimization is as follows. Let $P_N(x,y)$ be the finite profile, let P be its limit, and let $y = u_N(x)$, $y = u(x)$ be their analytic stationary branches through zero. Put $F_N = \partial_y P_N$, $F = \partial_y P$, and

$$H_N = D_y F_N(0,0), \qquad H = D_y F(0,0). \tag{SH4}$$

The preceding all-order convergence implies convergence of every jet of F_N, in particular $H_N \to H$. The stationary gap makes these vertical Hessians invertible. Inversion is continuous near the invertible H, so $H_N^{-1} \to H^{-1}$. This is the inverse needed here, acting from vertical covectors to vertical vectors.

Differentiate $F_N(x, u_N(x)) = 0$ repeatedly. With $g_N(x) = (x, u_N(x))$, the exact diagonal chain rule is a finite sum over partitions π of the derivative slots:

$$0 = \sum_{\pi} D^{|\pi|} F_N(0)[\, D^{|B|} g_N(0)[v^{|B|}] : B \in \pi\,]. \tag{SH5}$$

The single-block term contains $H_N D^k u_N(0)[v^k]$ and the known horizontal derivative $D^k \,\mathrm{id}(0)[v^k]$. Every other block has size less than k. Thus (SH5), solved by H_N^{-1}, proves convergence of the k-th derivative by induction from the lower derivatives. In an ordered-partition convention identical single-block contributions are grouped with their positive multiplicity; dividing by that multiplicity gives the same recursion. Polarization of symmetric real multilinear forms supplies mixed derivatives. The case $k = 0$ is $u_N(0) = u(0) = 0$. Composing the convergent graph jets with the profile jets proves convergence of all stationary value jets.

There is also a precise cancellation at the first unmatched profile order. On a stationary graph $g = (\mathrm{id}, u)$,

$$DP(0)[(0,w)] = 0, \qquad D^2 P(0)[Dg(0)v, (0,w)] = 0. \tag{SH6}$$

The first equation is stationarity at zero; differentiate stationarity once to obtain the second, and use Hessian symmetry for the reversed order. Suppose P, Q agree through order $k+2$. Their stationary graph jets agree through order $k+1$, by the recursion above. In the chain rule at order $k+3$, any term involving a graph derivative of order at least $k+2$ has either one block or two blocks with the other block of size one. These are exactly the terms killed by (SH6). All remaining graph derivatives match. Moreover the difference $P - Q$ has zero jets through order $k+2$, so its composition at order $k+3$ has only the all-singleton partition. Consequently

$$D^{k+3}(P \circ g - Q \circ h)(0)[v^{k+3}] = D^{k+3}(P - Q)(0)[(Dg(0)v)^{k+3}]. \tag{SH7}$$

This is the stationary value identity used for the cubic and quartic response comparisons ($k = 0, 1$); higher unknown graph derivatives have not been discarded without compensation.

In the lattice profile, the linear map in (SH7) is the derivative of the actual fine logarithmic inclusion of the slice graph. Its derivative is the fine stretching map applied to $(v, Du(0)v)$. A different tree-gauge representative changes it by an exact lattice gradient. Gauge invariance, together with vanishing of all lower jets of the profile difference, makes its first nonzero derivative insensitive to such a gradient: differentiate the gauge identity and observe that every correction involving a nonlinear gauge derivative multiplies a lower, zero profile jet. Hence the stationary matrix response and its tree-extracted version give the same first unmatched value. Passing an eventual one-sided bound through the convergent branch derivatives proves the corresponding bound for the limiting coarse profile. Interchanging the two profiles gives the reverse bound. This is the response

comparison used together with the decay and base matching estimates in the stationary-history construction.

The one-loop profile also has a literal decomposition before passage to the limit:

$$\left[\text{loop action} + \tfrac{1}{2}\log\det(\partial_y^2 P_N) - \log \prod_{e \text{ off tree}} J(y_e)\right] \circ g_N. \tag{SH8}$$

The determinant is nonzero at zero by the same positive vertical Hessian. Determinant is polynomial in its entries, logarithm is analytic near this positive value, and the Haar logarithms are analytic in the common chart. The jet convergence just proved therefore applies to (SH8), including its second jet. A finite inequality comparing this second jet to the adjacent draft (the coupling coefficient times the quadratic action plus its loop action) passes to the limit by continuity of absolute value. Field-independent scalar offsets are not identified by this argument.

Finally the extra physical covariance and response rows have their own quantitative, full-sequence convergence. In their weighted row and column norms the finite suppliers give, for $i < N \le M$,

$$\|U_{N,i} - U_{M,i}\|_{\mathrm{row}},\ \|U_{N,i} - U_{M,i}\|_{\mathrm{col}} \le B\ \mathrm{clock}_k(i+1)^r \left(\frac{a_N}{a_{i+1}}\right)^c, \qquad c > 0. \tag{SH9}$$

For the covariance row the distance weight has exponent σ/L_i at unit output scale; for the response row it has exponent σ at output scale L_i. Since $a_N \to 0$, (SH9) gives an entrywise Cauchy limit. Fix a finite set in a weighted row or column, pass its absolute sum to the limit $M \to \infty$, and then take the supremum over finite sets. This retains (SH9) with $U_{\infty,i}$ in place of $U_{M,i}$, in both norms and at precisely these scales. The limits thus supply the physical rows on the same common witnesses, with quantitative tails; they are not an additional independently chosen stationary subsequence.

There is a further full-sequence assertion for the actual metric rows, which we separate from book extraction. At every fixed torus, the unamended diagonal matrix has the limit given by flat quadratic convergence. Its inverse matrices have entry bounds from one common positive strip allowance. A finite-dimensional compactness argument shows that the limiting diagonal matrix is invertible: extract any convergent subsequence of the bounded inverses, pass the identity $D_N D_N^{-1} = 1$ to the limit, and obtain a right inverse of D. A square matrix with a right inverse is invertible. Continuity of inversion then proves convergence of the entire sequence of inverses, independently of the extracted subsequence.

For the actual amended rows the prescribed construction, not merely agreement on curls, gives at that torus

$$M_N^{-1} = D_N^{-1} + u(d_2 D_N^{-1})^{\mathsf{T}}(d_2 D_N^{-1}), \tag{SH10}$$

where u is the fixed amendment coefficient at this level. All operations are continuous finite matrix operations. Thus M_N^{-1} converges. Its inverse M_N converges by the same bounded-inverse argument, using the metric's original strip allowance. This explains why the amendment must be retained from the same construction; agreement of potential Hessians alone would not imply (SH10).

To recover cover coefficients from these all-period limits, take two possible coefficient limits R, R' of a uniformly exponentially bounded row. Their folded kernels agree on every scheduled torus. Fix a cover displacement z and choose a torus whose centered box contains z. The folded coefficient at its residue is $R(z)$ plus the other congruent coefficients. Their sum is bounded by a constant times the exponential tail outside that box, uniformly in the history index. As the scheduled periods grow, this tail tends to zero, hence $R(z) = R'(z)$. Compact coefficient extraction therefore has only one possible limit; the full row converges coefficientwise. Splitting into a finite

displacement set and its exponentially bounded complement proves convergence in every strictly smaller strip norm, exactly as in (SH2)–(SH3). The argument applies to the diagonal, its dual and the amended metric. The physical covariance and response rows have in addition the rate (SH9).

The extraction (SH1)–(SH3) for nonlinear path books alone remains subsequential. The auxiliary response comparison is used through exactly the orders it supplies: base matching through the second jet, then cubic and quartic comparisons (SH7), and the loop second-jet comparison (SH8). In particular any two extractions share the second profile jet, since it is the same quadratic matrix limit just proved. The recursive first-unmatched-jet argument compares the next two jets with their stated bounds. It does not assert equality of arbitrary all-order path books or field-independent scalar offsets. Whenever the subsequent primitive-continuity argument uses common path coefficients, its two sequences are taken from the already selected common book refinement. Full quadratic-row convergence and retained nonlinear-book convergence thus have different, explicit scopes.

8.34 From retained coefficients to physical row comparisons

At a fixed step the retained-history comparison data consist of the limits of the upper and lower path books, hence of their quadratic/cubic/quartic atom coefficients, along the single common horizon refinement selected in (SH1)–(SH3). They also include the three cover-kernel limits and the coupling and activity limits. In this paragraph, two escaping subsequences mean two subsequences of that retained horizon refinement. Both therefore have the same book limits; this is not an assertion that arbitrary subsequential path books extracted before the refinement are equal. The coupling comparison is passed to the same tail ultrafilter, while the retained activity comparisons supply the regular-occurrence and forced-coefficient evaluations at arbitrary varying fine configurations. This constructs the comparison data for the local formulas. The full fixed-volume limit of the Wilson smeared integrals is a separate conclusion proved in “Extending the full fixed-volume mesh limit”; it is not obtained by silently declaring path-book compactness to be full convergence.

The cutoff makes each physical linear expression finite. There are only finitely many canonical path labels of length at most k: their plane type and integer displacement lie in a finite coordinate box. In shortening with length multiplier $d > 0$, every retained label has length at most k. Thus the cutoff shortening is a finite sum of its compared scalar coefficients times the corresponding transported curvature slots. The direct plotted branch is the fixed curvature function and needs no coefficient limit.

Likewise a coarse routed kernel at cutoff k uses only cover displacements of coordinate length at most k. For a fine-to-coarse route from a lifted residue to Lz, route length is $|Lz - \mathrm{lift}(r)|_1$. Since $|\mathrm{lift}(r)|_1 \le 4L$, the retained displacements lie in the box of radius $k + 4L$. Its endpoint and cutoff predicates are the same in both histories. Each nonzero summand is its compared kernel coefficient times adjoint transport along the same finite route. Smooth finite word multiplication, inverse and adjoint action therefore compare these summands and their derivatives. Summing proves comparison of all three physical kernels.

Coarse vision is a finite contraction of these kernels with the compared plotted coarse fields. The linear field on a logarithmic fine direction is either the fixed discrete curl, a finite cutoff path sum of such curls with the compared path coefficients, or the fixed coarse noise linear map. An oriented row additionally applies the common signed displacement rule and carrier adjoint transport. A reversed link uses the negative adjoint-transformed displacement; its base shift is the one already fixed by the oriented link convention. These operations are smooth finite sums and products and hence preserve comparison.

The nonlinear energy rows are then their original finite scalar expressions: the direct clipped

fine part plus the two coarse-response half-squares and cross terms; the wire curvature deficit and effective noise square; the negative coarse curvature deficit together with the two departure terms; and the tight-response half-square. The clipped or kept fine tests act on the same fine configuration in the two histories. Their varying numerical values can therefore be retained as bounded common factors, without taking a continuity limit through the test. The inherited regular source uses the regular-occurrence comparison directly; the atom and path-tail sources use their retained scalar coefficients and the same actual masked words. These prove the physical row and inherited-source comparisons feeding the packet and mean calculations above. The underlying retained stationary and activity convergence is the history input, whose extraction must use the common choices described in the construction.

8.35 Mean, local source and chain expressions

The bill vector comparison can be made explicit. A selected row's strain is its free residual vision plus the physical row operator applied to the saved log field and the hot part of the varying fluctuation. Selection and hot-coordinate sets are fixed by the draft. The saved log is the same all-argument local function on both sides; its uniform derivative bounds give its comparison. On hot coordinates use the compared fluctuation; on dormant coordinates its lens contribution is identically zero. Physical row comparison therefore compares the strain.

For the fine free residual, compare the actual base-stroke map and subtract the compared coarse vision. For a noise residual, the base-stroke part is common; its derivative bounds hold in the clean hull ensured by viability and selection. Adding and subtracting the same compared vision then gives the residual comparison. Thus the clean-hull condition is used exactly on selected noise rows, rather than postulating clean fields everywhere.

The bill vector is the transpose form applied to the selected zero-fluctuation strains, with their original scaling. Finite matrix-vector multiplication compares it. The protected mean is minus the transpose window-column matrix applied to this bill; hence it compares as asserted above. Its scaled stationary defect is the plate times this mean plus the bill, multiplied by $-\sqrt{\beta}$, and therefore compares as well. The positive common coupling limit is guaranteed by the lower coupling-window endpoint; taking square roots and reciprocal scales is consequently valid.

The chain terms are now finite expressions in compared matrices and vectors. Their successive factors are the appropriate negative plate–covariance defect while the index is within the prescribed itinerary, and the plate factor at its terminal position. Matrix products compare these factors, with their order unchanged. The quadratic fluctuation contraction and the linear stationary-defect contraction compare by finite coordinate sums; the local cone cutoffs multiplying them are smooth functions of squared fluctuation norm. The blank-reference chain uses the empty store and identity output field, whose viability is automatic. Its plate has the same common lower bound, so the same inverse and determinant comparisons apply. The condensation source is one half the difference of the blank and actual local determinant expressions.

The Haar source is smooth for all fluctuation arguments after its cutoff. Inside the cutoff support its argument lies strictly within the positive Haar chart radius, so the chart logarithm is smooth. Outside that support the source is locally zero. The strict chart margin covers the cutoff boundary and proves the all-argument claim. This justifies its use in products without extending an uncut logarithm beyond its domain.

Finally the packet logarithm is the original finite combination

$$-\beta(\mathrm{chalky}+\mathrm{crispy})+\sum_{\mathrm{visited}}\mathrm{condensation}+\sum_{\mathrm{plaza}}\mathrm{HaarSource}-|\mathrm{pens}|\log(\mathrm{chartMass}(\beta)).$$

Here chalky is the posted root energy minus its selected quadratic costs; crispy is the selected quadratic cost at the protected mean. Each cost is a finite sum of squared strains over the actual parent rows. The root-energy comparison is the corresponding nonlinear energy comparison minus that same cost. The chart mass is positive and continuous at the positive coupling limit, so its logarithm compares. All posted, visited, plaza and pen sets are fixed finite sets in the draft. This proves the packet-log comparison from the strain, root-energy and determinant comparisons, with the original normalization retained.

8.36 The four vertex comparisons

After fixing the draft, the ordinary integrand has the same four vertex types as its absolute activity bound. The pointwise comparison treats their predicates explicitly.

For a connected block packet, connectedness and the containment/empty-flag test are fixed discrete predicates. If a predicate fails the corresponding factor is identically zero in both histories. Otherwise the packet is the exponential of its compared energy expression, multiplied by the common contract factor, tickets and profiles, and the compared smooth local cutoff on its actual enlarged plaza. The fine-configuration ticket and profile factors may vary with the history index but are the same on both sides and uniformly bounded. Their output-field dependence uses the already chosen smooth wire profile. Thus the finite product comparison applies without requiring a limit of the ticket indicator.

For an old forced vertex, compare its actual forced coefficient evaluated on the same fine configuration, as supplied by the retained forced-activity comparison, and multiply by the compared cutoff on its true footprint. For a regular exponential vertex, compare its logarithmic expression first and then apply the smooth exponential and subtract one. Its branches are inherited regular terms, the short local energy source, the undercurrent condensation term, the Haar term and the chain-row terms. The tests for being inside the posted region, for voyage membership and for hot exterior coordinates are fixed by the draft; their zero branches agree exactly. The nonzero branches use the protected mean, covariance/window, local logarithm and row comparisons already needed for these expressions, with the actual signs and factors of β retained.

For a warning vertex, its mandatory/hot-region predicate is also fixed. The factor is a compared cutoff times the actual fine profile or contract charge minus its constant reference value, or an exterior cone minus one. Fine profiles and finite label charges are bounded common factors; the exterior cone is a smooth function of the compared fluctuation direction. Its mandatory branch is identically the stated constant in both histories. Thus all warning comparisons retain their original hard predicates.

Finite products over the molecule give the integrand comparison, and multiplication by the compared density gives the transformed integrand comparison. The derivative version uses the same four factors in their local smooth extended formulas, the same fixed fine configuration, and the compared steered fluctuation and corrected Gaussian coordinates. Its full finite product is identified with the original differentiated integrand by the exact local formula on a neighborhood of the healthy base. These statements compare the true vertices rather than replacing hard indicators by smooth approximations. The underlying energy, row, cutoff and Haar comparisons are the finite-coordinate ingredients of these four cases.

8.37 The marked primitive comparison

The marked premium contains a test that its distinguished badge occurs in the molecule. The badge is fixed by the common mark key, so this test agrees in the two histories. If it fails, both premiums

vanish. If it holds, the integrand replaces precisely that vertex by the paid marked ingredient and retains the product over the molecule with that badge erased.

At fixed physical fluctuation and common draft, the reconstructed fine configuration is the same on both sides. The ingredient comparison is therefore the comparison of the two mark entries evaluated on this common, possibly history-dependent, fine configuration. Multiply it by the common finite enforcement cutoff, the compared ordinary product over the erased-badge molecule, and the compared Gaussian density. This proves the pointwise comparison of the marked amplitude. The exact reconstruction identity between the paid ingredient and this expression uses viability; it is applied on the eventual viable set.

For the derivative integrand the marked entry remains evaluated on the fixed reconstructed fine input. Output-field derivatives act on its smooth steered cutoff and on the remaining ordinary vertices, as prescribed by the local reweighting formula. The local chart identity identifies the actual marked ingredient with this expression near the healthy base. The compared steering map gives comparison of the cutoff; the four ordinary vertex comparisons give comparison of the remaining product; the local Gaussian ratio and Jacobian give comparison of the reweighting factor. Their product is exactly the derivative integrand, not a derivative of a newly substituted marked function. This explains why convergence of mark entries under varying-configuration evaluation is sufficient here.

The same Gaussian-polynomial majorant bounds these marked derivatives and permits the integral passage already proved. Finite boundary and forest averages then compare the actual marked grove. Its all-field bound also gives convergence of its uncut amplitude under arbitrary varying output configurations, using the fixed-draft argument separately at each gauge-slice point. This supplies both the regular marked factor in the one-mark cluster stream and the forced marked comparison used in the later reset.

8.38 Uniform derivative comparison of a single cluster stream

A one-mark connected stream with r occurrences is its fixed Ursell coefficient divided by $r!$, multiplied by

$$\sum_{\nu=1}^{r} \ell_\nu \prod_{\mu\neq\nu} u_\mu,$$

where u_μ is an ordinary primitive entry and ℓ_ν its marked replacement. Its actual support is the occurrence union when the mark speaks to that label, and its zero alternative otherwise. The speaking condition depends on the fixed mark key and label inventory, so corresponding histories have exactly the same support. A silent stream is identically zero.

The stream's derivative bill is summable uniformly in the history index. To see the essential bound, let a_μ, b_μ bound ordinary and marked entries through derivative order j. The ordered product rule bounds the term with mark ν by

$$(1+r)^j b_\nu \prod_{\mu\neq\nu} a_\mu.$$

Replace a_μ, b_μ by $2^j a_\mu, 2^j b_\mu$. Since $1+r \le 2^r$, this absorbs that derivative factor. The ordinary bounds may be taken as the smaller of their uniform primitive envelope and their original purchased budget; taking the smaller preserves summable tails. A positive reserve d, selected by the one-mark budget argument, makes the inflated array $f_\mu = 2^j a_\mu + d\,2^j b_\mu$ fit the connected-cluster reserve.

Expanding its product with nonnegative entries gives

$$d\sum_{\nu}\left(\prod_{\mu\neq\nu}2^{j}a_{\mu}\right)2^{j}b_{\nu}\leq\prod_{\mu}f_{\mu}.$$

Thus a valid stream bill is d^{-1} times the ordinary absolute connected-cluster weight for f. The rooted cluster estimate makes this family uniformly summable with small tails. Multiplication by the fixed Ursell factor and $r!^{-1}$ is retained exactly.

For a fixed speaking stream, a healthy union makes every occurrence support healthy. The primitive comparison along any allowed configuration family then compares all its ordinary factors and its single marked factor through derivative order j. Finite product and sum differentiation compares the stream itself. For a silent stream the assertion is immediate. Passing through every ultrafilter above the tail filter turns this into ordinary convergence of the derivatives of any fixed-order allowed test family.

Here is why such test-family convergence gives a uniform paid derivative bound, including on the closed healthy set. If an order-k derivative bound failed uniformly over the allowed maps from $\mathbb{R}^{k}$ and healthy base configurations, select a violating map at each bad history index. The assumed convergence for that sequence of maps contradicts the violations. Take the finite intersection for $0\leq k\leq j$. This yields a single late cutoff for all these orders.

To extend the test to an arbitrary finite-dimensional parameter space and directions $v_{1},\dots,v_{k}$, define

$$Lz=\frac{1}{k+1}\sum_{\nu=1}^{k}z_{\nu}\,\widehat{v}_{\nu},\qquad\widehat{v}_{\nu}=\begin{cases}v_{\nu}/\|v_{\nu}\|, & v_{\nu}\neq0,\\ 0, & v_{\nu}=0.\end{cases}$$

With the coordinate maximum norm, $\|L\|\leq1$. The vectors $a_{\nu}=(k+1)\|v_{\nu}\|e_{\nu}$ satisfy $La_{\nu}=v_{\nu}$. Pulling an allowed map back by $z\mapsto Lz+p$ preserves its allowed derivative bounds. A test bound $b/(k+1)^{k}$ then bounds its derivative on the a_{ν} by $b\prod_{\nu}\|v_{\nu}\|$, proving the full operator bound b. Smoothness for this subtraction is furnished by the separate primitive derivative bills at the healthy base; the mask itself is never differentiated.

Finally convert these arbitrary-map derivative bounds to the anchored paid direction convention. At a fixed period $1+\operatorname{dist}(x,y)$ is bounded by one finite Y. The conversion costs at most $Y^{p.\mathrm{power}\,j}$, independent of the history index. Divide the desired tolerance by this constant before applying the uniform comparison. The identity configuration makes the healthy test set nonempty. This proves the regular comparison for a stream at its actual fixed support. Real and imaginary streams are treated separately, multiplied by their common normalization and recombined, then embedded by zero into the common enumeration. This supplies the microscopic regular estimate used above.

8.39 Passing microscopic comparisons through projection

The microscopic regular comparison means that both arrays are regular on one fixed family of supports, have uniformly summable paid bounds with uniformly small tails, and, at each fixed label, their difference eventually has paid derivative bound below every positive tolerance. The forced comparison means both arrays have absolute bounds with uniformly small summable tails and, for each fixed label and every sequence of configurations U_{n}, the difference of their values at U_{n} tends to zero. The latter implies eventual uniform smallness over configurations: otherwise choose a violating configuration on each bad index to contradict that sequence test. Neither notion assumes continuity of the masks.

The tail hypothesis gives a useful dominated-series principle. If two arrays have uniformly summable absolute tails and their fixed-label differences tend to zero, split the sum into a finite set

and its complement. The finite difference sum tends to zero, and the two tails have an arbitrarily small uniform total. Hence the difference of the full series tends to zero. The same proof permits evaluation at arbitrary varying configurations. Forced comparisons are therefore preserved by summing a series, by collecting labels into finitely many barcodes, by a fixed injective reindexing with zero outside its image, and by finite sums and scalar multiples. They are also preserved by products: each fixed label is uniformly bounded, and $fg - f'g' = (f - f')g + f'(g - g')$; a global bound on one absolute-budget family times the other's tail budget controls the product tails.

Let Π denote the fixed weighted local section projection, with projection price C_{read}. The core of the collected raw insertion is its seed plus the sum of Π applied to its bulk terms. For regular germs this projection is linear, including subtraction. Its paid bound gives

$$|\Pi(f - g)|_1 \le C_{\mathrm{read}}\,\mathrm{paidBound}(f - g).$$

At each label the right side tends to zero, and the projected rows have uniformly summable tails by that same price. Thus the projected core difference tends to zero whenever the seed difference does.

The residual and rebate operations used to remove this core are likewise linear on regular germs at the common threshold and support. Their fixed paid conversion price bounds the residual in the enlarged support and bounds the rebate everywhere. Apply these operations to the difference of the two bulk terms. Arbitrarily small paid difference bounds therefore give arbitrarily small residual paid bounds and arbitrarily small all-field rebate differences. Their uniform tail budgets are multiplied only by the fixed conversion price. Placing residuals in the common canonical support enumeration, with zeros at unused indices, preserves the regular comparison.

The new forced profile is the old forced profile plus the threshold hook of the collected rebates. The hook is a finite sum over support classes, multiplied by the actual compatible products of forced activities. These products agree in their discrete support and flag predicates because the two representations have the same forced address data. Their numerical factors have the forced comparison by the assumed comparison of the original forced activities, finite-product stability, and summable product budgets. Multiplying by the compared rebates and summing therefore preserves the forced comparison. Finally the common keeper barcode embedding collects these profiles and extends them by zero outside its image; this preserves both the convergence and tail bounds. Thus projection and collection pass the microscopic regular and forced comparisons to the output formatted insertion.

8.40 The common input decomposition for a shot

For fixed input capture x and frame index i, a shot is the merged sum of three kinds of contributions: the primary frame insertion, one contribution for each regular remainder label, and one for each forced barcode label. The input is the same book in both histories. The common incoming orientation fixes its threshold, forced addresses and barcode profiles. Stowing the forced input through the keeper therefore gives one common function of the fine configuration for every fixed forced label, independent of the history index and of which of the two histories is used.

Each of these contributions has a summable common tariff. For the regular inputs choose their paid regular bounds v_j, and use tariff $\psi\, p.\mathrm{mass}(D_j, v_j)$. For forced inputs choose all-configuration bounds c_j, and use $\psi\, p.\mathrm{mass}(K_j, c_j)/\kappa_1$, with the same positive incoming forced scale κ_1. The input capture's wealth bound makes both arrays summable. The finite primary contribution uses the larger of its heading tariff and its heading-plus-core tariff. The paid shot estimates bound not only each remainder but also its output core coefficient by this array. Thus it is a valid majorant for the countable comparison argument.

For a regular label the actual input ingredient is the real or imaginary part of its fixed masked regular function; the preparation construction has exactly that ingredient, its actual logarithmic key, and weight at most v_j. For a forced label the ingredient is the real or imaginary part of the common stowed function and has its actual gas key. If that function is zero, the output contribution is identically zero in both histories. Otherwise its input bounds and eligibility allow the same preparation construction, with weight at most c_j. These are genuine fixed ingredients across the two histories, although their auxiliary preparation data may vary.

The nonconstant primary frame term likewise uses the fixed masked section insertion at the common incoming threshold, split into real and imaginary parts. For the constant frame index, preparation is unnecessary: the output is the same scaled unit insertion in both histories. Its forced profile is zero, and stowing that zero profile through any keeper leaves it zero; this proves equality even if the two keepers have different presentations.

We now establish the microscopic comparison of these preparations. Fix the scale, its block ratio L, the degree d, and the incoming dictionary. Both histories have that same dictionary and budget at this scale. Their couplings approach each other, their regular and forced activity coefficients have the weighted comparison already proved, and their stationary kernels and path responses have the common limits. Thus their local Gaussian covariances, means, saved fields and vertex factors have precisely the comparisons established earlier in this chapter.

A preparation of a fixed real ingredient retains its wing type and a common bound on its input weight. Converting the preparation to the common local chart preserves its ingredient, wing and weight. Consequently the two converted real ingredients agree exactly, as do the two imaginary ingredients. This is stronger than convergence of a possibly changing representative. The support address of a local integrated term depends on the wing and its fixed combinatorial fragrance address. It therefore agrees in the two histories and throughout the comparison sequence. Denote this support by E_α. The insertion format ensures $0 \in E_\alpha$.

Apply the mixed marked-integral comparison to each such address. Its hypotheses are now the identical ingredient just identified, the compared Gaussian data and vertices, and their common summable derivative envelopes. The change to a fixed Gaussian integration variable, the two dominated-convergence arguments, and the finite-order derivative estimate give: for every address α and every $\epsilon > 0$, the difference of its two real bulk terms eventually has paid bound ϵ, on E_α, at the common output threshold and derivative order. The same statement holds for the imaginary bulk terms. Their separate envelope arrays have the uniform tails of the marked cluster majorant. Multiplying by the fixed block factor L^d and recombining real and imaginary parts preserves these conclusions. Extending the address enumeration by zero preserves them as well. We have therefore proved the bulk comparison, including both regularity on the common support and the summable tails needed for projection.

The forced part also retains its actual address data. Its construction first restores the missing ordinary logarithms, then restores the marked slots, and finally applies the coupling reset to the common shelves. For the ordinary logarithms, the finite-set/tail comparison of the reset families proves convergence and common tails. The marked slots have the mixed-integral comparison just obtained; inserting the prescribed finite subsets and extending by zero preserves it. Products of the restored ordinary exponentials with the marked factors preserve convergence and their summable envelopes. Finally the reset comparison applies to the same path family, threshold, shelves and rebate operations on both sides. The forced output profiles consequently have uniformly summable bounds and satisfy

$$F_{n,\alpha}(U_n) - F'_{n,\alpha}(U_n) \longrightarrow 0 \quad \text{for every sequence of configurations } U_n. \tag{HL1}$$

This is the strong all-configuration comparison used here; no continuity of the masks is involved.

For clarity, (HL1) implies the corresponding uniform comparison. Otherwise there are $\epsilon > 0$, a subsequence n_j, and configurations U_{n_j} where the difference exceeds ϵ. Complete these choices to a sequence by arbitrary configurations at the other indices. It contradicts (HL1). Thus the supremum of the difference tends to zero. Combined with the uniform tails, the finite-set/tail argument also proves convergence of each summed forced profile.

The resulting real and imaginary forced comparisons recombine as

$$L^d(F_n + iG_n) - L^d(F'_n + iG'_n) = L^d\big[(F_n - F'_n) + i(G_n - G'_n)\big],$$

with the same fixed L and d. This proves the required forced comparison for the raw complex preparation. Its primary seed has the already compared coefficient. The linear projection estimate, residual and rebate bounds, and common keeper extension from the preceding section now give convergence of the projected core and arbitrarily small bound on the collected remainder. The countable ingredient summation and finite-frame assembly therefore use proved local comparisons, all on the original common input book.

8.41 Summing local insertion comparisons

The local comparison feeding the transfer metric preserves the same coefficient supports in the two histories. It consists of convergence of their finite core coefficients to each other and, for every $\epsilon > 0$, an eventual remainder bound ϵ for their difference in the common output orientation. The difference is formed by appending the negative of the second labelled remainder to the first, so it remains a legitimate formatted insertion on that common support.

Here is the summation principle needed for countably many local comparisons. Suppose the i-th pair of insertions each has a common summable bound $f_i \geq 0$, and its difference has a remainder bound tending to zero with the history index. The coarse bound on that difference is $2f_i$. Given $r > 0$, choose a finite set E with $\sum_{i \notin E} 2f_i \leq r/2$. On the finitely many indices in E, eventually choose the common bound $v = r/[2(|E|+1)]$. Use v on E and $2f_i$ outside. This is a valid per-index bound and its total is at most r. The paid merge bound therefore gives an eventual remainder bound r for the full merged difference. Absolute convergence and the same label enumeration identify this merged difference with the difference of the two merged insertions. The finite core sums converge by the same finite-set/tail split. Thus countably many local comparisons assemble into the full comparison.

The marks of a fixed coefficient input are obtained by a finite sum over the frame indices and the prescribed motion chips. Their shot inputs are the same captured coordinate-axis data in both histories. The local comparison for each such shot passes through its fixed motion and through this finite sum by the formatted-insertion operations. Hence the resulting site marks have converging core difference and arbitrarily small common remainder bound.

Finally each actual coefficient readout equals that site's core, and its recorded capture equals capture of that site mark under the common output keeper. Each site mark has a finite paid remainder bound. The identity collecting a difference under the keeper identifies the difference of their recorded captures with capture of the formatted difference mark. Its arbitrarily small remainder bound bounds the capture norm by the same small amount. This proves convergence to zero both of the coefficient-isomorphism outputs and of their recorded capture outputs. These are precisely the fixed-input comparisons used in the transfer argument; the preceding microscopic comparison supplies their local analytic input.

8.42 The transfer comparison metric

Let $|x|_1 = \sum_\nu |x_\nu|$ be the coefficient size in the fixed finite frame. A Gram datum P contains a complex-linear capture map g_P into the insertion-data seminormed space, with $\|g_P x\| \le a|x|_1$. The comparison relation $P \sim_\epsilon Q$ means $\epsilon \ge 0$ and

$$\|g_P x - g_Q x\| \le \epsilon a|x|_1 \quad \text{for every } x.$$

It need not separate distinct Gram data. It is symmetric, satisfies the triangle inequality in ϵ, and every pair is comparable with $\epsilon = 2$. For coefficient maps, the corresponding comparison is $|Sx - Tx|_1 \le \epsilon|x|_1$. Both uniform comparisons follow from bounds on the finitely many coordinate axes: expand x in that basis and sum absolute values.

An adjacent transfer consists of a flow $P \mapsto \Phi_n(P)$ on Gram data and a coefficient isomorphism $S_n(P)$. Its established paid inequalities are

$$|x|_1 \le b|S_n(P)x|_1, \quad |(S_n(P) - S_n(Q))x|_1 \le c\epsilon|x|_1, \quad \Phi_n(P) \sim_{\epsilon/2} \Phi_n(Q)$$

whenever $P \sim_\epsilon Q$. At fixed input P, both the flow outputs and the maps $S_n(P)$ are Cauchy in their respective comparisons. These statements concern the actual integrated transfer maps; outside the eventual admitted schedule they may be completed by the zero capture map and the identity coefficient map. This harmless completion supplies total sequences while preserving all eventual statements.

To recover flow comparison from coefficient-normalized quantities, suppose $g_X V = T$, $g_Z W = U$, where V, W are coefficient isomorphisms, their difference has comparison cost c_0, the maps T, U differ by at most $c_0|x|_1$, and $|x|_1 \le b|Vx|_1$. Put $k = V^{-1}i$. Then

$$g_X i - g_Z i = (Tk - Uk) + g_Z(Wk - i).$$

The two terms are bounded by $c_0 b|i|_1$ and $ac_0 b|i|_1$. Thus $X \sim_\ell Z$ whenever $(1 + a)c_0 b \le \ell a$. This proves the normalized comparison without an inverse estimate on an untested ambient space.

Pointwise convergence of the normalized maps along every pair of escaping history subsequences implies the uniform Cauchy comparisons. First apply it to each of the finitely many coefficient axes, obtaining a common bound and hence a uniform map bound. Next use the preceding inverse-normalization estimate for the flow. Finally, if an eventual pairwise tail bound failed, one could choose bad indices $u(n), v(n) \ge n$. Both escape to infinity and contradict the assumed subsequence convergence. This turns the subsequence statement into an actual double-tail Cauchy estimate.

If inputs P_n themselves are Cauchy, the flow outputs are Cauchy too: insert one fixed late input P_N between $\Phi_n(P_n)$ and $\Phi_m(P_m)$. The two outer errors are halved by contraction and the middle error is small by fixed-input convergence. Similarly insert P_N in the coefficient maps; their two outer errors cost c times the input error. Taking the input tolerance to be $\epsilon/[4(c + 1)]$ and the middle tolerance $\epsilon/2$ proves the uniform coefficient-map comparison. These establish the precise transfer properties used in the finite-depth approximation and coefficient Cauchy argument below.

8.43 Fixed-level convergence of transported insertions

We explain the compactness mechanism for insertion coefficients before taking the continuum limit. At a fixed physical period, let $W_{i,n}$ be the integrated Gram data at level i in the n-th retained history. Adjacent levels satisfy, eventually in n,

$$W_{i,n} = T_{i,n}.\mathrm{flow}(W_{i+1,n}).$$

The paid transfer bounds make each flow contract the comparison error by one half and carry the two-part Cauchy property of its input to its output. All admitted Gram data at one level have mutual comparison error at most two. Iterating d levels and replacing the last input by fixed admissible data therefore approximates $W_{i,n}$, eventually, with error at most $2\,2^{-d}$. For fixed d the approximating sequence has the two-part Cauchy property by the finite transfer bounds. Choose d so the approximation error is below a quarter of the desired tolerance, and then use this Cauchy property. The triangle inequality proves that $W_{i,n}$ itself has that property at every fixed level. Only finitely many eventual identities are used for any such approximation.

Here is the inverse-normalization argument for coefficient transport. In the finite-dimensional coefficient norm, suppose $x_n = S_n y_n$ eventually, x_n is Cauchy, the maps S_n are Cauchy uniformly on bounded inputs, and

$$|v|_{\mathrm{coeff}} \le b\,|S_n v|_{\mathrm{coeff}}$$

with fixed $b > 0$. First x_n is bounded, so y_n is eventually bounded by the same inequality. Next

$$S_n(y_n - y_m) = x_n - x_m + (S_m - S_n)y_m.$$

The first term tends uniformly to zero for large n, m; the second does so by boundedness of y_m and the uniform map comparison. The inverse bound proves y_n Cauchy. Equivalence of the finite-dimensional coefficient norm and the ambient norm converts this to the usual Cauchy statement.

Apply this successively to the actual adjacent transfer relations for the coefficients normalized by the reference inverse. Their lowest retained coefficient is eventually the fixed terminal coefficient y, by the based transfer condition, and is therefore Cauchy. The Gram comparison just proved supplies the Cauchy comparison of each transfer map; its paid inverse bound supplies the uniform b. Induction upwards gives Cauchy coefficient sequences at every fixed level. This uses inverse bounds on the actual graph input at that level, without assuming a bounded product through arbitrarily many future levels.

The corresponding captured insertion data are Cauchy as well. If g_n is the integrated capture map at that level and x_n its Cauchy coefficient, write

$$g_n x_n - g_m x_m = g_n(x_n - x_m) + (g_n - g_m)x_m.$$

The uniform map bound controls the first term, and the second part of the Gram comparison controls the second on the bounded coefficient sequence. This constructs the Cauchy sequence of captured data used for the fixed-level density limit.

In the retained-history construction the terminal horizons tend to infinity. Thus every fixed level eventually lies below the horizon, all coefficients and captured data above refer to the original common history choices, and the path at that level is the prescribed stationary path. The history conditions provide the corresponding uniform lawful, scalar, regular, forced and path budgets. At a fixed level the Cauchy coefficient sequence is bounded; multiplying its bound by the fixed normalization scalar gives the uniform insertion payment bound. The exponential insertion and each forced profile are measurable: they are sums of masked measurable regular terms, followed by continuous pullback and either complex conjugation or identity according to the chosen motion. These are the measurability inputs for the fixed-level integral passage.

8.44 Uniform convergence of fixed-level insertion integrands

At a fixed level the comparison required for integration is uniform over all configurations; no continuity across changing masks is assumed. More precisely, a sequence of functions has the needed property if it is eventually uniformly bounded and is uniformly Cauchy over its entire

argument set. This property is preserved by eventual equality, fixed pullback, finite sums and products, and continuous functions on the finite-dimensional bounded target ranges. For the last assertion the common bounded range lies in a compact ball, on which the continuous function is uniformly continuous.

For captured insertion data x, let $\mathcal{L}_x(U)$ be its summed masked regular remainder and let $\mathcal{F}_{x,B}(U)$ be its forced profile with barcode B. The capture norm and paid local mass bounds give absolute convergence and

$$|\mathcal{L}_x(U)| \le \|x\|, \qquad |\mathcal{F}_{x,B}(U)| \le k\|x\|,$$

uniformly in U. Masks have absolute value at most one; for the regular terms apply the healthy paid bound when the mask is one and use zero otherwise. For the forced profile restrict the summable forced budget to the specified barcode. Both maps are linear in x: absolute convergence permits termwise subtraction, including subtraction under the fixed masks. Thus the same inequalities hold for differences. A Cauchy sequence of captured data consequently gives uniformly Cauchy regular remainders and every fixed forced profile, with eventual uniform bounds.

The finite core insertion is also controlled uniformly in the configuration. For each coefficient-axis vector, its placed core function is continuous on the compact finite lattice configuration space and hence bounded; its healthy mask cannot increase that bound. Expanding a general coefficient in these finitely many axes gives a uniform Lipschitz bound in its coefficient norm. The constant coefficient is handled in the same way. Therefore a Cauchy sequence of core coefficients gives uniformly Cauchy core insertion functions, including their masks.

The format and capture identities identify the actual exponential insertion entry with the sum of this core function and the captured regular remainder, eventually in the history index. They identify the actual forced profile with the captured barcode profile as well. The preceding uniform bounds therefore prove the fixed-level convergence assertions for both actual entries and all profiles. A fixed translation, signed motion, complex conjugation when required, and fixed normalization scalar preserve these assertions. This proves the translated insertion convergence at any fixed anchor from the origin construction, without choosing a new normalization there.

For later use, integration of such a sequence against a probability measure is Cauchy whenever its terms are eventually measurable: the uniform bound gives integrability, and the norm of the difference of two integrals is at most the uniform norm of their difference. Applied to the finite polynomial restore expression, finite-dimensional continuity and finite sums/products preserve the same convergence. The common positive lower bound on its unmarked denominator then permits the normalized quotient limit. This connects the captured-data convergence to the fixed-level normalized expectations.

8.45 Exact unmarked density and component regrouping

Here is the unmarked version of the molecule expansion, including its coverage condition. Fix one forest throughout, one probability Haar measure in its gauge coordinates, and the common output configuration. A packed address consists of a store, a finite warning-label set, and the remaining labelled lists. The map sending this data to its molecule uses the prescribed block layout of that store. Reading the roster, warning ownership and lists from the molecule is a left inverse. Conversely, every molecule satisfying that layout is recovered from the read data. Thus this map is injective and misses only addresses whose activity is zero.

For each store and boundary choice, the local expansion has summable terms. Its absolute majorant is the scalar tribute times a product of nonnegative vertex tariffs. The tariff family is summable; hence its finite-subset products are summable, since their total is bounded by the exponential of the tariff sum. Restricting to the injectively encoded addresses preserves summability.

Integrating the expansion and then summing the finite boundary choices gives the store's exact cut entry.

The scalar

$$O_X = \mathrm{levy}_X \ \mathrm{deducted}_X$$

is unchanged by replacing the store and assigning a boundary choice: the neutrality identity cancels their apparent effect on the scalar correction, with the forest held fixed. If there are b boundary choices, the tribute for each is exactly O_X/b. Their sum is therefore O_X times the averaged molecule receipt. Summing over stores, and using the full discrete sector coverage, yields the absolutely convergent identity

$$\mathrm{push}(R; V) = O_X \sum_p \mathrm{receipt}_X(p).$$

This is a density relative to probability Haar in the output variable; no unit partition normalization has been introduced.

To factor this sum into connected activities, partition each molecule into the connected components of its actual enlarged supports. These components have pairwise disjoint grounds and their union reconstructs the molecule. A compatible family contributes only if its roster flags cover exactly the output flag set, and each component satisfies its own layout, eligibility and flag-read conditions. These conditions are equivalent to the union molecule's layout and approval. For example, intersecting the union flag set with one component's ground leaves precisely that component's flags: flags from another component cannot enter, by disjointness and the fact that a component's flags lie in its ground.

Define a component activity by planting all output flags in its local data and then applying its own eligibility test. On its ground, this planted data agrees with the data from the union roster. The same holds for the selected regular keys: a key belonging to another component cannot meet this ground. Locality of the integrated activity therefore identifies the planted component activity with the union-data activity. Independence on disjoint grounds factors the latter product, including the independent finite boundary averages. Hence

$$\mathrm{receipt}_X\Big(\bigcup_{p\in A} p\Big) = \mathbf{1}_{\{\text{compatibility and exact flag coverage}\}} \prod_{p\in A} a_X(p).$$

If the union approval fails, either flag coverage fails, making the displayed contribution zero, or at least one component fails its eligibility condition and its factor vanishes. Thus this equality includes all zero cases.

The connected-component map is injective because union is its left inverse. Every family outside its image has zero compatible contribution: a nonzero compatible family consists exactly of the components of its union. Reindexing the absolutely convergent sum proves

$$\mathrm{push}(R; V) = O_X \sum_{A \text{ finite}} \mathbf{1}_{\{\text{compatibility and exact flag coverage}\}} \prod_{p\in A} a_X(p).$$

Finally $O_X > 0$: its Gaussian chart mass and kernel normalization factors are positive, and its other factors are exponentials or positive powers of those masses. The extensive scalar stays outside the connected gas. The argument has not asserted locality of that scalar or performed the subsequent coupling reset.

The frozen data in this identity are the store and boundary choice after fixing the gauge slice; the threshold, width and selector belong to the fixed paid chart. The comparison of these expressions uses the common drafts, principal covariance matrices and Gaussian densities constructed earlier in

this chapter, together with the bill and protected-mean comparison and the four vertex cases. For a marked product it uses the same comparison with its finite marker coefficient retained. The next subsection states how these addresswise comparisons pass through the common source, its products and the reset; the signed line estimate (LCN11) gives the corresponding parameter control. No varying discrete approval predicate is differentiated.

8.46 Comparison of the common source and its marked products

Fix a period, a derivative order j, and a budget B of order at most j, with threshold equal to the common temperature. All addresses and all finite trial tori in this paragraph are fixed before either history is varied. A regular comparison means common sparse summable envelopes in the B-seminorms and convergence, for each address, in those seminorms. A forced comparison means common sparse envelopes in the uniform norm and addresswise uniform convergence. These are stronger than pointwise comparison of the total action; the distinction permits the subsequent products and reset operations.

For the connected logarithmic stream, the local Gaussian comparison already proved gives convergence of every derivative through order j, along every sequence of healthy configurations and admissible test maps. This implies uniform convergence in the healthy-configuration derivative seminorm: otherwise choose a violating configuration and test map for each member of a bad subsequence, contradicting that sequential statement. On a fixed torus choose $Y \geq 1 + \operatorname{dist}(x, y)$ for every pair of sites. Conversion from admissible derivatives to anchored B-derivatives costs at most

$$T = Y^{B.\mathrm{power}\, j}.$$

Apply the uniform derivative result with tolerance ϵ/T. The common stream envelopes, multiplied by T, remain sparse and give the required regular comparison. The forced connected stream is handled by the same bad-sequence argument with order zero; its separate sparse envelopes were obtained from the forced integral estimate. In particular, the ultrafilter statements used to express local Gaussian limits here imply ordinary sequential convergence, since every free ultrafilter on every bad subsequence would give the same contradiction.

The path contribution is the sum of four already defined terms: the low path term with coefficient $-v_n$, the cubic and quartic terms with coefficient $-\beta_n$, and the loop quadratic term with coefficient -1. Their coefficient arrays converge in the common weighted path norms. The scalar sequences v_n, β_n in the two histories converge to the same respective limits and are uniformly bounded. For any such scalar pair a_n, a'_n and regular pair F_n, F'_n, use

$$a_n F_n - a'_n F'_n = a_n(F_n - F'_n) + (a_n - a'_n)F'_n.$$

The uniform scalar bound preserves sparse envelopes, while the two terms on the right tend to zero at each address. Thus all four path terms compare. Finite frame averaging preserves this conclusion: transport by a frame costs the fixed factor $\mathrm{turnFee}(B)$, and multiplication by the averaging weight costs $\mathrm{Rim}(c)$. For a target tolerance ϵ, use the original comparison with tolerance $\epsilon/(\mathrm{Rim}(c)\mathrm{turnFee}(B))$. The finite frame index multiplies each sparse envelope only finitely many times.

Consequently the full regular source, the negative connected logarithmic stream plus these path terms and their prescribed frame average, has a regular comparison. The hard source has a forced comparison, including its zero branch and its canonical label embedding. The enumeration used to place these terms in the common dictionary is the same for both histories.

The finite trial correction requires one additional observation. At the fixed period both selection panels are eventually contained in a single finite label set, and their cutoff parameters eventually

exceed every label in that set. Their panels therefore agree eventually. Removing the selected terms from the regular source preserves its sparse envelopes and addresswise convergence. Likewise the selected site contribution is a finite sum over this common set with the same fixed selection weights. Its comparison follows term by term. There is no passage through a discontinuous moving selection rule.

A trial on the fixed small torus is pulled back through its prescribed padded shape map. Its support lies within the fixed radius z, and a healthy neighborhood of that radius maps into the all-healthy trial torus. The derivative pullback costs at most

$$M = 2^{B.\mathrm{order}}(1+z)^{B.\mathrm{power}\,B.\mathrm{order}}.$$

Using tolerance ϵ/M on the trial torus proves comparison after pullback. The ornament is a finite sum over the trial types and sites, with fixed ray weights, so it has a regular comparison. The corresponding failed-trial contributions have a forced comparison by the same finite-sum argument. The site correction is exactly the selected regular contribution minus its value at the identity configuration minus the failed-trial contribution. The identity is healthy; hence evaluating there preserves the regular comparison. Subtraction proves the uniform site comparison, and the finite set of sites gives the forced comparison required for the correction.

These facts assemble the common source: its regular part is the unselected regular source together with the trial ornament; its hard part contains the original hard source and the negative frame average of the site correction. Every sum uses a common disjoint label set. This proves comparison of the complete source, rather than merely of its total evaluated function.

For clarity, the marked-product step can also be written without the source notation. A marked family $u_{n,i}(t,V)$ carries sparse bounds for

$$|u_{n,i}(t,V)| + |\partial_t u_{n,i}(t,V)|$$

and convergence of both quantities along arbitrary configuration sequences, at the prescribed parameter values. An unmarked compared logarithm $\ell_{n,i}(V)$ becomes a constant marked family. Its exponential factor is $e^{\ell_{n,i}} - 1$; on a common bounded interval the mean-value theorem transfers the logarithmic comparison to this factor and preserves the sparse envelopes with the fixed exponential multiplier.

Extend both families by zero through the same label embeddings. A resulting forced atom is a finite product of the embedded marked factors and these unmarked exponential factors, multiplied by its eligibility indicator. The two histories have exactly the same eligibility predicate on the common configuration. Product differentiation,

$$\partial_t \prod_{i\in A} u_i = \sum_{i\in A} (\partial_t u_i) \prod_{k\in A\setminus\{i\}} u_k,$$

therefore proves comparison for both the atom and its mark derivative. No continuity of the eligibility indicator is needed: it multiplies both differences by the same number in $\{0,1\}$. The sparse product bound established earlier controls the sum over these finite subsets. Widening the source is another common zero extension, so the same argument applies after widening. Finally evaluate ∂_t at $t = 0$. Its absolute value is bounded by the marked envelope, and its comparison is the derivative part just proved. This supplies the forced marked reset comparison used in the two-history output limit.

8.47 Analytic control of omitted insertion terms

One estimate needed in the continuum comparison is obtained by scaling all restored insertion errors with a single complex parameter. The elementary analytic fact is this: if f is holomorphic on $|z| < S$, $S > 2$, and $|f(z)| \le B$ there, then

$$|f(1) - f(0)| \le \frac{2B}{S}. \tag{IC1}$$

For every z on the segment from zero to one, its closed radius-$S/2$ disk lies strictly inside the radius-S disk. Cauchy's derivative estimate gives $|f'(z)| \le B/(S/2)$. Integrate along that segment to obtain (IC1).

Here the normalized insertion moment at level j is

$$\mathrm{Stamp}_{j,t}(P) = \frac{\displaystyle\int (P(U))_t \, d\,\mathrm{prior}_j(U)}{Z_N},$$

where t is an insertion-polynomial coefficient and $Z_N > 0$ is the original Wilson partition function. The denominator is the same at every intermediate level, by the exact transport identity.

Suppose P is served by the insertion family at level j, and its restoration formula is $P = \mathrm{restore}(P_0, P_1)$. Define

$$P(z) = \mathrm{restore}(zP_0, zP_1), \qquad f(z) = \mathrm{Stamp}_{j,t}(P(z)).$$

Then $P(1) = P$, while $P(0)$ is the restoration with the two error inputs removed. The retained insertion estimates give holomorphy on the disk of radius $S = \mathrm{brisk}(B_j) > 2$. They also give integrability for every z in that disk. The scalar complex parameter uses the available support-exponent margin of two; no geometric support is changed.

Transport $P(z)$ to a fixed earlier level $k < j$. The retained transport theorem keeps its served insertion bound there. Integration is unchanged by transport, including after taking the finite polynomial coefficient t; coefficient extraction is a continuous linear map and commutes with this integrable finite-dimensional integral. The root small-field estimate at k therefore gives the uniform disk bound

$$|f(z)| \le \frac{\exp\bigl(|\,\mathrm{Site}_k\,|\,\mathrm{treasury}_k + n(x+y+2)\bigr)}{\mathrm{toll}(w_k, t)} =: \mathcal{B}_{k,t}.$$

Here n is the number of insertion slots, x, y are their two root norm allowances, and the toll uses their positive weights. All volume dependence is at the fixed reporting level k; there is no exponential cost involving the increasingly fine level j.

Applying (IC1) gives

$$\Bigl|\mathrm{Stamp}_{j,t}(P) - \mathrm{Stamp}_{j,t}(\mathrm{restore}(0,0))\Bigr| \le \frac{\mathcal{B}_{k,t}}{\mathrm{brisk}(B_j)/2}. \tag{IC2}$$

This is the uniform omission estimate. It follows from the exact insertion transport and the common analytic disk, and does not require convergence of the intermediate representations of the mixed pin terms.

9 Physical labels and their detection

Physical labels are first finite Taylor-polynomial coefficient objects. Radial gauge fixes their potential representative before equations are imposed; the curl and radial reconstruction maps below

are exact inverses on the Bianchi subspace. Three degrees must be distinguished: polynomial position degree, engineering weight and field-count degree. Later scaling formulas use the engineering weight, whereas Wick products also keep track of field count.

9.1 Radial Taylor variables and the free equation space

The jet variables in this chapter are Taylor coefficients, with the factorial already divided out. At order $j \geq 0$, a radial potential coefficient is a family A_i of symmetric $(j+1)$-linear maps from $\mathbb{R}^4$ to the real color space. Its associated homogeneous potential is $a_i(x) = A_i(x, \dots, x)$, and the constraint is

$$\sum_i x_i a_i(x) = 0 \quad \text{for every } x. \tag{RJ1}$$

This finite-dimensional real space is the radial Taylor space at order j. Its polynomial-position degree is $j+1$, its engineering weight is $j+2$, and its field-count degree is one. Polynomial labels inherit their weight and field count by addition over their factors; constants have both degrees zero. Complex labels are obtained by complexification. No field equation has yet been imposed on these radial variables.

Define the curvature coefficient by the homogeneous polynomial

$$Q_{il}(x) = \partial_i a_l(x) - \partial_l a_i(x). \tag{RJ2}$$

On symmetric tensors this means freezing one input at e_i, subtracting the reversed expression, and multiplying by $j+1$. The curvature is antisymmetric and satisfies

$$\partial_a Q_{bc} + \partial_b Q_{ca} + \partial_c Q_{ab} = 0, \tag{RJ3}$$

since the second derivatives commute. Conversely let Q be a homogeneous degree-j antisymmetric two-form satisfying (RJ3). Set

$$a_i(x) = \frac{1}{j+2} \sum_k x_k Q_{ki}(x). \tag{RJ4}$$

Antisymmetry proves (RJ1). Differentiating (RJ4), subtracting the reversed derivatives, and using (RJ3) gives

$$\partial_i a_l - \partial_l a_i = \frac{2Q_{il} + \sum_k x_k \partial_k Q_{il}}{j+2} = Q_{il},$$

where the last equality is Euler's identity for degree j. Conversely, for a radial potential, differentiation of $\sum_k x_k a_k = 0$ gives $\sum_k x_k \partial_i a_k = -a_i$. Hence $\sum_k x_k Q_{ki} = (j+1)a_i + a_i = (j+2)a_i$. Thus curl and (RJ4) are inverse linear maps. Equality of homogeneous diagonal polynomials implies equality of their symmetric tensors by polarization, so these are identities of the coefficient spaces themselves.

The free solution space consists of those Q which additionally satisfy

$$\sum_a \partial_a Q_{ab} = 0 \quad \text{for every } b. \tag{RJ5}$$

The free gauge-jet map sends this subspace to radial potentials by (RJ4). It is injective because curl is its left inverse. In particular, saying that a radial coefficient is in its range is exactly saying that its curvature satisfies the free equation (RJ5); Bianchi is already built in by (RJ2). Taking normalized Taylor coefficients of a smooth Maxwell solution on an open set preserves (RJ3) and (RJ5): differentiate these equations at the point, commute the derivatives, and divide by the

Taylor factorial. Conversely the polynomial attached to an element of the solution space obeys these equations everywhere. This justifies the solution-jet substitutions used in the equation-ideal argument below.

Color conjugation acts on these coefficients by the adjoint action. A rotation g acts by

$$a_i(x) \longmapsto \sum_k g_{ik} a_k(g^{-1}x).$$

The radial constraint, curl, Bianchi and (RJ5) are preserved, directly by orthogonality and the chain rule. Thus the reconstruction (RJ4) intertwines both actions. The polynomial ideal is the kernel of restriction from radial coefficients to this free solution range, taken at each finite degree horizon. A choice of projection onto that range is a finite-dimensional linear choice; the polynomial telescoping argument below shows explicitly how its kernel is generated by the linear equation coordinates. This explains both the projection P_j used there and the radial reconstruction used in the Gaussian covariance comparison.

9.2 Polynomial labels, coherent frames and local sections

The polynomial variables in this construction are coordinates of real radial jets, rather than coordinates of arbitrary connections. Write V_j for the real space of radial homogeneous potentials of polynomial degree $j+1$, described in (RJ1)–(RJ4), and choose a real basis of each V_j. If $z_{j,a}$ are its coordinate functions, the algebra of labels before imposing gauge invariance is

$$\mathcal{P} = \mathbb{C}[z_{j,a} : j \geq 0,\ 1 \leq a \leq \dim_{\mathbb{R}} V_j]. \tag{PF1}$$

Each polynomial uses finitely many variables. Its evaluation on a real jet $J = (J_j)_{j\geq 0}$ substitutes the real coordinates of J_j, regarded as complex numbers. Such evaluations separate polynomials: reduce to the finitely many variables present, and successively use the fact that a nonzero one-variable polynomial cannot vanish at every real number.

Give $z_{j,a}$ the bidegree $(j+2, 1)$. Thus the bidegree of a monomial z^m is

$$\left(\sum_{j,a} (j+2) m_{j,a},\ \sum_{j,a} m_{j,a} \right) = (r, f). \tag{PF2}$$

The first entry is engineering degree and the second is number of fields. The horizon-d label space consists of the adjoint-invariant polynomials whose monomials have $r \leq d$. It is finite dimensional: only variables of weight at most d occur and every exponent is at most d. Adjoint transformations and orthogonal transformations mix coordinates within a fixed V_j; therefore they preserve both entries of (PF2). Taking a homogeneous component preserves adjoint invariance and the horizon. Conjugation of coefficients preserves these spaces as well, because evaluation takes place on real jets.

There is a basis b_i, with $b_0 = 1$, consisting of real polynomials of individual bidegrees (r_i, f_i). Indeed decompose into the finitely many bidegrees, and then into real and imaginary parts. The resulting real homogeneous elements span the complex space; choose an independent spanning subfamily containing 1. At successive horizons extend the previous independent family inside the same spanning set. This gives coherent frames: previous basis elements remain present as elements, although their nonzero index numbers need not remain fixed. The only elements with field number zero are constants. Hence $f_i > 0$ for $i \neq 0$.

Coefficients $c = (c_i)$ mean the polynomial $\sum_i c_i b_i$; use $\sum_i |c_i|$ as a coefficient norm. There are two different scaling operations:

$$J_j \mapsto h^{j+2} J_j, \qquad J_j \mapsto t J_j. \tag{PF3}$$

On b_i these multiply the reading by h^{r_i} and t^{f_i}, respectively. In particular coefficient scaling by field number is the diagonal map $c_i \mapsto t^{f_i} c_i$, invertible for $t \neq 0$, and fixes the unit. The coefficient action of $s \in O(4)$ is obtained by applying the tensor action to $\sum_i c_i b_i$ and taking its coordinates. This action is an action on the polynomial label space. Its compatibility with the physical quotient requires the separate rotation argument below; it does not follow merely from this definition.

Here is the actual construction of a local lattice representative of each label. At horizon d put

$$A_i^J(x) = \sum_{j<d-1} J_{ji}(x), \qquad a_{x,i}^{J,h}(t) = h A_i^J(h(x + te_i)). \tag{PF4}$$

An empty sum is zero. Along this edge solve the matrix group equation

$$u(0) = 1, \qquad u'(t) = u(t) a_{x,i}^{J,h}(t), \tag{PF5}$$

where the color is embedded in the skew-adjoint matrix Lie algebra. Set the link equal to $u(1)$. Existence does not assume a global logarithm. In an exponential chart, differentiating $\log(\exp b \exp v)$ at $v = 0$ converts (PF5) to a smooth ordinary differential equation in b. Solve locally, left translate to any initial value, and glue. Solutions agreeing once agree throughout an overlapping interval: the derivative of uv^* is zero. The times reachable from zero form an open set with open complement, by the same local extension and gluing argument; connectedness of the real line gives the solution at every time. This also proves uniqueness.

For a finite collection of edges these links depend analytically on the finite jet J near zero. One can see this without differentiating an infinite configuration map. The integral equation in matrix path space is $u = 1 + \mathcal{V}_J u$, where $\mathcal{V}_J$ is integration followed by multiplication by the connection. Near $J = 0$ its norm is less than one, so $(1 - \mathcal{V}_J)^{-1} = \sum_{n \geq 0} \mathcal{V}_J^n$ is analytic. Uniqueness identifies it with (PF5). Differentiating at zero gives the edge tangent $\int_0^1 A_i^J(x + te_i)\, dt$. Products and inverse links therefore have tangents equal to their signed edge sums.

A sampling slot consists of an oriented elementary plaquette based at an integer point and a finite path from the origin to that point. Evaluate its curvature and transport the color back along that path. A gauge transformation changes every slot by the same adjoint transformation at the origin. At the identity configuration, the connector derivatives cancel: differentiating ZPZ^{-1}, with $Z = P = 1$, leaves just P'. Consequently the linearization of the slot is the circulation of the integrated polynomial potential around the plaquette.

These linear samples separate the truncated radial jets. For one homogeneous component $A \in V_k$, define

$$N_{ij}(x,t) = \int_0^1 \int_0^1 \partial_i A_j(x + tse_j + tre_i)\, dr\, ds - \int_0^1 \int_0^1 \partial_j A_i(x + tse_i + tre_j)\, dr\, ds. \tag{PF6}$$

At $t = 1$ the fundamental theorem of calculus identifies this with the square circulation. At $t = 0$ it is $\partial_i A_j - \partial_j A_i$. Homogeneity gives $N_{ij}(tx, t) = t^k N_{ij}(x, 1)$. If all integer square samples of a sum of components through k vanish, multiply the rescaled component of degree ℓ by $t^{k-\ell}$. Choose integer x_t with $tx_t \to x$. The resulting finite sum is zero for $t > 0$; continuity at zero leaves only $N_{ij}^{(k)}(x, 0)$. Thus the top curl vanishes. Radial reconstruction (RJ4) makes the top potential zero. Descending induction eliminates all components. No interchange of the two iterated integrals in (PF6) is required.

Only finitely many slots are needed. Among kernels of finite sampling families choose one of minimal dimension. Appending any further slot cannot lower that dimension, hence cannot lower the kernel. A vector in this kernel is therefore killed by every slot and, by the preceding argument, is zero. Write $X = \prod_{j<d-1} V_j$, $Y = \mathfrak{g}^m$, and $T : X \to Y$ for the resulting injective linearization.

There is an adjoint-equivariant left inverse $C : Y \to X$. Equip Y with the sum of the invariant color inner products. The range of T is invariant under the adjoint action in both directions, so its orthogonal projection commutes with that action. Compose this projection with T^{-1} on its range. This proves the equivariance of C without needing an invariant inner product on X.

Let $B : X \to Y$ denote the nonlinear slot readings of the links (PF5) at $h = 1$. Then $B(0) = 0$, $DB(0) = T$, and $D(CB)(0) = 1$. The finite-dimensional analytic inverse function theorem supplies a local inverse g^{-1} of $g = CB$. Define

$$R(y) = g^{-1}(Cy). \tag{PF7}$$

It is analytic near zero, $R(0) = 0$, and $R(B(x)) = x$ near zero. For each fixed gauge transformation its equivariance follows by applying g to both candidate expressions, using equivariance of B, C, and using injectivity on a sufficiently small common neighborhood. The neighborhood is allowed to depend on that transformation.

Replace this shared recovery map by its Taylor polynomial

$$R_{\le d}(y) = \sum_{q=0}^{d} \frac{1}{q!} D^q R(0)[y, \dots, y], \qquad P_i(y) = b_i(R_{\le d}(y)). \tag{PF8}$$

All derivatives through order d of the first expression equal those of R; differentiating each homogeneous term proves this directly, including its factorial. Expanding multilinear forms in real panel coordinates makes P_i an actual complex polynomial. Its evaluation on real panels is real, and it is adjoint invariant. Because $R_{\le d}(0) = 0$, each substituted coordinate has no constant term. Every monomial of b_i has exactly f_i factors. Hence every monomial of P_i has at least f_i slot factors. The substitution is not truncated again: terms of degree greater than d are retained.

The geometric scaling identity for (PF5) is

$$U(J, h) = U(\mathrm{trunk}_d(h^{j+2} J_j), 1). \tag{PF9}$$

It follows from homogeneity in (PF4) and uniqueness of the edge equation. Along this curve the slot panel is $B(\mathrm{trunk}_d(h^{j+2} J_j))$. Equations (PF7)–(PF8), and the chain rule through order d, therefore give

$$\frac{d^q}{dh^q} P_i(\text{slots of } U(J, h))\Big|_{h=0} = \frac{d^q}{dh^q} b_i((h^{j+2} J_j)_j)\Big|_{h=0}, \quad 0 \le q \le d. \tag{PF10}$$

This is equality of finite Taylor readings, not equality at nonzero mesh.

To impose the signed lattice symmetries, extend P_i linearly in the coefficient vector, precompose configurations by each signed coordinate permutation, apply the inverse tensor action to the coefficient input, and average over this finite group. Equation (PF10) is preserved: the holonomy of the transformed connection is the transformed holonomy, including reversal and inversion of negatively oriented edges, by (PF5) and uniqueness. The tensor action preserves field number. A transformed slot is a signed slot; combining all these finite slot families into one list and substituting signed variables therefore preserves the lower bound f_i on monomial degree. The averaged section fixes the unit and obeys the exact signed symmetry and reality laws.

All paths and squares used are finite. Enclose them and their signed images in one origin-containing symmetric integer box. Each transported curvature slot is a polynomial in matrix entries

of its links and their conjugates: inverse unitary links use conjugate transpose, and curvature uses its fixed linear projection of the skew part. Thus the section is a finite polynomial cylinder, gauge invariant, continuous everywhere and smooth along link variations. On a torus it is placed by translating and reducing this finite stencil. Formula (PF10) is used when the stencil does not wrap; the cylinder itself is defined even when it wraps.

Finally fix a derivative order and the direction power of (AB4). On this fixed stencil each transported curvature slot, including its normalized ordered left derivatives, has an envelope bounded by $C\theta$ on the healthy domain. The zeroth estimate is the healthy plaquette bound and adjoint isometry; differentiated plaquettes and connector products have bounded coefficients on the compact group and each allowed direction contributes θ times a fixed stencil constant. For a differentiated product of at least f_i slots the Leibniz rule distributes the derivatives; an undifferentiated slot pays θ, and a differentiated slot pays at least one θ. Further derivatives only improve this power since $\theta \le 1$. Finite sums and a maximum over the finite basis give a constant A, independent of torus size, placement and the remaining budget parameters, such that the section of b_i has normalized envelope

$$A\theta^{f_i}. \tag{PF11}$$

This is the local section estimate used in training and reset budgets. It includes nonlinear terms of (PF8); it is not an estimate only of the linearized sampler.

The actual finest-lattice insertion includes two coefficient changes in a specified order. Let $\theta_N > 0$ be the healthy-field threshold at level N, let $S_d(c; o, U)$ be the section just constructed, and let $\mathcal{T}_{d,N,p}$ be the terminal coefficient isomorphism of the selected history. Put

$$I_{d,N,p}(v; o, U) = a_N^{-d} S_d\Big(\big(\theta_N^{-f_i} (\mathcal{T}_{d,N,p}^{-1} v)_i\big)_i ; o, U \Big). \tag{PF12}$$

Thus calibration is applied first, threshold scaling second, and engineering normalization last. The exponent in the threshold scaling is the field number f_i, not the engineering degree r_i. A general horizon label contains several bidegrees, so neither diagonal scaling may be replaced by one scalar on that entire space.

Although the coefficient vector in (PF12) depends on N, p, its stencil is the single stencil of S_d. Linear combination and multiplication by constants do not enlarge it. Consequently (PF12) is continuous, gauge invariant and supported on that translated stencil for every cutoff, including cutoffs at which it wraps around the torus. The finite-Taylor normalization and the calibrated symmetry statements are used only after the reference level fits below N; this distinction is needed when constructing a family defined at every cutoff.

At such a cutoff the terminal map commutes with conjugation and signed orthogonal actions. Its inverse does too, by applying the terminal map to the two candidate values. Threshold scaling commutes with those actions because they preserve field number. Together with the corresponding identities for S_d, these facts prove the reality and signed covariance of (PF12). If e_0 denotes the coefficient vector of the constant polynomial, the physical unit at level m is $a_m^d e_0$. The based-history identity is

$$\mathcal{T}_{d,N,p}(a_N^d e_0) = a_{k_d}^d e_0. \tag{PF13}$$

Substituting its inverse in (PF12), using $f_0 = 0$ and $S_d(e_0) = 1$, proves that the reference physical unit gives the constant cylinder one. Translating the placement translates precisely the same coefficient vector; no new history or new calibration is chosen at the new site.

9.3 Calibrated coefficient changes and exact solutions

Calibration is a finite-dimensional coordinate change on the specified coefficient space. Equation cancellation is proved on the entire relevant cylinder, with masks and gauge averaging retained, before passing to a quotient. A detector separates the physical classes only after this equation image has been shown to be silent against every allowed spectator reading. These are algebraic and analytic steps with separate hypotheses.

The inverse of the local slot sampler in (PF8) is an inverse on truncated radial jets. The coefficient inverse required for a physical insertion is a different map. We construct it in the norms and units in which it is used. Fix a degree horizon d, its coherent frame, and a reporting level $k \geq k_d$, where k_d is the reference level for that frame. All assertions in this paragraph are uniform in $N \geq k$ and the period p, with k fixed.

Let $\mathcal{T}_{d,N,p}$ be the calibrated terminal coefficient isomorphism of the insertion history. If v is its reference coefficient, transport $\mathcal{T}_{d,N,p}^{-1}v$ along that same history to level k and take its primary coefficient. Denote the result by

$$C_{d,k,N,p}v, \qquad A_{d,k,N,p}v = a_k^{-d}C_{d,k,N,p}v. \tag{PC1}$$

At the reference level $C_{d,k_d,N,p}v = v$, by the normalization of the history. At an adjacent step its coefficient estimates give, in the frame norm $|\cdot|_1$,

$$|C_iv|_1 \leq R_{i+1}L_i^d|C_{i+1}v|_1, \qquad |C_{i+1}v|_1 \leq R_{i+1}|C_iv|_1, \tag{PC2}$$

where $R_{i+1} \geq 1$ is the actual column allowance and $L_i \geq 1$ the block side. These are bounds on the primary coefficients of the original transported entries. They do not replace the primary coefficient by its class in the physical quotient.

Define the finite number

$$O_{d,k} = \prod_{i<k} R_{i+1}L_i^d \geq 1. \tag{PC3}$$

Iterating the first inequality of (PC2), starting at the reference level, and enlarging the product by its factors before that level, proves

$$|v|_1 \leq O_{d,k}|C_{d,k,N,p}v|_1 = O_{d,k}a_k^d|A_{d,k,N,p}v|_1. \tag{PC4}$$

Iterating the second inequality gives a finite upper bound as well. Thus C and A are injective and hence bijective on the finite coefficient space, with both bounds uniform in the remaining cutoff and period at this fixed level. No bound on $O_{d,k}$ as $k \to \infty$ is asserted.

The readout map is constructed by polynomial substitution followed by Gaussian contraction. Start with the linear map sending a reference coefficient v to the local polynomial representative of its transported insertion. Rename its finitely many slot variables as the actual link-coordinate variables, multiply each by the threshold factor, and embed them in the one common stencil. Let $M(v)$ denote the resulting polynomial. For its actual symmetric covariance Q, Gaussian contraction means the finite linear operator

$$\mathcal{W}_Q = \exp\left(\frac{1}{2}\sum_{\alpha,\beta} Q_{\alpha\beta}\partial_\alpha\partial_\beta\right). \tag{PC5}$$

The exponential terminates on each polynomial. Apply $\mathcal{W}_Q$, substitute the linearized slot expressions in the radial-jet variables, and retain monomials of engineering weight at most d. The resulting

polynomial is invariant and real when v is real: substitution intertwines the color action, the covariance is invariant, and real coefficients and real covariance commute with coefficient conjugation. It therefore has a unique coefficient vector in the coherent frame.

There is a last field-number rescaling in this definition. Put

$$\xi_{k,N,p} = \frac{\beta_{N,p,k}^{-1/2}}{\vartheta_k} > 0, \qquad (D_\xi z)_i = \xi^{f_i} z_i, \tag{PC6}$$

where ϑ_k is the original barrier threshold, not a restarted threshold. Define the linear readout $T_{d,k,N,p}$ by requiring that the polynomial obtained from (PC5) have coefficient vector $D_\xi T_{d,k,N,p} v$. This determines T because every diagonal entry of D_ξ is nonzero. It also explains the exact calibration identity: the retained polynomial, including its Gaussian contraction, is the polynomial with those coefficients. There is no assertion that the whole untruncated exponential insertion is polynomial.

The analytic comparison estimate for this construction is measured before the rescaling D_ξ:

$$|A_{d,k,N,p} v - T_{d,k,N,p} v|_1 \le \frac{1}{2} |A_{d,k,N,p} v|_1. \tag{PC7}$$

We derive this relative estimate before solving any calibrated columns. Let n bound the ordinary polynomial degree, let $z = \vartheta_k / g \ge 1$, $g = \beta_{N,p,k}^{-1/2}$, and multiply every variable of the unweighted insertion polynomial by z; call that polynomial q. Its coefficient sum obeys $\mathrm{coin}(q) \le B|Av|_1$. The actual nonlinear transported-variable substitution has the form

$$X_\alpha \longmapsto f_\alpha + \vartheta_k^{-1} \mathcal{S}_{\vartheta_k} A_\alpha, \tag{PC7a}$$

where f_α is linear, $\mathcal{S}_t$ multiplies every variable by t, and every monomial of A_α has ordinary degree at least two. Their coefficient sums are at most B. For $0 < t \le 1$,

$$\mathrm{coin}(t^{-1} \mathcal{S}_t A_\alpha) \le tB, \quad \mathrm{coin}(f_\alpha + t^{-1} \mathcal{S}_t A_\alpha) \le 2B. \tag{PC7b}$$

The first assertion follows term by term from $t^{m-1} \le t$ for $m \ge 2$; the second is the triangle inequality. Telescoping each monomial of degree at most n, replacing one variable at a time, bounds the nonlinear substitution error by $\mathrm{coin}(q) n (tB)(2B)^n$. For a covariance with entries at most $\epsilon \le 1$, Wick's finite pairing expansion bounds $\mathrm{coin}(\mathcal{W}_Q q - q)$ by $\mathrm{coin}(q)\epsilon[2(n+2)]^n$. After the linear substitution this costs a further factor B^n. These two errors give

$$\mathcal{A}(t, B, \epsilon, n) = n(tB)(2B)^n + \epsilon[2(n+2)]^n B^n. \tag{PC7c}$$

Engineering truncation can only decrease coefficient sum.

For the actual holonomy Taylor reading, the calibration before Wick contraction has error at most $\epsilon|Av|_1$ in polynomial coefficients. This is the Taylor remainder of the transported slots in (PF8)–(PF11), with their ordinary nonlinear part separated as in (PC7a). On the finite-dimensional image of the polynomial exposure map, fix a constant $c \ge 1$ bounding its inverse from polynomial coefficient sum to frame coefficient sum. The total frame error is therefore

$$|Av - Tv|_1 \le c\{\epsilon + B\mathcal{A}(\vartheta_k, B, \epsilon, n)\} |Av|_1. \tag{PC7d}$$

The covariance here is the rescaled covariance $z^{-2}Q$, since scaling variables by z and Wick contraction commute with that change of covariance. The band estimate $gH_k \le 2\vartheta_k$ implies $|(z^{-2}Q)_{\alpha\beta}| \le \epsilon$ when $|Q_{\alpha\beta}| \le B$ and $4B \le H_k$.

Choose $a = [4(n+2)]^{-1}$, $B = H_k^a$, $\epsilon = H_k^{-1}$. The original threshold obeys $\vartheta_k \le H_k^{-1}$ eventually. The terms on the right of (PC7d), including its outer B, have powers of H_k at most

$$-1, \quad -1 + a(n+2) = -\tfrac{3}{4}, \quad -1 + a(n+1) < -\tfrac{3}{4}. \tag{PC7e}$$

Their fixed prefactors depend on d, n, c, not on N, p. Thus $2c\{\epsilon + B\mathcal{A}\} \le 1$ eventually, proving (PC7) uniformly in those remaining indices. Also $B + B\epsilon[2(n+2)]^n \le H_k$ eventually, so the Wick polynomial fits its stated monitor price. These estimates use the original threshold and the actual rescaled covariance; no solved-column bound has been used to prove the inverse.

The unit is exact: its representative is the constant polynomial, contraction and truncation fix it, and both rescalings have weight zero on it. Thus $Tv_{\mathrm{unit}} = e_0$. We retain (PC7) in its relative form; an absolute error on a few selected fields would not imply the inverse assertion that follows.

The triangle inequality gives

$$|Av|_1 \le 2|Tv|_1, \qquad |Tv|_1 \le \tfrac{3}{2}|Av|_1 \le 2|Av|_1. \tag{PC8}$$

Combining the first inequality with (PC4) proves

$$|T^{-1}z|_1 \le 2O_{d,k}a_k^d|z|_1. \tag{PC9}$$

To justify existence of this inverse before using it, $Tv = 0$ implies $Av = 0$ by (PC8), then $v = 0$ by (PC4). An injective endomorphism of a finite-dimensional space is surjective. This proves both existence and the bound (PC9).

For a desired reference coefficient b and a position magnification ρ, let $(W_\rho b)_i = \rho^{r_i} b_i$. The exact solved calibration is

$$D_\xi T u = W_\rho b, \qquad u = T^{-1} D_{\xi^{-1}} W_\rho b. \tag{PC10}$$

The two diagonal operations use the distinct entries f_i, r_i of the bidegree. The solution is unique. Real b gives real u, because conjugating the equation gives the same equation and uniqueness applies. If $b = e_0$, the solution is v_{unit} by the exact unit identity. This constructs all the solved basis columns simultaneously for the one readout, rather than choosing a different readout for each column.

The finite cutoff bounds on β give

$$\xi^{-1} = \vartheta_k \sqrt{\beta_{N,p,k}} \le 2H_k, \tag{PC11}$$

where H_k is the actual monitor at level k. If $b_1, \dots, b_s$ are the finitely many chosen real representatives of a free quotient basis, (PC9)–(PC11) imply the uniform fixed-level bound

$$\|u_j\| \le 2O_{d,k}a_k^d \sum_{j'=1}^{s} |D_{2H_k} W_\rho b_{j'}|_1. \tag{PC12}$$

Indeed bound each component of $D_{\xi^{-1}}$ by the corresponding component of D_{2H_k}, use $\|\cdot\| \le |\cdot|_1$, and enlarge the single basis-column bound to their finite sum. This is the compact ball used in extracting the solved columns. It is allowed to grow with the reporting level.

For the subsequent scale-dependent prices write $H_k = e^{B_k}$, choose $\rho = H_k^b$, and put $q = \max_i f_i$. Formula (PC10) bounds

$$|Tu|_1 \le (2H_k)^q \rho^d |b|_1, \qquad H_k|Au|_1 \le 2^{q+1}|b|_1 H_k^{bd+q+1}. \tag{PC13}$$

For fixed b and d, the finite prefactor is eventually paid by one further factor H_k. The allowed polynomial coefficient price is therefore H_k^{bd+q+2}. This derives the degree-dependent price used for the calibrated polynomials. It also separates it from the fixed-level inverse bound (PC9) and from the inverse of the separating-test matrix used in (WS5).

9.4 The full free Maxwell polynomial hierarchy

The free comparison hierarchy is specified on all polynomial labels, with no restriction to its two-point function. Let $\mathcal{S}_j$ be the real space of homogeneous curvature Taylor coefficients of order j satisfying antisymmetry, Bianchi and the Maxwell divergence equation in (RJ3),(RJ5). Choose a real basis and write $X_{j,\alpha}$ for its coordinate symbols. The free polynomial algebra is

$$\mathcal{W} = \mathbb{C}[X_{j,\alpha}], \qquad \operatorname{weight}(X_{j,\alpha}) = j + 2. \tag{FM1}$$

A real solution tower $A = (A_j)$ evaluates these symbols in its real coordinates. Evaluation separates polynomials by the finite-variable argument of (PF1). The restriction map from radial labels substitutes the radial potential reconstructed from A_j into each original jet coordinate. Its kernel at horizon d is exactly the equation ideal: it consists of polynomials vanishing on every such solution tower. Its image $\mathcal{W}_d^{\mathrm{inv}}$ is the invariant sector of weight at most d. To see surjectivity onto this sector, use the color-equivariant projection of radial curvature coefficients onto their Maxwell subspace and then radial reconstruction to lift every linear solution coordinate. Polynomial substitution preserves weight and equivariance; restriction of this lift is the identity. This is an algebraic operation on coefficients, not a claim about interacting equations.

The basic covariance is fixed by

$$G_0(z) = \frac{1}{4\pi^2|z|^2}, \qquad C_{ab,kl}(z) = \delta_{bk}\partial_a\partial_l G_0 + \delta_{al}\partial_b\partial_k G_0 - \delta_{bl}\partial_a\partial_k G_0 - \delta_{ak}\partial_b\partial_l G_0, \quad z \neq 0. \tag{FM2}$$

Colors contract by their invariant Euclidean inner product. This is the curl at the first endpoint and the opposite curl at the second endpoint of the potential covariance; the opposite endpoint derivative accounts for its signs. Its two legs are closed and divergence free away from zero, by commutation of derivatives and harmonicity of G_0.

The contractions of higher solution coordinates use normalized Taylor coefficients, so the factorial convention must be retained. If $p \in \mathcal{S}_j^*$, first extend it linearly to the ambient two-form Taylor tensors. Apply it to the order-j normalized Taylor tensor of the first leg of (FM2), obtaining a two-form in the second indices and color. Pair its order-l normalized Taylor tensor with an extension of $q \in \mathcal{S}_l^*$, and multiply by $(-1)^l$. Denote the resulting scalar by $K_{p,q}(z)$. Both extensions are immaterial: at each point away from zero the differentiated leg is in the corresponding solution subspace, so their differences annihilate it. This defines the contractions on precisely the constrained variables in (FM1).

For monomials at distinct points $x_1, \dots, x_n$, regard every power of every symbol as that many distinguishable occurrences. Sum over perfect matchings of all these occurrences, allowing a pair only if its two occurrences belong to different insertion points. A pair $(i,p),(k,q)$ contributes $K_{p,q}(x_i - x_k)$. The product over its pairs, summed over the permitted matchings, is the monomial moment. The empty matching contributes one and an odd total number contributes zero. Extend by complex multilinearity in the n polynomial inputs:

$$M_n(x; P_1, \dots, P_n) = \sum_{\text{monomial choices}} \left(\prod_i c_i \right) \sum_{\substack{\text{perfect matchings} \\ \text{no pair within one insertion}}} \prod_{\text{pairs}} K_{p,q}(x_i - x_k). \tag{FM3}$$

The restriction on pairs is the Wick ordering of each composite insertion. There are finitely many terms for every tuple of polynomials. Pulling (FM3) back by the restriction map defines the free table on the original coherent label frames.

The elementary contraction is symmetric under exchange of endpoints and legs. Reindexing matchings therefore proves permutation symmetry. It depends only on differences, proving translation invariance. The chain rule applied to (FM2), orthogonality of the color action, and orthogonal covariance of the two curls give the full orthogonal tensor law. All entries are real on real coefficients, giving the reality law. Inserting the constant polynomial one introduces no occurrences, so it removes that slot. These prove the algebraic hierarchy laws for every arity.

Each derivative of a contraction of orders j, l is bounded by a constant times a power of the inverse separation; at order zero the power is $4 + j + l$. This follows by differentiating the rational homogeneous function (FM2), with the finite normalization factors included in its constant. Taking products and finite sums in (FM3) proves analyticity on the separated configuration space and a power bound there. For the product bound add the exponents and multiply constants; for a sum take their sum and enlarge to a common exponent using $r^{-1} \geq 1$ when $0 < r \leq 1$. These are the positional bounds needed for its smeared hierarchy.

9.5 Positivity and strictness of the reflected free pairing

We derive the reflected one-particle positivity needed for polynomial detection. Fix an axis a, let θ_a reflect that coordinate, and let R_a act on all tensor indices of a probe by this reflection. Take finitely many real probes p_i of possibly different jet orders d_i, at points x_i with $(x_i)_a > 0$. Coincident points among these sources are allowed. Float all sources by te_a, $t \geq 0$, and put

$$H(t) = \sum_{i,j} K(p_i, R_a p_j; x_i + te_a, \theta_a(x_j + te_a)),$$

where the last two arguments indicate the endpoints of the difference kernel. Let F_t be the sum of their curvature pulses. The derivative in the first endpoint and the opposite motion of the reflected endpoint give

$$H'(t) = -2L(t), \qquad L(t) = \sum_{i,j} -K(\partial_a p_i, R_a p_j; x_i + te_a, \theta_a(x_j + te_a)). \tag{GP1}$$

Here $\partial_a p_i$ is the raised probe, with the normalized derivative convention of (FM2). The factor two is the sum of the two endpoint motions.

On the reflecting plane the reflected pulse is $F_{t,bc} \mapsto \epsilon_b \epsilon_c F_{t,bc}$, where $\epsilon_a = -1$ and all other $\epsilon_b = 1$. Substitute these signs into the bilinear Maxwell stress. Its normal-normal component is $-\frac{1}{2}\sum_{b,c} \|F_{t,bc}\|^2$. The mixed terms with exactly one normal index change sign; skew symmetry gives $F_{t,aa} = 0$, and combining the remaining terms yields exactly this negative half-sum. Integrate its divergence over the positive half-box, removing small balls about its finitely many upper sources. The plane term has the negative outward normal; the small-sphere terms are the translation case of the pulse-current identity (FS2). The reflected sources are in the negative half-space and cause no extra sphere term. The other box faces disappear as their size tends to infinity. Indeed each pulse of order d is $O(|y|^{-4-d})$; its bilinear stress is $O(|y|^{-8-d_i-d_j})$, while the three-dimensional face measure is $O(R^3)$. Their fluxes are therefore $O(R^{-5-d_i-d_j})$. Summing the finitely many terms proves

$$L(t) = \lim_{R\to\infty} \frac{1}{2} \int_{[-R,R]^3} \sum_{b,c} \|F_{t,bc}(y_a = 0)\|^2 \, d^3y \geq 0. \tag{GP2}$$

The same argument bounds each finite-box integral by $L(t)$, by monotonicity of the nonnegative integrand. Since every source-to-reflected-source separation tends to infinity, the pole bounds give $H(t) \to 0$. Equations (GP1)–(GP2) make H decreasing, so $H(t) \geq 0$. This proves the real positive

semidefiniteness of every finite reflected one-particle matrix. Symmetry of that matrix follows from endpoint exchange and reflection covariance. Its complex extension is Hermitian positive semidefinite by splitting a coefficient vector into real and imaginary parts.

The equality case is stronger when all source points equal one point x with $x_a > 0$. Suppose $H(0) = 0$. Monotonicity and nonnegativity imply $H(t) = 0$ for all $t \geq 0$, hence $L(t) = 0$ for $t > 0$. Equation (GP2), together with continuity, forces every component of F_t to vanish on the entire reflecting plane: a nonzero value would have a neighborhood with positive squared integral inside some finite box. Translation covariance then makes the unshifted summed pulse vanish on $y_a = -t$, for every $t > 0$. Each scalar color component is real analytic on $\mathbb{R}^4 \setminus \{x\}$, which is connected. The analytic identity theorem extends this half-space vanishing to every $y \neq x$.

Apply now the radial version of (FS2) to this zero summed pulse and any polynomial Maxwell solution f. It yields

$$\sum_i (d_i + 2)\, p_i(j_{d_i} f(x)) = 0, \tag{GP3}$$

where j_k is the normalized solution jet of order k. For any homogeneous solution tensor A_k, the field $f(y) = \mathrm{asWave}(A_k)(y - x)$ has normalized jet A_k at order k at x, and zero jets there at all other orders. Thus (GP3) gives $\sum_{i:d_i=k} p_i(A_k) = 0$, since $k + 2 > 0$. Summing over the finitely many jet orders proves that a zero reflected norm annihilates every solution tower. In particular, for any finite alphabet of independent solution-coordinate probes, its one-point reflected matrix is positive definite: evaluate the last identity on a tower with a single basis coordinate equal to one to force each coefficient to be zero. This establishes strictness on the constrained solution spaces themselves, including mixtures of jet orders.

We pass from legs to polynomial words algebraically. On finitely many complex variables define

$$\langle P, Q \rangle_{\mathrm{F}} = \sum_\nu \left(\prod_i \nu_i! \right) \overline{P_\nu} Q_\nu. \tag{GP4}$$

Every weight is positive; hence $\langle P, P \rangle_{\mathrm{F}} > 0$ for $P \neq 0$, including the scalar polynomial. Factor a finite reflected leg matrix as $L^* L$, which is ordinary finite-dimensional positive matrix factorization. Associate to a leg i the creation polynomial $\ell_i(Z) = \sum_k L_{ki} Z_k$. To represent a word, process its occurrences successively: append its creation polynomial and, for each possible contraction to a previous occurrence on the same side, add that covariance times the polynomial for the remaining word. This finite recursion is Wick's pairing recurrence. Pairing the two resulting polynomials by (GP4) contracts their remaining occurrences across the reflection plane, since $\langle \ell_i, \ell_j \rangle_{\mathrm{F}} = (L^* L)_{ij}$. The factorial $\prod \nu_i!$ counts the bijections between equal creation occurrences. Induction on word length therefore identifies the reflected Wick moment with this polynomial inner product, including all permitted contractions internal to either half. A finite linear combination of words has reflected norm equal to the nonnegative squared norm of its associated polynomial. Expanding polynomial insertions into monomials proves reflection positivity at arbitrary arity, and smearing preserves it by integration.

At a single source point there are no contractions within either half because each insertion is Wick ordered. The recursion is then simply the substitution $Z_i \mapsto \ell_i(Z)$. The one-point leg matrix is positive definite by (GP3), so L is invertible; substituting L^{-1} gives the inverse polynomial substitution. Thus a nonzero polynomial remains nonzero, and (GP4) makes its reflected pairing strictly positive. This is the strict pairing used in the hosting argument and the detector construction. It proves that the free all-spectator null space has precisely the equation ideal already removed in forming the solution coordinates.

Accordingly the free physical field space is canonically

$$E_d/N_d \simeq \mathcal{W}_d^{\mathrm{inv}}. \tag{FM4}$$

Under this identification the horizon inclusion is literal polynomial inclusion. A derivative in direction v is the polynomial derivation determined on a linear symbol by

$$\mathsf{D}_v \ \mathrm{symbol}_j(p) = \mathrm{symbol}_{j+1}(p \circ \mathrm{lower}_v). \tag{FM5}$$

Here $\mathrm{lower}_v A_{j+1}$ is the order-j Taylor coefficient of the derivative of the homogeneous polynomial represented by A_{j+1}; its coefficient is $j+1$ times contraction with v. Thus (FM5) has the actual Taylor normalization. The Leibniz rule differentiates one occurrence at a time. Differentiating a matching product in (FM3) selects its incident covariance edge, while replacing an occurrence by (FM5) differentiates precisely that same edge. Summing proves the insertion-derivative identity. It raises the horizon by one and stays in the invariant sector. Conjugation and signed transformations obey (QD6). Hence the free differential system is the same type of all-spectator system as the interacting one.

For $s \in O(4)$ the free field action is the polynomial substitution induced by the tensor action on the solution tower. It preserves weight, equations and color invariance, with inverse given by s^{-1}; conjugating it through (FM4) gives the full field action. Infinitesimal skew rotations likewise act by the derivation obtained by differentiating that linear tensor substitution. Both descriptions are on the same field space and preserve the moment law (FM3).

The free symmetric traceless stress is the actual weight-four invariant polynomial

$$T_{ab}^{\mathrm{free}} = \sum_{c,i} F_{ai}^c F_{bi}^c - \frac{\delta_{ab}}{4} \sum_{c,i,k} F_{ik}^c F_{ik}^c. \tag{FM6}$$

Products here mean the Wick polynomial product used in (FM3). Adjoint invariance follows from preservation of the color inner product, symmetry by exchanging its two factors, and reality from real coefficients. Its trace is zero by summing the four diagonal entries. Its divergence is zero as a polynomial on solution jets. Indeed the derivative of the first term is $\sum_{c,i} F_{ai}^c \partial_a F_{bi}^c$, since $\partial_a F_{ai}^c = 0$. Bianchi and antisymmetry give

$$\sum_{a,i} F_{ai}^c \partial_a F_{bi}^c = \frac{1}{4} \partial_b \sum_{a,i} F_{ai}^c F_{ai}^c; \tag{FM7}$$

multiply the Bianchi equation by F_{ai}^c, sum, and interchange a, i in its second derivative term to obtain the factor two on the left and the half derivative on the right. This cancels the divergence of the subtracted term in (FM6). Evaluation on all real solution towers then proves the polynomial identity. Through (FM4) it gives a real, symmetric, conserved free field at horizon four, with the exact normalization used in the later free comparison.

9.6 The free low-degree stress and vector spaces

The low-degree classification uses simplicity of the actual color Lie algebra, without assuming that the matrix group is connected. Differentiate its adjoint action along the actual exponential curves. This gives

$$\frac{d}{dt} \mathrm{Ad}_{\exp(tx)} \, y \Big|_0 = [x, y]. \tag{FC1}$$

The span of brackets is a nonzero ideal, hence is the whole simple nonabelian color algebra. An invariant linear functional annihilates each bracket by (FC1), so it is zero. An invariant real bilinear

form satisfies $B([x,y],z) = -B(y,[x,z])$. Its alternating part is zero: three cyclic applications to $B([x,y],z)$, with bracket antisymmetry, identify this value with its negative, and brackets span. For the symmetric part write $B(x,y) = \langle Ax, y\rangle$ with A self-adjoint. Invariance says $A[x,y] = [x,Ay]$. Every eigenspace of A is therefore an ideal. A real self-adjoint operator has an eigenvalue; its nonzero eigenspace is the whole algebra, so A is scalar. Applying this argument to real and imaginary parts proves that every invariant complex-valued real bilinear form is a complex multiple of the fixed color inner product.

A monomial of weight at most five has at most two fields, since each variable has weight at least two. With two fields its jet orders are either $(0,0)$ or $(0,1)$; the latter has weight five. Decompose a polynomial into its constant, linear and these quadratic pieces by ordinary field number and engineering weight. Invariance of each piece follows from the linearity of the color action. The linear piece vanishes by (FC1). Polarizing the $(0,0)$ piece, and separating the two jet orders in the $(0,1)$ piece, gives invariant bilinear forms. Resolve their spatial tensor coordinates; (FC1) then makes every color contraction a multiple of the color inner product. Consequently all weight-four labels are spanned by the unit and $\langle F_{ab}, F_{cd}\rangle$; through weight five one adds $\langle F_{ab}, \partial_e F_{cd}\rangle$. This is a spanning assertion about the actual invariant labels before further free equation relations.

Put $Q_{ab,cd} = \sum_c F^c_{ab} F^c_{cd}$, with products interpreted as in (FM3), and define

$$N_a = \sum_i Q_{ai,ai}, \qquad R_a = \sum_{i,k:\ i\neq a,\ k\neq a} Q_{ik,ik}, \qquad P_{ab} = \sum_i Q_{ai,bi}. \tag{FC2}$$

The sums in R_a are over ordered pairs. Consider centered real symmetric tensor fields transforming under the signed coordinate group as rank-two tensors. Reflection in one axis multiplies F_{ab} by the product of its two axis signs. Thus a diagonal component permits only products in which every axis occurs evenly; for two antisymmetric pairs this forces the same unoriented plane twice. The permutations fixing axis a have just two orbits: planes containing a and planes disjoint from a. They give N_a, R_a. A constant could also occur, but centeredness removes it: the Wick quadratic one-point functions are zero and the unit expectation is one.

For the off-diagonal component ab, $a \neq b$, the parity character is the product of the signs of a, b. The two planes must therefore share exactly one other axis i, and be ai, bi, up to antisymmetry and exchange of the two factors. Permutations fixing a, b identify the possible remaining axes, giving precisely P_{ab}. The characters of the other products differ in at least one reflection, so averaging against the required character kills them. Permuting axes then determines every tensor component from 00 and 01. Reality makes the coefficients real. The whole centered signed tensor space is therefore spanned by

$$A_{ab} = \delta_{ab} N_a, \qquad B_{ab} = \delta_{ab} R_a, \qquad C_{ab} = 1_{\{a\neq b\}} P_{ab}. \tag{FC3}$$

For a real polar vector of weight at most five, fix component zero and average under its stabilizing permutations and reflection character. Constants and FF terms have the wrong parity. In an $F_{ab}\partial_e F_{cd}$ term, parity requires that the symmetric difference of the two plane-index sets be $\{0, e\}$. If $e = 0$, the planes coincide, giving the first two sums below. Otherwise they share a remaining index; the location of the derivative leg gives the last two sums. Thus its zeroth component is spanned by

$$\begin{aligned} n &= \sum_{i\neq 0} \langle F_{0i}, \partial_0 F_{0i}\rangle, \qquad c = \sum_{\substack{i,j\neq 0\\ i\neq j}} \langle F_{ij}, \partial_0 F_{ij}\rangle, \\ a &= \sum_{\substack{i,j\neq 0\\ i\neq j}} \langle F_{0i}, \partial_j F_{ji}\rangle, \quad h = \sum_{\substack{i,j\neq 0\\ i\neq j}} \langle F_{ji}, \partial_j F_{0i}\rangle. \end{aligned} \tag{FC4}$$

The free divergence equation gives $n + a = 0$. Bianchi gives $\langle F_{ji}, \partial_j F_{0i}\rangle + \langle F_{ij}, \partial_i F_{0j}\rangle = \langle F_{ij}, \partial_0 F_{ij}\rangle$. Summing over ordered distinct nonzero i, j gives $2h = c$. Hence only n, h are needed. A signed vector is determined by its zeroth component, since exchanging axis zero with any axis transports that component. The polar vector space has real dimension at most two. These relations hold on all solution towers, so they are zero in the actual free quotient by (FM4).

The lower rank bounds are obtained by explicit separated correlations, not by declaring the formal expressions independent. Let $H_{ab,cd}(z) = C_{ab,cd}(z)$ be (FM2) and set

$$\mathcal{Q}_{ab,ik;ce,fm}(z) = H_{ab,ce}(z)H_{ik,fm}(z) + H_{ab,fm}(z)H_{ik,ce}(z). \tag{FC5}$$

The correlation of $Q_{ab,ik}$ with $Q_{ce,fm}$ at $z, 0$ is $(\dim \mathfrak{g})\mathcal{Q}(z)$, because the two cross matchings force the same color. At $z = e_t$, writing $\nu = (4\pi^2)^{-1}$, the Hessian and third derivative of the Green function are

$$\begin{aligned} \nu^{-1}\partial_j\partial_i G_0(e_t) &= -2\delta_{ij} + 8\delta_{ti}\delta_{tj}, \\ \nu^{-1}\partial_r\partial_j\partial_i G_0(e_t) &= 8\delta_{ij}\delta_{tr} - 48\delta_{ti}\delta_{tj}\delta_{tr} + 8(\delta_{ri}\delta_{tj} + \delta_{rj}\delta_{ti}). \end{aligned} \tag{FC6}$$

These follow by differentiating $|z|^{-2}$. Substituting them into the four signed Kronecker terms in (FM2), and applying the two-factor product rule to (FC5), gives every sample below. The common nonzero factor removed is $\ell_G = (\dim \mathfrak{g})\nu^2 > 0$.

Take the divergence component zero of (FC3), correlate it at e_0 with $Q_{01,01}$ and $Q_{12,12}$ at zero, divide by ℓ_G, and take real parts. The two resulting linear functionals have columns

	A	B	C
$Q_{01,01}$	-256	0	128
$Q_{12,12}$	0	-512	-128

. (FC7)

For example the first A entry is $\sum_{k=0}^{3} \partial_0 \mathcal{Q}_{0k,0k;01,01}(e_0)/\nu^2 = -256$; the corresponding B entry is the same sum over ordered $i, k \neq 0$, which is zero. For C sum $\partial_r \mathcal{Q}_{rk,0k;01,01}$ over $r \neq 0, k$. The other row uses $12, 12$. These finite sums and (FC6) give respectively the three displayed columns. A third, undifferentiated sample correlates tensor component 01 at e_0 with $Q_{02,12}$ at zero; its columns are $(0, 0, -16)$, directly from (FC5). It separates the last tensor direction.

Thus A, B, C are independent, so the tensor space has dimension exactly three. The two divergence samples in (FC7) are independent; combined with the upper bound two from (FC4), they identify the entire polar vector space and prove surjectivity of divergence. Rank-nullity gives a one-dimensional conserved kernel. Solving (FC7) yields its generator

$$\tfrac{1}{2}A - \tfrac{1}{4}B + C = T^{\mathrm{free}}. \tag{FC8}$$

For its diagonal entry use $\sum_{i,k} Q_{ik,ik} = 2N_a + R_a$; its off-diagonal entry is P_{ab}. This is exactly (FM6), already shown conserved. The last sample is nonzero on it, so the generator is nonzero. Every conserved centered real signed tensor in this weight-four space is its real scalar multiple. Any later nonzero scalar normalization, including either sign, preserves the rank and the conserved-line statements; the normalization of shell fluxes is a separate argument.

There is a raw degree-five vector direction which must be kept until the equation quotient is taken. In polynomial notation its components are

$$E_b = \sum_{i,a} \mathrm{pairDeriv}(b, i, a, a, i), \tag{FC9}$$

that is, the contraction of F_{bi} with the equation $\partial_a F_{ai}$, in the normalized radial-jet coordinates used above. The two other raw vector directions are the polynomials called momentum1 and

momentum2 in the preceding classification. Every real signed-equivariant raw vector has a unique expression $uV_1 + vV_2 + wE$. To see the independence after imposing the Maxwell equations, take a unit color vector and the two solution jets with respectively

$$(F^{(0)}, F^{(1)}) = (\mathrm{leg}_{12}, \mathrm{leg}_{0;12}), \quad (\mathrm{leg}_{01}, \mathrm{leg}_{0;01} - \mathrm{leg}_{2;21}). \tag{FC10}$$

Here the legs are radial potential jets, with the normalization made explicit below. Both satisfy the linear divergence equation and Bianchi identity. The first reads $V_1 = 0, V_2 = 4/3$ in component zero; the second reads $V_1 = 1$. An equation polynomial reads zero on both. It follows first that $v = 0$, then that $u = 0$. Conversely E vanishes on every solution jet, directly by the divergence equation. A non-solution jet with nonzero divergence shows that E itself is a nonzero raw polynomial. Thus the raw vector space has dimension three, and the kernel of its map to the free vector quotient is exactly the one-dimensional space $\mathbb{R}E$.

This statement is independent of coefficient coordinates. For a coherent degree-five frame, apply its inverse coordinate map to each component of a real signed-equivariant coefficient vector. This is a real linear isomorphism to the raw polynomial vector space, because the coordinate map intertwines signed permutations and conjugation. Composing with the free quotient projection gives a surjection whose kernel is the inverse-coordinate image of $\mathbb{R}E$. A section is obtained by applying the corrected polynomial lift to each quotient component. That lift commutes with both symmetries, and its quotient projection is the identity. These facts justify the vector kernel subtraction used later: it removes precisely the equation direction, and does not discard either of the two nonzero free vector classes.

To make the raw coordinates completely explicit, write

$$P_{ab;kcd} = \langle F_{ab}, \partial_k F_{cd} \rangle, \quad (V_1)_b = \sum_i P_{bi;bbi}, \quad (V_2)_b = \sum_{i,j} P_{ij;bij}, \quad E_b = \sum_{i,a} P_{bi;aai}. \tag{FC11}$$

The pairings are the actual invariant polynomial labels, evaluated on radial potential jets; imposing the Maxwell equation is a later quotient operation. The basic legs in (FC10) are the homogeneous potentials

$$(\mathrm{leg}_{ab}\, e)_i(z) = \tfrac{1}{2}(\delta_{ib} z_a - \delta_{ia} z_b) e, \qquad (\mathrm{leg}_{k;ab}\, e)_i(z) = \tfrac{1}{3}(\delta_{ib} z_a - \delta_{ia} z_b) z_k e. \tag{FC12}$$

Taking their exterior derivatives defines the curvature legs, including the factor $1/3$ in the first curvature derivative. This factor explains the $4/3$ sample in (FC10). A constant curvature is automatically divergence free. For a linear curvature the Maxwell condition is exactly $\sum_i \mathrm{curl}_1(i, i, b) = 0$ for every b; direct differentiation of (FC12) checks it for the two combinations in (FC10). Higher homogeneous jets can be set to zero.

Here is also the raw upper-dimension argument, before using those solution samples. A weight-five label decomposes as a constant, an invariant linear term, an FF term and an $F\,\partial F$ term. Simplicity kills the linear term as above. Total axis inversion fixes the two F's but negates a polar vector, so both its constant and FF terms vanish. Resolve its remaining invariant bilinear color form into the color inner product. To recover its spatial coefficients it suffices to use the mixed evaluations

$$\mathfrak{m}(A, B; v) = \mathrm{Re}\{v(A + B) - v(A) - v(B) + v(0)\}, \tag{FC13}$$

where A occupies jet order zero and B order one. Take, in component zero, the three pairs

$$(\mathrm{leg}_{01}\, e, \mathrm{leg}_{0;01}\, e), \quad (\mathrm{leg}_{01}\, e, \mathrm{leg}_{2;21}\, e), \quad (\mathrm{leg}_{21}\, e, \mathrm{leg}_{2;01}\, e).$$

If all three readings vanish, every mixed coefficient vanishes. Indeed reflections require the symmetric difference of its two plane-index sets to equal $\{0, k\}$, where k is the derivative index. For

$k = 0$, equal planes containing zero and equal planes avoiding zero are the two possible types. For $k \neq 0$, the planes share one other index, and the two placements of the derivative factor are the other types. Permutations fixing zero identify all coefficients of each type; antisymmetry fixes their signs. The three displayed tests kill three types. The cyclic identity $\mathrm{leg}_{0;12} + \mathrm{leg}_{1;20} + \mathrm{leg}_{2;01} = 0$, paired with $\mathrm{leg}_{12}\, e$, kills the fourth. The legs span each radial jet space by the expansion

$$A = \tfrac{1}{2} \sum_{i,j} \mathrm{leg}_{ij}(\mathrm{curl}_0(i,j)A), \qquad B = \tfrac{1}{2} \sum_{i,j,k} \mathrm{leg}_{k;ij}(\mathrm{curl}_1(k,i,j)B).$$

Bilinearity therefore kills the entire mixed part. Reality recovers the complex coefficient from its real part. The three evaluations form an injective real map, proving the raw upper bound three. The two solution evaluations in (FC10) and the nonsolution pair $(\mathrm{leg}_{01}\, e, \mathrm{leg}_{0;01}\, e)$, on which $E_0 = 1$, prove independence of V_1, V_2, E. Thus (FC11) is a basis and every raw vector has unique coordinates.

For later use, the divergence of the raw centered tensor with coefficients $w = (w_0, w_1, w_2)$, respectively diagonal mountain, diagonal total and off-diagonal shear, has these coordinates:

$$\mathrm{Div}(w) = 2(w_0 - w_2)V_1 + (2w_1 + \tfrac{1}{2}w_2)V_2 + w_2 E. \tag{FC14}$$

The product rule gives $2V_1$ for the first tensor and $2V_2$ for the second. For the shear use Bianchi and interchange the ordered indices to obtain $\sum_{a,i} P_{ai;abi} = \tfrac{1}{2}V_{2,b}$; deleting the $a = b$ terms subtracts $2V_{1,b}$, while the other product derivative is E_b. This proves (FC14). Its coordinate matrix is invertible, with inverse

$$w_2 = y_2, \qquad w_0 = y_0/2 + y_2, \qquad w_1 = y_1/2 - y_2/4.$$

It defines the actual scalar contraction on raw vectors:

$$\rho(\mathrm{Div}(w)) = -4(w_0 + 4w_1) = -2y_0 - 8y_1. \tag{FC15}$$

If $Q = \frac{1}{4}\sum_{i,j}\langle F_{ij}, F_{ij}\rangle$, the tensor trace is $4(w_0 + 4w_1)Q$; hence ρ reads minus that trace in Q-units. It kills E and the traceless conserved tensor $w = (1, -1/4, 1)$. Since (FC10) proves that the raw equation kernel is exactly $\mathbb{R}E$, ρ annihilates that entire kernel. This is the scalar sensor used later, with its sign and normalization fixed. The identification of the free quotient here uses the Maxwell polynomial evaluation and the reflected Gram norm proved in (FM3)–(FM4) and in the detector's reflected polynomial calculation; it is not an additional assumption of nondegeneracy.

10 Free shell charges and separation of the Ward kernel

The free shell calculation is the comparison model for physical Ward charges. Compact cutoffs ensure that all shell integrals avoid the insertions, and the fixed exterior degree horizon controls the tests used to separate defects. The argument does not assume that a low-degree test alone detects an arbitrary high-degree label; the displacement and Wick calculations below provide the additional separating probes.

We give the polynomial calculations underlying the free charge separation used in the rotation and trace arguments. This also specifies the exterior tests needed when a charge acts on high-degree fields.

10.1 The compact stress shell identity

Let χ be a smooth compactly supported function equal to one near a distinguished point x, whose support avoids all spectator points. For a current J, define its outward shell functional by

$$\mathcal{E}_\chi(J) = -\sum_a \int (\partial_a \chi)(y) J_a(y)\, d^4y. \tag{FS1}$$

The integrand is supported in a compact annulus away from all insertions. All exchanges of finite Wick sums and this integral are consequently justified by smoothness there. A one-leg source p at x produces the Maxwell pulse obtained by applying its jet functional to the covariance. For a smooth source-free Maxwell field f on a neighborhood of the support of χ, Green's formula for the bilinear Maxwell stress current gives

$$\mathcal{E}_\chi J_v(\mathrm{pulse}_{p,x}, f) = -\,\mathrm{sample}_x(p, \delta_v f). \tag{FS2}$$

Indeed the divergence vanishes off the source by the Maxwell and Bianchi equations. Integrating the differentiated fundamental covariance against the cutoff transfers the derivatives to the smooth field. The fundamental-solution identity leaves precisely its jet at x. The boundary term at infinity is absent because the cutoff is compact. The variation of a two-form is $\delta_v f_{ab} = v^c \partial_c f_{ab} + (\partial_a v^c) f_{cb} + (\partial_b v^c) f_{ac}$. For the three vector fields used here this variation preserves the Maxwell equations. The stress is symmetric and traceless, so its contraction with the symmetric derivative of a constant, skew-linear, or radial vector field contributes zero to the divergence.

Here is a regularized derivation of the fundamental-solution step in (FS2). With the covariance normalization $c_4 = 1/(4\pi^2)$, put

$$g_t(y) = \frac{c_4}{t^2 + |y|^2}, \qquad \rho_t(y) = \frac{8c_4 t^2}{(t^2 + |y|^2)^3}, \qquad t > 0. \tag{FS2a}$$

Direct differentiation gives $-\Delta g_t = \rho_t$. The radial integral gives $\int \rho_1 = 1$: the sphere area is $2\pi^2$ and $\int_0^\infty r^3(1+r^2)^{-3} dr = 1/4$. Also $t^4 \rho_t(ty) = \rho_1(y)$. Hence for every continuous compactly supported h, $\int \rho_t(y-x)h(y)dy \to h(x)$, by the change of variable $y = x + tz$ and the integrable majorant $\rho_1(z)\|h\|_\infty$.

For a plane source il, form the potential $A_b = \delta_{ib}\partial_l g_t - \delta_{lb}\partial_i g_t$, and take its exterior derivative. It satisfies the Bianchi identity by commutation of derivatives, while its divergence is $\delta_{lb}\partial_i \rho_t - \delta_{ib}\partial_l \rho_t$. The divergence of its bilinear stress with the source-free field f is therefore the contraction of this expression with f; the terms from the symmetric derivative of v cancel by tracelessness and the other terms cancel by skew symmetry. Integration by parts against χ transfers the source derivatives to the smooth compact test made from χ, v, f. The preceding approximate-identity limit evaluates their difference at x, giving $-\delta_v f_{il}(x)$. On the shell the regulated exterior derivative converges smoothly to the pulse, since that compact shell avoids x. Its flux thus converges to the left side of (FS2). This proves the plane-source case with the stated sign.

Differentiate the source center in any direction while keeping the cutoff and vector field fixed. All differentiated shell integrands have a common bound on a fixed compact annulus, so differentiation passes through the integral. A center derivative raises the source leg by that direction; on the right it differentiates the sample of $\delta_v f$. Induction gives (FS2) for all derivative legs; linearity gives every solution-jet functional. The wave property of $\delta_v f$ for the three motions ensures these samples respect the Maxwell jet relations.

For a constant vector e_i, (FS2) gives minus the derivative-leg sample. For $v(y) = A(y-x)$, with A skew, it gives minus the infinitesimal tensor rotation. For $v(y) = y - x$, a homogeneous solution

jet of derivative order j has Euler contribution j, and the two form indices contribute two more. Thus the one-leg radial charge is $-(j+2)$ times that jet. This specifies the engineering weight, including the two contributed by curvature rather than its derivatives.

When both pulses are centered at x, use cutoff independence and homogeneity instead of multiplying distributions at their common pole. Two cutoffs equal to one near x have a difference supported away from x; source-free integration by parts makes their fluxes equal. Replace the cutoff by its dilation by a factor two about x. A jet-order-j pulse scales by 2^{-4-j}; the other scales by 2^{-4-l}; the volume and cutoff derivative together contribute 2^3, and a motion homogeneous of degree $e = 0$ or 1 contributes 2^e. Thus the common flux Z equals $2^{e-5-j-l}Z$. That coefficient is strictly smaller than one, so $Z = 0$. When neither pulse is centered at x, their fields are smooth throughout the support of the cutoff, and the source-free divergence identity gives zero. When precisely one pulse is centered there and the other is at a spectator point, (FS2) gives minus the covariance with the transformed distinguished leg.

Now insert the quadratic Wick stress in an arbitrary Wick monomial word. Double Wick extraction expresses its correlation as one half the sum over ordered choices of two occurrences, times the bilinear pulse current for those occurrences and the Wick pairing of all remaining occurrences. There are exactly the three cases just considered. The terms with both distinguished occurrences or with neither distinguished occurrence vanish. In each remaining term exactly one occurrence of the distinguished monomial is transformed. The two orders of the selected pair cancel the one half. Summing over its occurrences is the Leibniz rule for the derivation on that monomial. By multilinearity this proves, for arbitrary polynomial words,

$$\mathcal{E}_\chi\left(\sum_b v_b(y)\langle T_{ab}(y)P(x)\prod_\ell P_\ell(x_\ell)\rangle\right) = -\left\langle (\mathcal{D}_v P)(x)\prod_\ell P_\ell(x_\ell)\right\rangle. \tag{FS3}$$

Here $\mathcal{D}_v$ is respectively translation, rotation, or the engineering-weight derivation. This proves the three shell laws, including their common minus sign, and includes every contraction term. In particular conservation of the free stress alone was not being used as a substitute for its charge normalization.

10.2 Low-degree vector charges

Let $\mathscr{S}_j$ be the real space of homogeneous order-j Maxwell solution jets and $\mathscr{V}_j = \mathscr{S}_j^*$. Their elements are the real one-leg test functionals. A variable at order j has engineering weight $j+2$. For $p \in \mathscr{V}_j$, write $B_c(p)$ for its linear polynomial in color c, and set

$$[p,q] = \sum_c B_c(p)B_c(q).$$

This bracket denotes a color-contracted quadratic polynomial, not a Lie bracket. For a linear map $D : \mathscr{S}_j \to \mathscr{V}_l$, its polynomial is

$$\mathcal{C}(D) = \sum_\nu [e^\nu, De_\nu], \tag{FS4}$$

where (e_ν) is any basis and (e^ν) its dual. This is basis-independent, and $\mathcal{C}(p\otimes q) = [p,q]$. If $j \neq l$, assigning independent jets in these two orders in one nonzero color proves that $\mathcal{C}$ is injective.

Consider a vector row T_i on the invariant polynomial filtration, killing the unit, increasing filtration by at most one, real under conjugation, and signed-covariant. Suppose it is silent, meaning

$$\sum_\ell \langle P_1(x_1)\cdots(T_iP_\ell)(x_\ell)\cdots P_n(x_n)\rangle = 0 \tag{FS5}$$

for every allowed exterior word and separated configuration. At input weight four the invariant space is spanned by the unit and $[p, q]$ with $p, q \in \mathscr{V}_0$. At output weight at most five the only further possibility is a color-contracted pair of orders zero and one. Total inversion is even on weight four and odd on such a pair. The vector law anticommutes with total inversion, and therefore removes the weight-four part of the output. Reality permits real coefficients in (FS4). Consequently

$$T_i[p, p] = \mathcal{C}(D_i(p)), \qquad D_i(p) : \mathscr{S}_0 \to \mathscr{V}_1. \tag{FS6}$$

Let $K_z(p, q)$ denote the real single-color covariance between two test legs at displacement $z \neq 0$. To extract the constraint on (FS6), put two quadratic insertions at tu and zero, and a third at a fixed nonzero point, and multiply (FS5) by t^5. The covariance of jet orders zero and one is homogeneous of degree -5. The triangle terms where the transformed order-one leg is on either near insertion have a finite limit; the term with that leg at the distant insertion tends to zero. The double contractions contribute $4 \dim \mathfrak{g}$. Choose the distant one-leg functional so that its covariance with a zero-order test is evaluation on a prescribed $X \in \mathscr{S}_0$. Cancelling the common nonzero factor gives

$$K_u(D_i(p)X, q)\, q(X) = K_u(D_i(q)X, p)\, p(X). \tag{FS7}$$

The availability of these evaluations follows from covariance separation of solution jets. First suppose a test functional of order j has zero covariance with every order-zero plane test at every nonzero displacement. Its pulse is then zero off its source, in every color component. Pair its radial stress shell with the homogeneous polynomial Maxwell solution having any prescribed order-j solution jet. The shell integral is zero because the pulse is zero there; (FS2) makes it $-(j + 2)$ times the test evaluated on that arbitrary jet. The functional is zero. For partner order $l + 1$, vanishing against all such partners makes the radial derivative of its covariance with any order-l partner zero. That covariance is homogeneous of degree $-(4 + j + l)$; Euler's rule makes the radial derivative that nonzero scalar times the covariance. Thus it vanishes, and induction reduces to order zero. Finite-dimensional duality now makes these covariance evaluations span the required full dual. Thus no inverse of a covariance matrix at a single momentum is being assumed. This also proves that the solution-jet space is nonzero at every order: otherwise separation against that order would force the nonzero plane evaluation at order zero to vanish.

If $p(X) = 0$, (FS7) first makes $K_u(D_i(p)X, q) = 0$ for $q(X) \neq 0$. Given $q(X) = 0$, add a test with nonzero value at X, apply that result to the sum, and subtract. Covariance separation gives $D_i(p)X = 0$. Hence $\ker p \subseteq \ker D_i(p)$. Choose X_0 with $p(X_0) \neq 0$; subtracting $p(X)X_0/p(X_0)$ from X proves

$$D_i(p) = p \otimes f_i(p), \qquad T_i[p, p] = [p, f_i(p)]. \tag{FS8}$$

Set $f_i(0) = 0$. For two nonzero linear functionals there is an X on which both are nonzero: use a witness for one, a witness for the other, and their sum if needed. Substitution in (FS7) yields

$$K_z(f_i(p), q) = K_z(f_i(q), p). \tag{FS9}$$

Pair the difference $f_i(p + q) - f_i(p) - f_i(q)$ with an arbitrary one-leg test, use (FS9), and cancel by linearity in its other slot. Separation makes the difference zero. The same argument with a scalar proves homogeneity. Thus $f_i : \mathscr{V}_0 \to \mathscr{V}_1$ is real linear. Signed covariance of T_i, injectivity of (FS4), and cancellation of a nonzero test functional give simultaneous signed covariance of the family f_i.

We classify this family explicitly. Write c_{ab} for evaluation of a zero jet in the ab-plane and $s_{iab} = \partial_i c_{ab}$ for the corresponding first-jet functional. Under reflection of axis z, the latter has character equal to the product of the three signs belonging to i, a, b. Project onto a reflection character by the four commuting projections $(1 \pm R_z)/2$. The character of $f_0(c_{01})$ leaves only

$s_{001}, s_{212}, s_{313}$. Average the permutation interchanging axes two and three. The Maxwell relation gives $s_{212} + s_{313} = s_{001}$, so this fixed subspace is the line spanned by s_{001}. Similarly the character of $f_0(c_{12})$ leaves only $s_{012}, s_{102}, s_{201}$. Antisymmetrize the interchange of axes one and two; the Bianchi relation $s_{102} - s_{201} = s_{012}$ leaves the line spanned by s_{012}. Hence these two values have scalar coefficients a, b.

Use (FS9) with $p = c_{01}$, $q = c_{12}$, and $z = e_0 + e_2$. The explicit differentiated covariance (FM2), evaluated below after (WE4), gives the same nonzero value $-2c_4$ for s_{001}, c_{12} and s_{012}, c_{01}. Therefore $a = b$. Permutations fixing zero give all remaining plane evaluations, which span $\mathscr{V}_0$; permutations moving zero give the other directions. We conclude

$$f_i = a\partial_i, \qquad T_i[p,p] = a[p, \partial_i p]. \tag{FS10}$$

Since $\partial_i[p,p] = 2[p, \partial_i p]$, the scalar of the translation derivation is $a/2$. Squares span the symmetric quadratic polynomials by polarization, and both rows kill the unit. This proves the entire degree-four assertion, with its factor two, rather than just a classification of invariant labels.

10.3 Fixed-displacement evaluation and the triangular Wick calculation

Here are the three explicit evaluations used in (FS7)–(FS10). Put $c_4 = (4\pi^2)^{-1}$, and let $H_z = I - 2zz^{\mathsf{T}}/|z|^2$ for $z \neq 0$. Its action on two-forms is the solution action $\rho(H_z)$. If E_{ab} is the unit antisymmetric plane tensor, the differentiated covariance (FM2) gives

$$K_z(p, c_{ab}) = \kappa_z\, p(\rho(H_z)E_{ab}), \qquad \kappa_z = 4c_4|z|^{-4} > 0. \tag{WE1}$$

Indeed substitute the second derivatives of $c_4|z|^{-2}$ into the two curls: their radial terms combine to the product of the two Householder matrix entries on the antisymmetric indices, with the displayed common radial factor. This is an identity on the six-dimensional real space of antisymmetric order-zero solution tensors.

Given $X \in \mathscr{S}_0$, set $W = \rho(H_z)^{-1}X$. Antisymmetry gives the plane expansion $W = \frac{1}{2}\sum_{a,b} c_{ab}(W)E_{ab}$. Define the one fixed functional

$$w_{z,X} = \kappa_z^{-1}\frac{1}{2}\sum_{a,b} c_{ab}(W)c_{ab}.$$

Bilinearity and (WE1) now give

$$K_z(p, w_{z,X}) = \frac{1}{2}\sum_{a,b} c_{ab}(W)p(\rho(H_z)E_{ab}) = p(X). \tag{WE2}$$

Thus the evaluation witness used at displacement $-z$ in (FS7) exists at that displacement itself. No change of the distant point or sum over different displacements is necessary.

Write $\mathcal{T}_{x,y,z}(A,B,C)$ for the three-point Wick moment. At each point its two letters are Wick ordered, so a contraction cannot join two letters at that same point. In a pairing of the six letters, each point must therefore have one leg going to each of the other two points. For letters a, v at the first point, two equal letters p, p at the second, and q, q at the third, there are two choices of which first-point letter goes to the second point. For either choice there are two choices of its partner among the equal p's and two choices of the partner among the equal q's. The third edge is then forced. Hence

$$\begin{aligned}\mathcal{T}_{x,y,z}([a,v],[p,p],[q,q]) = 4\dim\mathfrak{g}\,\{&K_{x-y}(a,p)K_{x-z}(v,q)\\ &+ K_{x-z}(a,q)K_{x-y}(v,p)\}K_{y-z}(p,q).\end{aligned} \tag{WE3}$$

The color factor is exactly $\dim \mathfrak{g}$: the three edges give Kronecker deltas requiring all three summed colors to coincide. All unequal-color summands are zero and each common color gives the same real scalar expression.

Let $x(t), y(t) \to 0$, $x(t) - y(t) = tu$, $u \neq 0$, with the third point $z \neq 0$ fixed. The near covariance of two order-zero legs scales as t^{-4}, while that of orders one and zero scales as t^{-5}. Multiplication of (WE3) by t^5 therefore gives

$$4 \dim \mathfrak{g} \, \{tK_u(a,p)K_{x(t)-z}(v,q) + K_{x(t)-z}(a,q)K_u(v,p)\}K_{y(t)-z}(p,q). \tag{WE4}$$

Every covariance involving the distant point is continuous in this limit. The first term tends to zero; the second tends to $4 \dim \mathfrak{g} \, K_{-z}(a,q)K_u(v,p)K_{-z}(p,q)$. If the order-one letter is instead at the distant point, its two edges stay bounded and the only short edge is an order-zero covariance. Multiplication by t^5 leaves an extra factor t, so that whole term tends to zero. Expand the panel in (FS6) in a basis as in (FS4), use (WE2) for the distant test, and apply these two limits to the three silent terms. The common nonzero factor $4 \dim \mathfrak{g}$ cancels; the odd sign of the order-one covariance under $u \mapsto -u$ gives exactly (FS7).

Finally take $p = c_{01}$, $q = c_{12}$. Substitution in (FM2) gives, at every $z \neq 0$,

$$K_z(c_{01}, c_{12}) = 8c_4 z_0 z_2 (|z|^2)^{-3}.$$

Differentiation of the first leg in direction zero gives

$$K_z(s_{001}, c_{12}) = 8c_4 \{z_2(|z|^2)^{-3} - 6z_0^2 z_2 (|z|^2)^{-4}\}.$$

At $z = e_0 + e_2$, its squared norm is two and both displayed coordinates are one. The value is $8c_4(1/8 - 6/16) = -2c_4 \neq 0$. The order-zero covariance is symmetric in its two plane tests; differentiating that symmetry gives the same value for $K_z(s_{012}, c_{01})$. Thus (FS9) reads $-2c_4 a = -2c_4 b$, proving $a = b$ in (FS10).

10.4 Detection with a fixed exterior degree bound

We first establish the collision extraction that will be used in the polynomial argument. Put $R_a(x) = |x|^{2a}$. Say that F is cleared at order a when $R_a F$, initially defined off zero, extends to a real analytic function on all of $\mathbb{R}^4$. A harmonic pole of degree r is analytic off zero, homogeneous of degree $-r$ under positive dilations, harmonic there, and cleared at order $r - 1$. The differentiated single-leg covariance of jet orders k, j is such a pole with $r = 4 + k + j$. A cleared homogeneous term of degree j means a function homogeneous of degree $-j$ cleared at order $\max(j - 2, 0)$. A product of poles of positive degrees a, b is a cleared homogeneous term of degree $a + b$, because the product of their analytic numerators clears it at order $(a - 1) + (b - 1)$.

The pole assertion for the differentiated kernels has a short direct proof. The Green function $c_4|x|^{-2}$ is homogeneous of degree -2; multiplying it by $|x|^2$ leaves the constant c_4. Away from zero its Laplacian vanishes, since the radial Laplacian in dimension four gives $f''(r) + 3f'(r)/r = 0$ for $f(r) = r^{-2}$. Suppose F is a pole of degree $r \geq 2$ with numerator $g = R_{r-1}F$. For a fixed direction v, differentiation gives the numerator identity

$$R_r D_v F = R_1 D_v g - (r - 1)(D_v R_1) g. \tag{HC0}$$

Its right side is defined with the same regularity on all real space. Differentiating the dilation identity makes $D_v F$ homogeneous of degree $-(r + 1)$, and commuting ordinary derivatives with the Laplacian shows it is harmonic off zero. Thus differentiation raises the pole degree by one and

retains the required clearing order. Finite linear combinations of equal-degree poles retain all these properties. The two curls in (FM2) are a finite combination of second derivatives of the Green function, so have degree four. The two jet orders add $k + j$ derivatives, with the factor $(-1)^j$ for derivatives at the second endpoint. Color pairings and arbitrary probe coordinates are fixed linear maps. This proves exactly the degree $4 + k + j$ and clearing hypothesis used above. It also gives a global bound $|F(x)| \le C|x|^{-r}$: bound the continuous function on the unit sphere and use positive dilation homogeneity. All uses here are off the real collision point; no contact value or complex cone division is involved.

For a harmonic homogeneous F of degree $-r$, Euler's identity $x \cdot \nabla F = -rF$, together with the product rule for the four-dimensional Laplacian, gives

$$\Delta(R_{a+1}F) = 4(a+1)(a+2-r)R_aF. \tag{HC1}$$

For example $\Delta(|x|^2F) = 8F + 4x\cdot\nabla F$; the general formula follows either by the same differentiation or induction. If $r \ge 2$ and the pole is cleared already at order $r-2$, it is zero. Indeed, starting with its analytic numerator at that order, (HC1) repeatedly produces an analytic numerator one order lower. At each step $a \le r-3$, so the displayed coefficient is nonzero. At the last step F itself extends analytically. Its homogeneity then implies $F(x) = t^rF(tx) \to 0$ as $t \downarrow 0$. For $r = 2$, clearing at order zero gives this conclusion immediately.

Here is the exact finite expansion needed; it includes terms more singular than the pole being extracted. For integers r, M, write

$$t^r\left(F(tx) - \sum_{j<M} W_j(tx)\right) \longrightarrow K(x) \quad (x \ne 0), \tag{HC2}$$

where each W_j is a cleared homogeneous term of degree j, and K is a harmonic pole of degree r. If $F(tx) = 0$ for all sufficiently small positive t, then $K = 0$, for $r \ge 2$. To prove this, peel the finite sum from its largest degree. A nonzero term of degree $J > r$ would diverge after scaling by t^r; multiplying the convergent expression by t^{J-r} forces $W_J = 0$. Repeat until only degrees at most r remain. Degrees below r tend to zero; the limit is the degree-r term. Thus $-K = W_r$, or is zero if that degree is absent. In either case K is cleared at order $r-2$, and (HC1) makes it zero. This finite peeling, rather than a choice of a large scaling exponent, removes the more singular double contractions.

To apply this statement, let P contain jets through order j, let a pair of test legs of orders k, l, with $l < k$, be inserted at $x_0 + z$, and keep the other vertices separated. The contraction of an order-k probe p with the order-a variables of P is the differential operator

$$B_a(p,z)P = \sum_\nu K_z(p, e_{a,\nu})\,\partial_{a,\nu}P, \qquad B(p,z)P = \sum_{a=0}^{j} B_a(p,z)P. \tag{HC3}$$

The index ν includes the jet basis and color. Wick's rule applies the two operators to every ordered pair of vertices. Set $r = 4 + k + j$ and $M = 8 + k + l + 2j + 1$. There are four cases.

1. Both legs contract at x_0. The resulting finite sum consists of second partial derivatives of P times products of two kernels. Their degrees are $8 + k + l + a + b < M$, with $a, b \le j$. They are the cleared terms W in (HC2), with harmonic part zero.

2. The order-k leg contracts at x_0, and the order-l leg at another vertex. The distant contraction is analytic at $z = 0$. Multiplication by t^r selects only $B_j(p,z)$; the terms $a < j$ have smaller pole degree and vanish. The distant contraction tends to its value at $z = 0$.

3. The order-l leg contracts at x_0, and the other leg is distant. Its largest pole degree is $4+l+j < r$, so the scaled limit is zero.

4. Both vertices are distant. Every coefficient is regular at $z = 0$, so the scaled limit is zero.

These conclusions require only continuity of the distant Wick coefficients for the scaled limit; the coefficients are in fact analytic. Summing the finitely many placements proves (HC2). In case 2, summing the distant placements is Wick's first-leg identity at x_0. Summing colors in the invariant two-leg tester therefore gives harmonic part

$$K(z) = \langle (\mathcal{C}_j(p,q;z)P)(x_0) \prod_i V_i(x_i) \rangle, \quad \mathcal{C}_j(p,q;z)P = \sum_c B(q,c)\, B_j(p \otimes c, z)P. \tag{HC4}$$

Here $B(q,c)$ is the order-l, color-c jet variable. Orthogonality of the color action makes the sum invariant. The operator replaces a variable of weight $j+2$ by one of weight $l+2$, so it strictly lowers engineering weight whenever $l < j$.

If P is a ghost at exterior cap d, and $4 + k + l \le d$, the inserted invariant pair is an allowed spectator. For every fixed nonzero direction z, all the inserted points stay separated for sufficiently small positive t. Permuting the inserted pair past P, ghost vanishing makes the left side before subtraction in (HC2) zero. Finite peeling and (HC1) now give $K(z) = 0$ for every $z \neq 0$. Since the spectator word was arbitrary, $\mathcal{C}_j(p,q;z)P$ is itself a ghost at the same cap. This is the cluster extraction used below; it retains arbitrary spectator arity while keeping each spectator's degree bounded.

The passage to arbitrary input weight needs a uniform form of polynomial separation. An exterior bound d limits the weight of each spectator field, while allowing arbitrarily many spectators. Call a polynomial a ghost at bound d if all its separated correlations with all such spectator words vanish. We prove that one finite d, depending on the group, detects polynomials of every input weight. This is stronger than a Gram statement at a fixed input cap.

First, separated single-leg covariances with fixed partner jet order span the dual of each homogeneous solution-jet space. Their annihilator is zero by the differentiated covariance separation used in (FS7), and finite-dimensional duality proves the span assertion. Consequently if every single contraction of a polynomial with those tests is zero, all its partial derivatives in the corresponding jet variables are zero. Indeed a single contraction is the sum of those partial derivatives multiplied by the covariance coordinates; linear combinations of the spanning covectors isolate each derivative.

Next we must separate colors using invariant tests. Put $q = \dim \mathfrak{g}$. For every color index $c = 0, \dots, q-1$, choose a nonzero functional $p_c \in \mathscr{V}_c$ and a solution jet X_c with $p_c(X_c) = 1$. Such jets exist at every order. Form the polynomial matrix

$$M_{cb} = B_b(p_c). \tag{FS11}$$

Its determinant is not the zero polynomial: assign the order-c jet to color c with value X_c, and all other color components at that order to zero. This evaluation makes M the identity. The argument takes place in the polynomial ring, which is an integral domain. It does not assert that M is invertible at every field configuration.

Suppose a polynomial has no jet variables above order $j+1$. Extract a mixed cluster consisting of one highest-order contracted leg and one of the low-order color tests just chosen. Vanishing of these extracted correlations says $MD = 0$, where D_b is the single contraction in color b with the highest-order variables. Multiply by the adjugate: $\det(M) D_b = 0$. Since the determinant is a nonzero polynomial, $D_b = 0$. The preceding dual-span argument makes every highest-order partial derivative zero. Over characteristic zero a variable occurring with a positive power has a nonzero

partial derivative, so the polynomial contains no variable at that order. This removes the highest jet order one layer at a time.

We spell out the bounded-degree and induction aspects of the cluster extraction. Set $j = q$, $L = j+1$. The polynomial ring on jets through order j has finitely many variables and is Noetherian. Form its ideal generated by all positive polynomial-length homogeneous invariant polynomials. Choose finitely many such generators of this ideal, and let r bound their engineering weights. They generate every invariant polynomial as an algebra: for an invariant polynomial of length $k > 0$, express it in the ideal, take its length-k component, and average the coefficients over the compact group. Averaging is a projection onto invariant polynomials and satisfies $\mathcal{R}(hg) = \mathcal{R}(h)g$ for invariant g. For a generator of positive length l, its coefficient after this projection is the averaged length-$k-l$ part of the old coefficient; it has length strictly smaller than k. Induction on length expresses it as products of the same generators. Constants are empty products. Thus every invariant polynomial homogeneous of length k on this shelf lies in the span of products $\prod_i V_i$, with each V_i invariant of engineering weight at most r, homogeneous of length k_i, and $\sum_i k_i = k$. This is the finite hosting bound; the number of factors is unrestricted.

Here is the precise coalescence argument that uses it. Suppose a ghost P has all length components below k zero. Fix $x \neq y$ and a hosted product $Q = \prod_{i=0}^{n-1} V_i$ of total length k. Place P at x and V_i at $y + tie_0$, for $t > 0$ sufficiently small. These points are distinct and avoid x. Expand all polynomials into monomials and apply Wick's rule. A monomial of P of length $l > k$ has too many letters to pair with the k exterior letters: it has no internal pair at x, so its moment is exactly zero. If $l = k$, every exterior letter must pair with a letter at x, leaving no exterior-to-exterior pair. All the resulting kernel arguments approach the nonzero vector $x - y$; their finite products therefore converge. Terms $l < k$ are absent by hypothesis. The limit is exactly the two-point Wick pairing $\mathsf{B}_{x,y}(P_k, Q)$, where P_k is the length-k component. Ghost vanishing at every positive separated t makes this limit zero. Linearity extends it from products to their span. No estimate of kernels between coalescing exterior points is needed, since the pairing count removes those kernels before taking the limit.

Induct on k. Let P_k be the first possibly nonzero length component and reflect its complex conjugate across the time hyperplane, giving $Q = \overline{P_k}^{\,\mathrm{ref}}$. The shelf, invariance, engineering cap and length are preserved by this operation. The hosting result puts Q in the allowed span. At points $-e_0, e_0$ its reflected Wick pairing with P_k is strictly positive if $P_k \neq 0$, by the polynomial Gram calculation above. Coalescence has made the same pairing zero, a contradiction. Thus $P_k = 0$. Every polynomial has only finitely many length components, so this induction proves that a ghost lying on this fixed shelf is zero. Define

$$d = \max\{r, 4 + j + L\}. \tag{FS12}$$

Induct on the engineering weight of the candidate ghost. If its highest jet order is $J > j$, cluster a pair of jet tests of orders L and $l < q$. Its weight is $4 + L + l \leq d$, and $l < L$, $l < J$. Apply (HC2)–(HC4) with highest shelf order J: the finite cutoff is $8 + L + l + 2J + 1$, and the extracted pole degree is $4 + L + J$. The extracted mixed contraction is a ghost of strictly smaller engineering weight, hence is zero by induction. The color-matrix argument then removes the highest jet layer. Repeating leaves a polynomial on the shelf through j, where the hosting tests separate it. The weight-zero case is the constant, detected by the empty word. This proves

$$\exists d\ \forall w\ \forall P \text{ of weight at most } w, \quad P \text{ a ghost at bound } d \implies P = 0. \tag{FS13}$$

The same finite bound is used before the finite-dimensional Ward readers are selected. In particular one takes their cap at least $\max(d, 4)$; lower-cap conclusions then follow by the filtered

inclusions. Arbitrarily high output weight in (FS13) is permitted, but arbitrarily high individual spectator degree has not been silently introduced.

10.5 Elimination of higher-degree vector and angular defects

Subtract from a silent vector row the scalar translation obtained in (FS10), and denote the residual row by Q_i. It is silent, kills every degree-four field, and in particular kills the free stress. Insert that stress as an additional spectator in (FS5), apply its radial shell identity (FS3), and subtract the silent identity with the engineering-weight derivative applied to the distinguished field. The terms where the row acts on a spectator cancel in pairs. The stress term is zero. What remains says that

$$[E, Q_i]P \quad \text{is a ghost}, \qquad E = \sum_{\nu} w_\nu X_\nu \partial_{X_\nu}, \quad w_\nu = \text{jet order of } X_\nu + 2. \tag{FS14}$$

All additional stress fields have degree four and the spectator cap is the fixed one in (FS13). Thus (FS13) gives $[E, Q_i]P = 0$.

Every polynomial is a finite sum of engineering-weight components P_k, and $EP_k = kP_k$. A map commuting with E preserves each such eigenspace: projecting its image onto weight $l \neq k$ gives $(l - k)$ times that projection equal to zero. Total inversion acts on a weight-k monomial by $(-1)^k$, because a derivative-order-j curvature variable has sign $(-1)^j = (-1)^{j+2}$. But the vector covariance law gives $RQ_i = -Q_i R$. On a preserved weight component these two rules give simultaneously $RQ_iP_k = (-1)^k Q_iP_k$ and $RQ_iP_k = -(-1)^k Q_iP_k$. Hence $Q_iP_k = 0$, and $Q_i = 0$ on the whole cap. This proves the full one-dimensional free translation-charge kernel used after (RT17).

For the internal angular row the degree-four quarter-turn argument in (RT18)–(RT19) already makes it zero on the free stress. Apply the angular version of (FS3) in the same silent identity. It makes the commutator with each angular generator a ghost, and (FS13) makes that commutator zero. On every finite filtered space, exponentiating a finite matrix gives the unique solution of its linear differential equation. Commutation with the generator therefore gives commutation with its quarter-turn. Signed covariance of the two alternating Lie-algebra slots and the plane-basis calculation following (RT19) then make every component zero at that cap. This supplies the higher-degree angular separation with the required bounded exterior tests.

Finally the analytic segment argument following (RT19) converts the collinear scene readings and all their position derivatives into the full silent identity (FS5). Taking the cap in (FS12)–(FS13), the preceding arguments identify its kernel. Finite-dimensional decreasing-kernel selection gives a finite list with exactly that kernel. Its nonzero minors remain nonzero for the convergent interacting reading matrices. This proves the rank transfer used to normalize the actual translation charge and eliminate the actual internal angular defect.

10.6 Obtaining a detector on the physical coefficient space

The detector used below is finite dimensional. At fixed degree let E be the coefficient space, I its equation ideal, and r the dimension of the free-field quotient E/I. The foundation identifies the kernel of the physical projection π_i with an invertible image of I. Rank-nullity therefore gives $\dim \operatorname{ran} \pi_i = r$, uniformly in the cutoffs. In the notation of the detector trials $r = \dim + 1$. This dimension statement does not assert that every bare coefficient survives the equations.

Choose r trained trial vectors $e_{j,i}$, uniformly bounded by v, and let the detector matrix consist of their tested responses. The limiting response matrix has full column rank. In finite dimensions the least norm on its unit sphere is positive; compactness of that sphere makes the lower bound

stable under sufficiently small entrywise perturbations. Denote the allowed perturbation by $\delta > 0$ and the resulting left-inverse bound by L. Train with tolerance $\epsilon = \delta/4$. The actual tested trial vector differs from its limiting response by at most two errors of size ϵ: the finite training error and the short-distance probe error.

There is a third error when $e_{j,i}$ is replaced by $\pi_i e_{j,i}$. Indeed $e_{j,i} - \pi_i e_{j,i}$ has exactly zero projection. The equation-cancellation bound makes its probe tend to zero, hence eventually smaller than ϵ, for all finitely many j simultaneously. Thus the projected trial response matrix is within $3\epsilon < \delta$ of the full-rank reference matrix. Its columns are independent. Since $\operatorname{ran} \pi_i$ has dimension at most r, the projected trials span that range.

The projection has a uniform operator bound A, so each spanning vector has norm at most Av. Apply the stable left inverse to the detector response of any vector in the range and expand in these r spanning vectors. Summing their norms gives

$$\|\pi_i x\| \le ArvL\, \|L_i \pi_i x\|.$$

The choices are ordered: first choose the training index sufficiently large, then a sufficiently small positive probe scale, and only then pass sufficiently far out in the whole-mesh filter. Constants are independent of the remaining cutoffs and volume. The exact equation-cancellation input is obtained by applying the banked insertion estimate to the identified kernel; its cost is the vanishing flight of exponent (-1/32). The banked estimate is proved in (IC3)–(IC5); the separated reflected Wick pairing below constructs the full-rank reference matrix.

The stability assertion has an explicit elementary proof. Write $G : \mathbb{C}^r \to \mathbb{C}^s$ for the reference response matrix, with sup norms, and choose $K > 0$ such that $|u| \le K|Gu|$. Set $\delta = [2K(r+1)]^{-1}$. If every column of $H - G$ has norm at most δ, then

$$\|(H-G)u\| \le r\delta\|u\|, \qquad \|u\| \le K\|Hu\| + Kr\delta\|u\| \le K\|Hu\| + \tfrac{1}{2}\|u\|.$$

Hence $|u| \le 2K|Hu|$; one can take $L = 2K$. The injectivity of the reference matrix follows directly from separation of the free quotient by the chosen correlation functionals: a vector annihilated by all rows is zero. Expressing these functionals in a basis identifies their matrix with G. This explains both the rank lower bound and its stability on precisely the same varying projected trial columns used in the rank upper bound.

10.7 Bounding physical coefficients before taking limits

Here is the normalization argument that upgrades the available insertion estimates to estimates for lattice differences. Write x_i for the coefficient vector representing a difference insertion, and $X_i = pi_i x_i$ for its physical projection. The presentation is idempotent: $pi_i^2 = pi_i$. The detector construction supplies linear maps L_i, a constant $Z \ge 0$, and, eventually in the mesh limit,

$$\|X_i\| \le Z\|L_i X_i\|.$$

An insertion already bounded in the Wilson laws, together with a vanishing error in the tested insertion seminorms, gives $|L_i x_i| \le c$ eventually. This alone does not bound X_i, because the equation part $x_i - X_i$ need not have small coefficient norm. Instead put

$$f_i = (1 + \|X_i\|)^{-1}, \qquad y_i = f_i X_i, \qquad z_i = f_i(x_i - X_i).$$

Then $|y_i| \le 1$, $0 < f_i \le 1$, and z_i belongs to the same admissible coefficient class: that class is closed under bounded scalar multiplication and subtraction of bounded coefficient families.

Idempotence gives $pi_i z_i = 0$ exactly. The equation-cancellation estimate in the literal Wilson law therefore makes the tested observables associated with z_i tend to zero. In particular, eventually $Z|L_i z_i| \le \frac{1}{2}$. By linearity,

$$Z\|L_i(X_i - x_i)\| \le \tfrac{1}{2}(1 + \|X_i\|).$$

The detector inequality and the triangle inequality now give

$$\|X_i\| \le Zc + \tfrac{1}{2}(1 + \|X_i\|), \qquad \|X_i\| \le 2Zc + 1.$$

Thus the projected coefficients have a fixed bound before any limit is taken. Their primary insertions consequently have fixed cost. The unprojected remainder has zero physical projection, so the equation-cancellation estimate gives it a vanishing cost, hence also a fixed cost. Subtracting this remainder and the original vanishing approximation error proves the fixed-cost estimate for the difference insertion itself. This argument uses cancellation for the actual number of copies in the detector; it does not replace equation terms by a coefficient-norm estimate.

For completeness, the algebraic passage from differences of correlations to these insertions is exact. In a Wilson integral of a product of slot observables, a lattice difference in the position of slot j acts only on the j-th factor. Finite subtraction commutes with integration, since the factors are continuous and integrable on the compact configuration space. Iterating this identity expresses every prescribed mixed lattice difference as a Wilson integral of the corresponding shifted and differenced slot observables. The preceding fixed-cost argument therefore yields the stated bounds for those differences once the detector and cancellation estimates have been supplied.

10.8 Equation cancellation in the physical detector

The physical detector must annihilate the calibrated equation image exactly. The proof therefore corrects complete cylinders, including all lower jets read by the detector, rather than prescribing a single leading coefficient and ignoring its induced lower terms. Equivariance and reality are then restored by averaging operations that preserve this exact cancellation.

We supply the cancellation input used in the detector argument. At fixed grade d, let I_d be the free equation ideal, and let R_d be its marked lift into local lattice insertions. The marking identity says that an insertion evaluated on $R_d y$, $y \in I_d$, has zero Wilson expectation against any bounded measurable spectator depending outside a fixed finite neighborhood, once the lattice is large enough. Both the neighborhood radius and the minimum size can be chosen uniformly over $y \in I_d$: choose a finite basis of this ideal, take the maxima of its basis radii and size thresholds, and use linearity and finite summation. Multiplication by arbitrary complex coefficients does not enlarge the neighborhood.

This local identity makes the replicated root seminorm exactly zero. Fix a coarse reporting level k. Eventually a fine-scale chamber contains the common marking neighborhood strictly inside it, because the ratio of fine to reporting spacing tends to zero. Tile the fine torus by the prescribed chambers and put one translated copy of the insertion in each chamber. Keep the distinguished copy in chamber zero, and multiply all other copies to form the spectator. Distinct chamber interiors keep that spectator outside the marking neighborhood. Its measurability and boundedness follow from the local insertion bounds on the compact finite configuration space. The marking identity then gives

$$\mathbb{E}_\beta \prod_{\text{chambers } a} f_a = 0.$$

The distinguished copy is exactly the original local equation insertion, including its scale normalization. Taking the absolute-value root defining the root seminorm gives zero. The same conclusion holds on the whole mesh and for every volume, with the reporting level fixed before taking the fine tail.

To transfer this exact cancellation through the trained coefficient charts, write Z_i for the terminal chart at cutoff i and P_i for its calibrated invertible coefficient change. The physical projection has kernel equal to the image of I_d under the composite marked chart. For a coefficient family x_i in this kernel define

$$y_i = P_i^{-1} Z_i^{-1} x_i.$$

Then $Z_i y_i$ belongs to the marked equation carrier. The calibrated inverse estimate bounds the normalized coefficient size of y_i by twice that of x_i. Hence a coefficient family with the permitted subexponential growth retains that growth after this inverse change.

The trained-chart comparison is

$$\operatorname{upper}(Z_i P_i y_i) - \operatorname{hallmarks}(y_i) = O_{\mathrm{root}}(\operatorname{flight}_{-1/32}(i)). \tag{IC3}$$

Its finite-dimensional derivation is explicit. Decompose y_i using the fixed finite signed-orbit tally of frame basis vectors. For each tally component the trained local comparison has root bound $\operatorname{flight}_{-1/16}(i)$. Its scalar coefficient is bounded by the coefficient norm times the fixed orbit-coordinate constant. Sum over the finite tally. The sum of these scalar bounds is subexponential, so it is eventually at most $\operatorname{flight}_{-1/32}(i)/\operatorname{flight}_{-1/16}(i)$. This proves (IC3). The algebraic recombination is exact: the trained basis expansion is $P_i y_i$, whereas the original frame expansion recombines to y_i. Signed translations and conjugations commute with the local insertion by its covariance.

The hallmark term in (IC3) has zero root seminorm by the chamber argument. Since $Z_i P_i y_i = x_i$, we obtain

$$\operatorname{upper}(x_i) = O_{\mathrm{root}}(\operatorname{flight}_{-1/32}(i)) \quad \text{when } x_i \in \ker \pi_i. \tag{IC4}$$

This is the banked equation-cancellation estimate. In the detector proof apply it to $x_i - \pi_i x_i$, using $\pi_i^2 = \pi_i$. The projected trial vectors and their unprojected versions then have the same limiting tested responses. Their common rank is the free quotient rank: the equation carrier is an invertible image of I_d, and the free observable map has kernel precisely I_d. Rank-nullity gives this dimension statement without assuming independence of all bare coefficients.

The reference response matrix is injective on that free quotient. Each row is a free correlation functional: insert the variable quotient field in the first slot and the prescribed fixed basis fields in its spectator slots. Multilinearity makes it a linear functional. The separating family of these tests has zero common kernel, so its map into the finite row space is injective. Composing it with a quotient basis gives exactly the reference matrix used earlier. Thus (IC4) and the finite-dimensional stability calculation together apply to the same varying projected trial columns, completing the detector's cancellation step.

The local marking identity underlying (IC4) originates in Haar integration by parts. Vary one stored positive link e by $U_e \mapsto \exp(tA)U_e$, with all other links fixed, and denote this derivative by $\nabla_{e,A}$. The change preserves product Haar measure. The Wilson action derivative is

$$\nabla_{e,A} \mathcal{S}(U) = \langle A, \mathcal{E}_e(U) \rangle.$$

Here $\mathcal{E}_e$ is the sum of the incident plaquette curvatures, with a backward incidence transported by the inverse preceding link. More explicitly a plaquette type with axes $i < j$ contributes, at a link of axis i, the difference

$$F_{ij}(x) - \operatorname{Ad}_{U^{-1}_{x-e_j,j}} F_{ij}(x - e_j),$$

and at a link of axis j the negative of the analogous difference in direction i. Other plaquette types contribute zero. There is one contribution for each positive-axis plaquette type. Differentiating its four ordered factors gives this formula, including the inverse-link sign. The normalization follows from

$$-\operatorname{Re}\operatorname{tr}(\operatorname{colorMat}(A)\operatorname{mat}(g)) = \operatorname{repDim}\langle A, \operatorname{curvature}(g)\rangle.$$

This identity uses skew-adjointness of the color matrix and the orthogonal projection defining curvature.

Let W and $\nabla_{e,A}W$ be continuous, and let S be bounded measurable and unchanged by this link variation. Haar invariance gives

$$\mathbb{E}_\beta\left[(\nabla_{e,A}W - \beta\langle A, \mathcal{E}_e\rangle W)S\right] = 0. \tag{IC5}$$

To justify differentiation, translate the derivative formula from zero to any t by the one-parameter group law. On the compact configuration space the continuous derivative of the Wilson-weighted W is bounded; multiply its bound by the bounded envelope of S. This is an integrable uniform bound for difference quotients by the fundamental theorem of calculus. Differentiation under the Haar integral is therefore valid. The integral is constant by invariance, so its derivative is zero. Dividing by the positive partition function proves (IC5). The spectator need not be continuous or gauge invariant. With $W = 1$, $\beta \neq 0$, it gives the exact equation expectation zero.

The marking construction retains the derivative term when a multiplier uses the varied link. Its local equation identity means zero expectation against every bounded measurable spectator depending outside some lattice radius r, on every torus of side at least a fixed M, at every nonzero real β. The class of continuous local functions with this property is closed under finite sums, scalar multiplication and conjugation: enlarge the radii and side thresholds to common maxima; for conjugation use the conjugate spectator and reality of the Wilson law. These operations are exactly those used to assemble the equation section.

One must also arrange the desired continuum label while enforcing these exact identities. The finite graded construction supplies a linear pair

$$p \longmapsto (Qp, f_p),$$

where f_p is a gauge-invariant local polynomial insertion with leading jet Qp. At every graded piece, $Qp - p$ lies strictly above that piece in the finite lexicographic weight ordering. The construction retains the exact marking identity on the equation ideal. Triangularity makes Q invertible: a nonzero vector has a first nonzero graded piece, which Q leaves unchanged, so $Qp = 0$ is impossible. In finite dimensions injectivity implies surjectivity. The same first-piece argument shows that Q^{-1} preserves each descending weight filtration.

Define a linear local realization by $Y(p) = f_{Q^{-1}p}$. It has leading jet p and retains the required power filtration. Real conjugation and signed coordinate symmetry are imposed on the paired construction, so Q, its inverse, and Y obey these symmetries, and the unit is fixed. There is one finite stencil realizing all its frame basis vectors: choose their finitely many polynomial presentations, enclose all their underlying stencils in a common centered box, and rewrite them as polynomials in the common box coordinates. Define the section on arbitrary coefficients by linear extension from this basis. Equality with Y on the basis proves equality everywhere. The box contains zero and is invariant under signed coordinate permutations.

With $R = Q$, the resulting fixed section satisfies

$$T(\operatorname{coords}(Rp)) = f_p.$$

In particular the equation ideal is sent to the exact local marking class, as required above. The same section supplies its finite derivative-credit constants by compact polynomial bounds. The section and R are chosen once at grade d, before any block side, period, coupling, refinement or later derivative budget is specified. This explains why the local equation cancellation survives all subsequent scheduling choices.

The graded linear lift just used is assembled as follows. For each weight pair m, let E_m be its homogeneous label space and $J_m = E_m \cap I_d$. The ordinary polynomial section supplies a linear lift on E_m with the correct leading form and power bound. On J_m, choose a finite basis and lift each basis vector by the positioned equation construction. Its output is a pair (q, f) with $q - p$ of higher weight, the prescribed continuum form, the same power bound, gauge invariance, and the exact local marking identity. Extend these finitely many pairs linearly on J_m; all these properties are linear, including membership in the marking class.

Choose an algebraic projection $L_m : E_m \to J_m$. Define the lift on E_m by

$$F_m(p) = F_m^{\mathrm{eq}}(L_m p) + F_m^{\mathrm{ord}}(p - L_m p).$$

For $p \in J_m$ the second term is zero. For arbitrary $p \in E_m$ both summands have the required grade and higher-weight difference, so their sum has leading form p. At the zero weight, $J_{(0,0)} = 0$, since the equation ideal contains no nonzero unit, so this construction sends the unit to $(1, 1)$. The whole lift is the finite sum $F(p) = \sum_m F_m(p_m)$ over the homogeneous components of p. The equation ideal is preserved by taking these components, so its entire image has the marking property. This proves the linear assembly used in the triangular argument.

Its equivariant version involves only finite averages. Let θ_s act on labels by a signed coordinate transformation, and let

$$\phi_s(q, f) = (\theta_s q, f \circ \mathrm{pull}_s).$$

Replace F by

$$F_{\mathrm{av}} = |\mathcal{S}|^{-1} \sum_{s \in \mathcal{S}} \phi_s F \theta_{s^{-1}}.$$

Each summand has the same leading label and higher-weight remainder, because the signed action preserves grades and the equation ideal. The local marking class is preserved by this coordinate change. The averaging coefficients sum to one, so the leading label and the unit stay fixed. Reindexing by left multiplication in $\mathcal{S}$ proves signed equivariance.

Finally let $\mathfrak{c}(q, f) = (\overline{q}, \overline{f})$, with the prescribed semilinear real structure on labels. Set

$$F_{\mathbb{R}}(p) = \tfrac{1}{2} \left(F_{\mathrm{av}}(p) + \mathfrak{c} F_{\mathrm{av}}(\overline{p}) \right).$$

The composition of the two semilinear maps in the second term is complex linear. The real structure commutes with the signed action and preserves every grade, higher-weight space, form and marking identity. Thus this midpoint is both signed-equivariant and real, retains the unit, and still has the triangular leading-label map. These averages are performed before constructing the common finite stencil or choosing any cutoff family.

Gauge invariance of a positioned equation lift is obtained by averaging its constant-color orbit, while retaining a finite polynomial realization. If f is a powered local polynomial, its orbit $f_g(U) = f(\mathrm{gauge}_g U)$ spans a finite-dimensional space: each factor uses a fixed finite matrix representation and finite polynomial degree, and conjugation acts linearly on its finitely many coefficient tensors. The marking property is preserved by g: in its expectation change variables by the constant gauge transformation, use invariance of the Wilson law, and pull the spectator back by the inverse transformation. Its locality radius and bounded measurable envelope are unchanged.

The paired vectors $(1, f_g)$ therefore lie in a finite-dimensional space. Their first coordinate and all configuration-evaluation coordinates are continuous functions of g. These functionals separate that space, so their span is the whole dual; hence every linear coordinate is continuous. Integrate the coordinates of $(1, f_g)$ against probability Haar measure to get a vector in their span. Its first coordinate is one and its second coordinate is exactly

$$\bar{f}(U) = \int_G f(\mathrm{gauge}_g U)\, dg.$$

Membership in this span expresses the pair as a finite linear combination of orbit pairs. The leading-form, higher-weight, power and marking properties are all preserved by those finite linear combinations, with the first coordinate keeping the leading label equal to the original p. Thus the average is still represented by an actual finite polynomial; no infinite-dimensional integral closure is needed.

For these transported local polynomials, a general gauge transformation acts through its value at the root. Right invariance of Haar measure therefore proves full gauge invariance of $\bar{f}$. The polynomial leading form is color-invariant as well: compute its constant-color transformed form by equivariance, then use uniqueness of the leading-form polynomial to identify it with the original. This puts that form in the label space and proves the gauge-invariant version of the positioned lift used above.

The generation of the equation ideal in this construction is ordinary polynomial algebra. On the degree-j jet space let P_j be a linear projection onto the range of the injective free gauge-jet map. Substituting P_j in every jet variable defines an algebra homomorphism $\mathcal{P}$. A label polynomial belongs to the equation ideal precisely when $\mathcal{P}p = 0$: restriction to the free solution range is zero if and only if the polynomial after this projection is zero on every jet. Polynomial evaluation separates polynomials.

For a linear jet functional ℓ, split

$$\ell = (\ell - \ell P_j) + \ell P_j.$$

The first part vanishes on free solution jets, the second is its projected part, and both have the same bidegree as ℓ. For a product of linear functionals, telescope the difference between that product and the product of their projected parts. Every resulting term contains one literal factor vanishing on free solution jets. Expand an arbitrary polynomial into its finitely many monomials. If $\mathcal{P}p = 0$, this telescoping writes p as a finite linear combination of such terms. When p is homogeneous, every term has its same bidegree, since the projection preserves each jet degree. It also follows that each homogeneous component of an equation polynomial is an equation polynomial. A polynomial of field-count degree zero is a constant; its restriction to free solution jets is that constant, so the equation ideal has no nonzero such element.

To lift a term $\ell_{\mathrm{eq}}\ell_1 \cdots \ell_r$, the source first represents the equation functional ℓ_{eq} as a linear combination of linearized transported lattice forces. For each individual force, it constructs the other factors on stencils avoiding that force's varied edge. Their product is therefore unchanged by that edge's one-parameter variation, and (IC5) makes the resulting force product an exact local equation insertion. Different force summands may use different avoiding stencils; they are combined only after their individual identities have been obtained.

The leading-jet bookkeeping keeps all linear terms through the retained degree. Each avoiding ordinary factor agrees with its desired linear functional modulo terms of field-count degree at least two. Their product agrees with the desired r-factor monomial modulo field-count degree at least $r + 1$, by the finite product expansion. Multiply by the force factor, truncate to the retained jet

degree, and the error has field-count degree at least $r+2$, strictly above the target count $r+1$. The force's retained linear functional may have higher jet weights than ℓ_{eq}; these are also retained and lie strictly above its target weight. Multiplication by the other factors preserves this strict inequality. Adding the finitely many force terms therefore gives a polynomial with target leading term and only higher-weight remainder, while the actual lattice function remains exactly locally silent. Finite linear combination and the ideal generation just proved yield the positioned raw lift for every homogeneous equation label. This accounts for the triangular rather than exactly homogeneous form of its retained continuum polynomial.

The force-span assertion is a finite-dimensional annihilator argument with a concrete kernel calculation. Linearize a force transported from a lattice source to the root along a path avoiding that source edge. At linear order its value is a finite sum of transverse differences of the sampled curvature jets:

$$\sum_{\ell=0}^{d-2} \sum_{a<b} \mathrm{incidence}_{\nu,a,b} \left(F_{ab}^{(\ell)}(x) - F_{ab}^{(\ell)}(x-e_i) \right).$$

The transport path drops out at this order because the curvature vanishes at the identity. The incidence signs are the same positive-link signs as in (IC5).

The leading homogeneous content of a transverse difference is its ordinary derivative. More precisely the difference of the two face averages equals a transverse interval integral of the derivative of their integrands. Together with the two face integrals, it is an integral over three unit intervals. This is the fundamental theorem of calculus, before any limiting argument. Rescale its base point and mesh by t; for a homogeneous jet of order ℓ the rescaled transverse expression is t^{ℓ} times its unit-mesh expression, after multiplication by the extra transverse factor t.

Suppose all transported force linearizations vanish on a finite jet array $A_0, \dots, A_j$. For each source address choose a path to it avoiding its outgoing source edge. The preceding formula gives vanishing of the sum of transverse differences at every integer address. Multiply the order-ℓ terms by $t^{j-\ell}$ and use the rescaling identity. At points $t\,\mathrm{index}_t(x)$, which tend to any prescribed $x \in \mathbb{R}^4$, the resulting continuous expression vanishes. Let $t \downarrow 0$. All terms except the top order vanish, and the three interval averages shrink to a point. The remaining equation is

$$\sum_i \partial_i F_{i\nu}^{(j)}(x) = 0 \quad \text{for every } x, \nu.$$

Its Bianchi identity already holds because these are curvature jets of radial potentials. Thus A_j is a free solution jet. If a functional on that top jet annihilates the free gauge-jet range, it annihilates the common kernel of all force probes. In finite dimensions the annihilator of that common kernel is the span of the probes. This proves the force-span assertion and yields an actual finite linear combination, not an infinite sum over source points.

The edge-avoiding ordinary probes span the whole finite jet dual. Fix one forbidden edge. Except at a finite set of curvature-sampling addresses, one can reach the sampling plaquette along a path that avoids this edge and whose plaquette boundary avoids it as well. If every such safe sample vanishes on a jet array, all its curvature polynomials vanish at every integer point outside that finite exceptional set. The same scaling argument isolates its top homogeneous part: for nonzero x, the sampled addresses $\mathrm{index}_t(x)$ eventually leave every fixed finite set, since their physical rescalings tend to $x \neq 0$, while fixed addresses rescale to zero. Continuity gives vanishing at $x = 0$ too. Hence the top homogeneous curvature jet is zero. Descending induction gives zero for every jet component. Radial reconstruction is injective, so the original jet array is zero. The safe samples therefore have zero common kernel; finite-dimensional duality proves that every desired ordinary jet functional is a finite linear combination of them.

Finally a transported force times an edge-avoiding multiplier has the exact marking identity claimed above. Move its root color test through the carrier by adjoint invariance:

$$\langle X, \mathrm{Ad}_g \mathcal{E}_e \rangle = \langle \mathrm{Ad}_{g^{-1}} X, \mathcal{E}_e \rangle.$$

The carrier avoids e, so its transported color vector is unchanged under variation of e. Expand that vector in a fixed real color basis. In each term its scalar coordinate, the avoiding multiplier, and the remote spectator are all unchanged by the varied edge. Apply the $W = 1$ case of (IC5) to each such product and sum over the finite basis. The spectator locality radius contains the fixed force edge; a sufficiently large side prevents its periodic identification with a forbidden remote edge. This proves the exact local marking property of each force term, completing the positioned equation-lift argument.

10.9 Uniform placement and testing of detector trials

Detecting a coefficient with actual correlations A detector is a finite list of linear readings of a distinguished insertion against fixed spectator insertions and tested positions. To use its inverse on the whole relevant coefficient space, the same placement and continuity bounds must work for every column in a fixed coefficient ball. The proof below obtains that uniformity from one separated compact neighborhood and a finite coordinate telescoping argument. Its conclusion concerns the actual whole lattice cylinders, including their corrections, rather than only their leading Taylor polynomials.

The detector probes are actual lattice correlations, linear in the tested first insertion. Fix the trained spectator insertions $F_1, \dots, F_n$. At cutoff $i = (N, p)$, define

$$\mathcal{M}_i(f; x) = \mathbb{E}_{\beta_i} \left[f(x_0) \prod_{\nu=1}^{n} F_\nu(x_\nu) \right].$$

Linearity follows by expanding the first factor under the finite Wilson probability integral; all local polynomial insertions are continuous on the compact configuration space. A forward lattice difference in the first insertion equals the corresponding difference in its first position. Iterating proves this identity for every finite direction list. The detector applies the corresponding finite averaging operator to $\mathcal{M}_i$ at its prescribed placement. It is therefore a complex linear map on the first coefficient vector.

The trained placements lie a positive distance inside a common physical box and are generic: every pair of points differs in every coordinate used by the separation convention. For each fixed row and training index choose a small closed neighborhood of the placement inside this generic set and inside $|x_{\nu,a}| < 1$. Its pair-coordinate gaps have a positive minimum by compactness. These gaps persist for every allowed torus volume. Indeed all physical periods are at least four; a difference of absolute value less than two is no closer to a nonzero period translate than it is to zero. Thus the same positive gap works for every volume, including the smallest admitted torus.

The unextracted insertion estimates consequently give a common local continuity bound for all coefficient columns in a fixed norm ball. To derive it, connect two nearby lattice configurations of positions by replacing their $4(n+1)$ coordinates one at a time. Every intermediate replacement stays within their original physical distance and within the chosen neighborhood. The one-coordinate telescoping estimate gives, if the physical distance is R, a correlation change at most $A^{n+1}R/r$ for each replacement, where $r > 0$ is the common separation reserve and A the insertion allowance. Summing gives

$$|\mathcal{M}_i(f; x) - \mathcal{M}_i(f; y)| \le 4(n+1)A^{n+1}R/r. \tag{IC6}$$

The bounds for reflected labels are uniform as well, since their finite-dimensional real-linear action maps a fixed coefficient ball to a bounded set.

A detector average with probe scale $b > 0$ samples points displaced by at most $2b$ per direction in its fixed direction list. Equation (IC6) makes its difference from the unaveraged monitor arbitrarily small by taking $b \downarrow 0$, uniformly on the subsequent whole-mesh tail and over every volume. The lattice placements have exactly the prescribed physical positions, because the refinement ratio satisfies $a_N \operatorname{count}(N, k) = a_k$. There is no additional moving-position error at that step.

For the trained columns, the unaveraged monitor tends to its free reference response in the order: training index first, then the whole mesh. Split its error into the finite-level training error and the free hierarchy's evaluation at the approximate rather than target generic placement. The former tends to zero by the trained insertion estimates. The latter tends to zero by smoothness of the free hierarchy on the distinct-position configuration space. There are only finitely many rows and columns, so all their tails can be intersected.

Take the training error below $\delta/4$, then take b so that the averaging error is below $\delta/4$, and finally use (IC4) so that projecting each trial changes its response by less than $\delta/4$. This recovers the three errors used in the detector argument, with one common mesh tail and no extraction needed to establish rank. The trial coefficient bounds and projection norm bounds are already uniform over N, p, so the resulting detector inequality has exactly the uniformity required before passage to continuum limits.

The training errors used above are bounds on complete corrected cylinders. Introduce commuting square-zero markers $\epsilon_1, \dots, \epsilon_n$. Let

$$F(\epsilon) = \sum_{A \subset \{1,\dots,n\}} S_{\mathrm{free}}(x_A; q_A)\epsilon_A, \qquad D = \log F,$$

where the empty moment is one and $\epsilon_A = \prod_{a \in A} \epsilon_a$. The logarithm and exponential are finite polynomials in the nilpotent marker ideal. For actual lattice insertions P_a, form

$$H(\epsilon, U) = e^{-D(\epsilon)} \prod_a (1 + \epsilon_a P_a(U)).$$

The training bound controls every nonempty coefficient of H in the replicated root seminorm by a common $\mathrm{gap}_k \to 0$, not just its singleton coefficients.

The finite identity $FH = \prod_a (1 + \epsilon_a P_a)$ gives, coefficient by coefficient,

$$\prod_{a \in A} P_a = F_A + \sum_{B \subsetneq A} F_B H_{A \setminus B}.$$

Thus the Wilson moment error from the free moment at the actual rescaled placement is bounded by the finite absolute coefficient price of F times gap_k. This price stays bounded as the placements tend to the fixed distinct configuration: every subconfiguration stays distinct and every free submoment is continuous there. This proves the training-error convergence used after (IC6).

In particular the singleton coefficient is $P_a - F_{\{a\}}$, and for two distinct slots the pair coefficient is

$$(P_a - F_{\{a\}})(P_b - F_{\{b\}}) - (F_{\{a,b\}} - F_{\{a\}} F_{\{b\}}).$$

The detector rows use at most these two slots. The pair correction is therefore its actual free cumulant at the rescaled placement; it cannot be replaced by a product of separately centered one-point terms.

All trial coefficient limits can be taken on a common refinement. For each fixed grade and training index the finitely many lifted coefficients lie in a fixed compact coefficient ball, uniformly

in the two cutoffs. The set of triples consisting of grade, training index and column index is countable. Apply successive compact subsequence extraction to these triples and use the diagonal refinement of the two cutoff orders. This preserves all previous hierarchy and projection limits. Replacing the trial's placeholder limiting columns by the limits so obtained changes none of its finite-cutoff estimates or placements. Passing the corrected-cylinder moment bound through the fixed-placement hierarchy limit proves convergence of every limiting trial response to the same free reference matrix. Its stable rank, together with the equation dimension bound, makes these limiting projected columns span the continuum physical quotient. The same trial data are retained when null labels are removed, so the cleaned hierarchy has exactly the free quotient dimension at every grade.

The detector's free columns can all be taken from the same reference frame, and one separated point pair suffices. The images of the frame basis span the free quotient, since the coefficient projection is onto. Its unit image is nonzero. Extend that unit to a basis by selecting further vectors from these finitely many spanning images, and enumerate the basis with the unit first. The original basis is real and homogeneous, and projection preserves conjugation and each bidegree. The selected quotient basis $b_0, \dots, b_m$ therefore has all these properties; no orthonormality at a probe separation is required.

For any $x \neq y$, let $B_{x,y}(p, q)$ be their free two-point Wick pairing. The free quotient identifies injectively with its polynomial realization on Maxwell solution jets. Reflect every tensor index of the conjugate polynomial p through the hyperplane perpendicular to $x - y$; call the resulting quotient field $p^\sharp$. The reflected Wick positivity estimate gives

$$B_{x,y}(p, p^\sharp) > 0 \quad (p \neq 0).$$

Hence the map

$$p \longmapsto (B_{x,y}(p, b_0), \dots, B_{x,y}(p, b_m))$$

is injective. If all its entries were zero, bilinearity and the basis expansion of $p^\sharp$ would make its positive pairing zero. The same positivity shows that the free all-spectator null space has no residual beyond the equation ideal: any nonzero solution polynomial is detected by this particular reflected partner at one pair.

Choose $x = 0$, $y = (1, 1, 1, 1)$. This pair differs in every coordinate and is therefore a generic detector placement. There is one row per basis partner and exactly one spectator in every row. Its response matrix is the injective two-point map just displayed. This specifies the finite free detector used in (IC6) and the subsequent rank argument.

The reflected positivity used here holds jointly for different jet orders, not just separately on each order. Fix a nonzero separation z, write $r(z) = \|z\|^{-2}$, and let R_z be reflection perpendicular to z. A real solution-tensor probe P of homogeneous order j has the Kelvin response

$$\operatorname{reply}(P)(y) = s_j r(y)^2 R_y^{\otimes 2} P(\operatorname{inv}(y)), \qquad s_j = 4(j+1)(j+2)(-2)^j \text{ normal}.$$

Tensor indices and arguments transform together. Translate this polynomial probe in its inverted argument by $-\operatorname{inv}(z)$ and decompose it into homogeneous degrees $0, \dots, j$. Multiply its order-k component by s_j/s_k; denote these lower-order probes by $T_{jk}(z)P$. The diagonal component is P.

The sum of their responses is consequently

$$r(y)^2 R_y^{\otimes 2} s_j P(\operatorname{inv}(y) - \operatorname{inv}(z)).$$

Along $y = z + tv$, the inversion difference is $tV(z, v, t)$, where V is smooth near zero and $V(z, v, 0) = r(z) R_z v$. Homogeneity shows that this response has a factor t^j. Thus all its normalized Taylor jets

at z below order j vanish, and its order-j jet is

$$s_j r(z)^{j+2} R_z P.$$

Polarization turns the equality along every line into equality of the symmetric derivative tensors. Pairing with a reflected probe introduces the derivative sign $(-1)^j$. Hence the transformed reflected Gram form, on a tower of orders $0, \dots, n-1$, is

$$\sum_{j=0}^{n-1} c_j(z) \langle P_j, Q_j \rangle_{\text{tensor}}, \qquad c_j(z) = (-1)^j s_j r(z)^{j+2} > 0.$$

The change of probes is triangular with identity diagonal. Alternatively its displayed positive Gram identity immediately proves injectivity, and finite dimension gives invertibility. The original simultaneous reflected form is therefore positive definite on the entire tower.

Choose a real orthogonal basis $e_1, \dots, e_M$ for this positive form and let $d_i > 0$ be its squared norms. Every finite Wick polynomial is a polynomial in such a finite tower. The cross-insertion Wick rule gives

$$\langle e^\alpha, e^\beta \rangle_{\text{refl}} = \mathbf{1}_{\alpha=\beta} \prod_{i=1}^{M} \alpha_i!\, d_i^{\alpha_i}.$$

Indeed remove one factor on the left. Only factors of the same index on the right can pair with it, giving $\beta_i d_i$; induction proves the formula, starting with the empty monomial. These are Wick products at each insertion, so within-insertion contractions are absent. Therefore for a polynomial $P = \sum_\alpha p_\alpha e^\alpha$,

$$\langle P, P \rangle_{\text{refl}} = \sum_\alpha |p_\alpha|^2 \prod_i \alpha_i! d_i^{\alpha_i} > 0 \quad (P \neq 0).$$

This proves reflected positivity, including arbitrary mixtures of engineering weights and field-count degrees and the constant monomial.

For completeness the Kelvin response follows from differentiation of the scalar Green function with the source's normalized Taylor tensors. For a symmetric order-j tensor P, define

$$\mathcal{D}_P f(x) = \sum_{\nu_1, \dots, \nu_j} \frac{\partial_{\nu_1} \cdots \partial_{\nu_j} f(x)}{j!} P(e_{\nu_1}, \dots, e_{\nu_j}).$$

A solution tensor S obeys the trace constraints corresponding to harmonicity. With $G(x) = \text{normal } r(x)$,

$$\mathcal{D}_{S_{im}} G(x) = (-2)^j \text{ normal } r(x)^{j+1} S_{im}(x, \dots, x). \tag{IC7}$$

Induct on j. The order-zero identity is the definition of G. The normalized derivative recursion is $\mathcal{D}_P f = (j+1)^{-1} \sum_a \partial_a \mathcal{D}_{P(e_a, \cdot)} f$. Use $\partial_a r^k = -2k r^{k+1} x_a$. The derivatives falling on the tensor polynomial sum to its frozen trace and vanish. The remaining sum resolves $\sum_a x_a e_a = x$; division by $j+1$ cancels that factor from the derivative of r^{j+1}, leaving the next factor -2. This proves (IC7).

Put $V_i(x) = \sum_a x_a S_{ai}(x, \dots, x)$. The Bianchi identity and Euler's homogeneous identity give

$$\partial_i V_m - \partial_m V_i = (j+2) S_{im}(x, \dots, x).$$

Indeed differentiating x_a contributes twice S_{im} by antisymmetry; the remaining difference is $\sum_a x_a \partial_a S_{im} = j S_{im}$. Applying (IC7) and the vanishing divergence of S gives

$$\sum_a \partial_a \mathcal{D}_{S_{ai}} G = -2(j+1)(-2)^j \text{ normal } r^{j+2} V_i.$$

The contraction of the two-curvature kernel with S is twice the curl of this last vector, with the order dictated by the opposite-side derivative. Differentiating once more yields

$$s_j r^{j+2} \left[S_{im} - 2r(x)(x_i V_m - x_m V_i) \right].$$

The bracket equals reflection of both antisymmetric form indices through R_x; its term quadratic in the two radial contractions vanishes by antisymmetry. Since S is homogeneous of degree j, replacing its argument x by $\mathrm{inv}(x) = r(x)x$ converts the prefactor r^{j+2} into r^2. This is precisely the Kelvin response used above, with $s_j = 4(j+1)(j+2)(-2)^j$ normal. The sign $(-1)^k$ in an order-k probe at the second point comes from differentiating $x - y$ there; it is included once. All these identities hold off the origin, where the Green function is smooth.

10.10 Selecting the physical detector trials

The trials used in the detector argument can be selected on the physical history already constructed. Here is the selection, with its quantifiers. Fix a finite detector panel, its real reference representatives, and its positive comparison width b. At arbitrarily late comparison indices k, the nonlinear construction supplies its metering bounds, the restarted representations and their admissible paths, the one- and two-insertion comparisons, and the marked estimates. For every sufficiently large refinement index $N \geq k$, these data are available for every volume parameter p. The representation and the readout may depend on N, p; the comparison index and its bounds do not. Only insertion arities at most two are needed, since each panel row has at most one additional insertion.

The finite label transport is chosen once for all rows at these k, N, p. It is a linear map on the finite label space with the previously proved filtration and reality properties. Its inverse transports the reference representatives to a real tuple $x_{N,p}$, with its zeroth member equal to the distinguished unit. The coefficient bound for this inverse gives

$$\sup_{N,p,\nu} \|x_{N,p,\nu}\| \leq D_k < \infty. \tag{IC8}$$

To make this statement include the finitely many refinements preceding the availability threshold, put every member of the tuple equal to the unit there, and enlarge D_k to include its norm. No estimate at those exceptional refinements will be used.

The comparisons at the middle of the admissible scale interval propagate to its starting width along the actual representations at every intermediate scale. On this interval the reference coupling lies between one half and one times its value at k. Thus the same metering bounds apply throughout. The conditioned insertion estimate, followed by the path propagation estimate, bounds each requested correlation by its assigned bill at the starting width. Its error is the common shipping budget $\varepsilon(k)$, which tends to zero. This is a statement about the transported physical insertions: it requires neither an assertion of a stronger density comparison nor a newly chosen scale history. There are finitely many panel rows and coefficient requests, so all their eventual thresholds can be met together.

For completeness, the passage from arbitrarily late indices to a sequence makes all the geometric limits simultaneous. Write $w(k)$ for the starting width and $g(k) > 0$ for the testing scale. Along admissible indices,

$$w(k) \longrightarrow \infty, \qquad g(k) \longrightarrow 0, \qquad \varepsilon(k) \longrightarrow 0.$$

The first assertion follows from the depth property of the scale calendar: if an admissible width stayed below a fixed j, its plateau identity would contradict the eventual strict increase of the

calendar beyond the depth at j. Given integers i, j and $\eta > 0$, intersect the arbitrarily late supply set with the eventual conditions

$$k > i, \quad w(k) > j, \quad g(k) < \eta,$$

the banking bound, the reference-grade threshold, and the finitely many geometric fit conditions. This intersection is nonempty. Recursively select k_{n+1} from it with

$$k_{n+1} > k_n, \qquad w(k_{n+1}) > w(k_n), \qquad g(k_{n+1}) < g(k_n).$$

The resulting indices and widths are cofinal, the positive testing scales decrease to zero, and the error budgets tend to zero. At each selected index choose the bounded real lifts in (IC8).

The trial marks are the prescribed lattice approximations to the panel's fixed physical configurations. Their fit bounds place them within the fixed quarter box, give their required separation and stencil admissibility, and imply convergence to those configurations. For each selected trial index, all sufficiently fine N, and every p, apply the physical conditioned estimate to these same lifts and marks. The stencil realization identity identifies the resulting correlations with the trial matrix entries. Consequently each entry has the required corrected comparison bound with error $\varepsilon(k_n)$, uniformly in p. This supplies the actual trial sequence used in the compact extraction and positive free-Gram argument, including the bounded lifts, common readout, unit and conjugation conditions, and the uniform volume quantifier.

11 The common continuum hierarchy

Two bound families enter compact extraction. One is available at each fixed physical period, with constants allowed to depend on that period. The other holds uniformly over sufficiently large periods, eventually in the mesh at each such period. The second is supplied separately; it is not a consequence of the first. The full inner Cauchy limit and the compact selection are identified before the outer period refinement is taken.

11.1 From convergent transported coefficients to fixed-level integrals

Fix the intermediate level m, period p, degree d, and an insertion coefficient y. Along the chosen increasing finest cutoffs, the transport construction gives a Cauchy sequence of primary coefficient vectors and a Cauchy sequence of captured interaction coefficients on one common dictionary at level m. The latter statement includes equality between each captured coefficient's realization and the interaction part of that same transported insertion. Thus coefficient convergence is applied to the actual transported insertion, rather than to an independently chosen representative.

The coefficient-to-observable estimates convert these two Cauchy statements into uniform boundedness and uniform Cauchy convergence over all configurations, both for the primary insertion and for every forced local profile. Spatial pullback, conjugation when required by the motion, and multiplication by its fixed scalar preserve these properties. Boundedness of a Cauchy coefficient sequence also bounds its paid coefficient norm, since that norm is a fixed scalar times a continuous finite-dimensional norm. Measurability follows from the local continuous terms, the measurable masking functions and the convergent countable sums in the insertion formulas. Continuity of the masks is not used.

To see how this enters the normalized integral, remove the scalar vacuum factor from the restored density and denote the result by $T_n(U)$. Its mode coefficient factors as the cold density times the selected coefficient of a product: the polynomial exponential of the primary insertion, and the field

polynomial built from the background activities plus the forced profiles with the prescribed flags. The fixed-level flux bounds give uniform Cauchy convergence for the cold densities and background field polynomials. The preceding insertion bounds do the same for primary and forced profiles. Finite sums, products and continuous maps on common compact bounded ranges preserve uniform Cauchy convergence, so every mode of T_n has that property.

Integration against the probability Haar prior therefore gives Cauchy numerator integrals. The stripped partition functions are Cauchy as well and have a common strictly positive lower bound, as proved earlier. Completeness gives their limits; the lower bound makes the denominator limit nonzero. Taking the quotient proves that the normalized mode integrals are Cauchy. Finally, the scalar vacuum factor cancels exactly in numerator and denominator: this quotient equals the common-normalization integral of the restored insertion. This is the link between the fixed-level coefficient construction and the earlier fine-to-intermediate comparison. The transport recurrence and captured-coefficient estimates are proved in the fixed-level convergence sections below.

11.2 Identification of the public hierarchy with the reconstructed one

The final public projection introduces no second quantum field system. Its Hilbert space is the reconstruction's completed state space, its common domain is the reconstructed word domain, its vacuum is the reconstructed vacuum, and its field operator is the same operator on that domain. The label space is the union of the graded field spaces already obtained from the continuum table. The phase space and probability measure are those of the character decomposition. For a character χ, the public separated correlation is the table's evaluation in that character; the unselected correlation is the table's unselected evaluation. The OPE reader supplies the channels, channel fields and coefficients. These identifications keep the lattice limits, OPE terms and reconstructed fields attached to the same table.

The equality of vacuum moments can be checked without any limiting interchange: recursively evaluating a list of smeared fields in the public definition gives the identical reconstructed word vector, starting with the same vacuum. Taking its inner product with the vacuum proves equality of the moments. Similarly translations are the original strongly continuous translation operators and the energy-momentum measure is the original projection-valued measure. A Poincaré transformation is represented by its Lorentz matrix and translation vector; its test action is the original inverse affine pullback. Thus the public covariance and spectral assertions follow by literal substitution once their reconstructed counterparts have been proved.

Each graded field injects into the label space by the table's pin map. Every label has a graded representative by the field-covering property; every graded field has an insertion-token representative by the table's quotient representative map. These are separate surjectivity assertions. Nontriviality follows from the degree-zero unit: its image is nonzero by injectivity and the unit's nonvanishing; faithfulness then supplies a Schwartz test with nonzero field operator. The boundary label-quotient argument in the reconstruction chapter proves faithfulness.

For separated configurations, a whole word can be placed in one of the decomposition's resolved configurations. The character evaluation of that resolved word is continuous in the character by the decomposition's weak continuity property. Its integral equals the unselected word expectation by the barycenter identity. Reading that word through the table gives both character continuity and the public mixture formula. The compression-valued moments and character readings in the reconstruction chapter provide the resolved-word placement and barycenter identities used here.

The analytic continuation used publicly is the reconstruction's holomorphic wave evaluated on the pinned graded labels. At strictly ordered Euclidean times, assembling the spatial coordinates and times reconstructs exactly the original four-dimensional configuration, so the heat-root iden-

tity gives its selected Euclidean value. For the boundary limit, materialize the product of Schwartz tests as one joint test. The public complex shift $t_i + irh_i$ is exactly the reconstruction's shift by the ordered height vector rh. Its boundary theorem therefore gives the required vacuum moment. The substantive holomorphy, heat-root and boundary-existence arguments remain part of the reconstruction proof, rather than consequences of these notational identifications.

11.3 Compatibility of the chosen family and its refinements

The construction chooses its Wilson family before a volume subsequence or an observable is requested. The explicit order of its pre-length floors, post-length numerical bills and countable width demands is (CAL12)–(CAL18). Their consumption gives the flat delay condition at entrance index zero on the run's actual horizon, as proved there. The same fixed family is used below.

At the continuum extraction stage there are two successive refinements of a requested order f. The first chooses a subsequence g along which the running reference parameters have limits at each fixed level. The second extraction gives a subsequence w, one continuum table, its final structure F and scaling fields. Use the composite refinement $g \circ w$. Every convergence assertion from the first extraction survives the second, because a strictly increasing natural sequence tends to infinity. Composition of the subsequence maps gives $(f \circ g) \circ w = f \circ (g \circ w)$, so all conclusions have precisely the same final volume order.

On that final table choose the reconstruction for every character, one equivariant OPE reader, and the trace test. The reader's coefficients and local fields are retained as one prescription at every truncation degree; the contact devices are then chosen with that same reader basis. This is the quantifier order used by the public theorem and its detailed extension. The preceding construction estimates and the subsequent reconstruction, operator-product and trace arguments establish these services on the retained table. The compatibility argument here proves that using these supplied results together preserves their actual family, table and subsequence, rather than treating unrelated existence witnesses as interchangeable.

11.4 Extending the full fixed-volume mesh limit

At fixed intermediate level the configuration space is fixed, and the source's Confluence condition means eventual uniform boundedness and uniform Cauchy convergence in that configuration variable. Finite sums and products preserve it. A continuous map preserves it when the range lies in a common compact set: uniform continuity on that compact set controls the Cauchy modulus. In particular continuous polynomial coordinate maps preserve it for the finite insertion algebra. For eventually measurable functions on the fixed Haar probability space, uniform boundedness gives integrability and

$$|\int f_N - \int f_M| \le \int |f_N - f_M| \le \sup_U |f_N(U) - f_M(U)|.$$

Thus the integrals are Cauchy. The finite polynomial insertion numerator is built from the scalar-normalized density, transported primary insertions and local error insertions by exactly these operations. Its Confluence follows from the corresponding fixed-level field and density statements. The scalar-normalized partition integrals are likewise Cauchy.

Their common positive lower bound is kept before division. At fixed level j and volume p, let E be the number of sites, v the positive small-field threshold, $H = 2\widehat{\beta}_j$, and let K be the uniform upper bound for the action and activity costs on the selected small-field neighborhood. Write D_{mat} for the fixed finite real matrix-coordinate dimension entering this estimate, denoted $\dim_{\mathbb{R}} \mathsf{Mat}(G)$

in the supplied notation. Its precise value is not used in the quotient argument; the lower bound is positive for every such fixed finite dimension. The denominator estimate gives

$$\exp[-K - \#\mathsf{Edges}\, D_{\mathrm{mat}}\, \log(16/v)]$$

for the scalar-stripped partition function, independently of the finer cutoff $N \geq j$. Its equality to the fixed-level normalized denominator follows from the selected trajectory's partition identity. The numerator and denominator limits therefore form a valid convergent quotient. Write c_{vac} for the represented scalar vacuum term. The factor $e^{-c_{\mathrm{vac}}}$ cancels exactly in this quotient, recovering the represented expectation. This proves the fixed-level ratio convergence once the uniform density and insertion convergence and the small-field denominator estimate have been established by the nonlinear construction.

The Taylor bound used here follows from bounds on two further lattice differences. On a compact collision-free K, choose a compact thickening still disjoint from collisions. All first differences of a fixed difference word are bounded there. A coordinate-by-coordinate lattice path joining two nearby mesh points stays in this thickening and has at most a fixed dimension factor times their physical distance divided by a_N steps. Telescoping the first differences along the path gives a local Lipschitz bound for the word. Apply this to each first difference, using its own first differences, to obtain a common Lipschitz constant A for the discrete gradient. Subtract the affine function with that gradient frozen at the starting point. Each unit difference of the residual is a_N times the change in the gradient, hence bounded by a_N A times the total endpoint distance. A second path telescope gives a dimension factor times A times the squared endpoint distance. This is the uniform quadratic Taylor estimate above. FixedBounds supplies every fixed finite difference order, so the two extra derivatives are available without changing the physical volume.

One can obtain the next-difference limit without selecting another mesh subsequence. Fix x, a coordinate a and $\varepsilon > 0$. Choose a positive physical displacement q small enough to stay inside the compact neighborhood and to make Bq small compared with ε, where B is its Taylor constant. Put $k_N = a_N\lceil q/a_N\rceil$, so $k_N \to q$ and eventually $q \leq k_N \leq 2q$. The long difference quotient of the known convergent word between its address near x and that address shifted by $\lceil q/a_N\rceil e_a$ is Cauchy, because both endpoints have cluster limits and $k_N \to q > 0$. The quadratic Taylor estimate places its distance from the unit next-difference value at x below $Bk_N \leq 2Bq$. Choose q first to make this error $\leq \varepsilon/4$; then take N,M large enough that the two long difference quotients differ by $< \varepsilon/2$. The next unit-difference sequence is Cauchy at x. Its eventual oscillation bound follows from the preceding local Lipschitz estimate and the $O(a_N)$ error of mesh addressing. The dense-set extension lemma, now applied to this derivative sequence at every x, gives its locally uniform limit along the original mesh sequence. Repeating proves the existence of all next limits used in the derivative-identification argument.

The discrete-to-classical derivative identification uses a quadratic Taylor bound. Suppose f_l is the locally uniform limit of the lattice difference word l, and $g_{a,l}$ is a locally uniform limit of its next coordinate difference. The fixed bounds supply, on a compact neighborhood, a uniform estimate

$$|F_{N,l}(Y) - F_{N,l}(X) - \sum_a (Y_a - X_a) F_{N,a::l}(X)| \leq B|Y - X|^2$$

for physical lattice points sufficiently close together. Approximate real points x,y by their mesh addresses. Pass to the limit in the inequality to obtain

$$|f_l(y) - f_l(x) - \sum_a (y_a - x_a) g_{a,l}(x)| \leq B|y - x|^2.$$

After division by |y-x|, the right side tends to zero. Therefore f_l is Fréchet differentiable, with derivative $h \mapsto \sum_a h_a g_{a,l}(x)$. The continuous limits $g_{a,l}$ identify its coordinate derivatives. Starting with the undifferenced limit and repeating for every finite word proves smoothness to all orders, since each derivative matrix is a finite linear combination of the continuous next limits. The constants and mesh thresholds may depend on the order and compact set. The coefficient-history convergence and uniform derivative estimates provide these next limits and Taylor bounds at every fixed finite order.

The result is independent of the particular fixed forward, backward or centered stencil. For nonnegative integers b,f with $e = b + f > 0$, the normalized difference

$$\frac{F(x + fe_a) - F(x - be_a)}{(b + f)a_N}$$

is the arithmetic average of the e unit forward differences based at $x - be_a + je_a$, $j = 0, \ldots, e-1$. Unit differences commute with these finite translation averages. Iterating expresses any finite mixed stencil word as a finite average of the corresponding unit-difference word. Every averaged argument has physical displacement $O(a_N)$ and hence the same limiting physical point. The cluster limit is therefore the same derivative, with no extra stencil-length factor. This proves the claimed normalization for all fixed reach stencils.

For the actual interpolant, take the surrounding right-closed mesh cube. If θ_a is the coordinate fraction measured from its right endpoint, assign a vertex choosing the left endpoint in coordinates I the weight $\prod_{a \in I} \theta_a \prod_{a \notin I}(1 - \theta_a)$. These weights are nonnegative and sum to one. Every vertex is within $O(na_N)$ of the physical point. Cluster convergence at those vertices, continuity of the limit and compactness imply uniformly small vertex errors on each compact collision-free set. The error of their weighted sum is bounded by the maximum vertex error. Thus the multilinear interpolants converge locally uniformly, also at cell ties. This assertion does not require or claim that a finite-mesh interpolant is smooth at its cell boundaries; lattice differences are taken before interpolation.

The dense-position Cauchy argument uses a second uniform approximation. Let p_N denote the fine-cutoff moment for a fixed separated configuration. At a fixed intermediate level j, replace it by its transported moment $V_{j,N}$. The fixed-level confluence estimates make $V_{j,N}$ Cauchy in N. The comparison theorem bounds $|p_N - V_{j,N}|$ by a fixed finite cost divided by a scale function $g(j) \to \infty$, uniformly for all sufficiently large N. To use that theorem simultaneously for a finite word, choose positive weights $x_i = \zeta_i/(|w_i| + 1)$, where w_i bounds the charged coefficient cost and ζ_i is its allowed margin. Then x_i times the charged cost is at most ζ_i. The finite collection of transported observables and localized errors has one eventual weighted bound by summing its individual bounds. This supplies the common comparison cost. Choose j making the uniform error below $\varepsilon/4$, then N,M large enough that $|V_{j,N} - V_{j,M}| < \varepsilon/4$. The triangle inequality proves the Cauchy property of p_N. The convergence-of-histories chapter establishes the confluence and comparison estimates used in this argument.

The equicontinuity lemma used in the fixed-volume construction does not require that the approximating functions themselves be continuous. Let U be open in a locally compact metric space and let $T_N : U \to \mathbb{C}$ have eventual uniform oscillation bounds on each compact $K \subset U$: for every $\varepsilon > 0$ some $\delta > 0$ makes $|T_N(x) - T_N(y)| < \varepsilon$ for all sufficiently large N and $x, y \in K$ with distance less than δ. Assume $T_N(x)$ is Cauchy on a dense subset $D \cap U$. For an arbitrary x choose a compact neighborhood $K \subset U$ and a dense point d within the oscillation radius. Then

$$|T_N(x) - T_M(x)| \le |T_N(x) - T_N(d)|$$
$$+|T_N(d) - T_M(d)| + |T_M(d) - T_M(x)|.$$

Taking all three terms below $\varepsilon/4$ proves the Cauchy property at x. Completeness of $\mathbb{C}$ gives a pointwise limit f. Passing to the limit in the oscillation estimate proves continuity of f. If convergence failed uniformly on a compact K, choose $N_j \to \infty$ and $x_j \in K$ whose errors stay above a fixed ε. A convergent subsequence $x_j \to x$, the oscillation bound for T_{N_j}, pointwise convergence at x and continuity of f contradict that error. Thus $T_N \to f$ locally uniformly on U, with no mesh refinement.

In the source U is the periodic collision-free position set at a fixed physical volume. The nonlinear bounds supply its oscillation estimate. The dense set is the schedule-compatible set of ancestor positions. At those positions the common-history comparison gives the Cauchy property for the full mesh sequence; the two-level approximation and fixed-level confluence arguments give this Cauchy property. The preceding lemma then extends this full-sequence convergence to all of U. The Taylor and next-difference arguments above identify all derivative limits and prove smoothness.

To pass a known cluster limit of lattice values to tests, express the lattice sum as the integral of the piecewise constant function on the physical lattice cells. Each cell has volume a_N^{4n}, exactly the mesh prefactor in the sum. The compact test support stays inside a suitable physical period box and away from collisions. Cluster convergence gives an eventual uniform bound on all lattice values whose physical positions lie in this compact support: otherwise choose an unbounded sequence of such values, extract a convergent sequence of physical positions in the compact set, and contradict the cluster limit. The test is uniformly bounded, so these piecewise constant integrands have one integrable bound on the fixed volume box. The sampled physical positions converge to each point, and continuity of the test and the cluster limit gives pointwise convergence. Dominated convergence therefore yields the integral of the test against the limiting correlation. The same argument applies to every already identified derivative cluster limit. This proves the smearing passage once those cluster limits have been established.

The scalar coupling refinement has a direct compactness proof. For each level m the represented action supplies the interval bound $\widehat{\beta}_m/2 \le \beta_{N,p,m} \le 2\widehat{\beta}_m$ whenever $m \le N$. Extend $\beta_{N,p,m}$ to $N < m$ by the value $\widehat{\beta}_m$, solely for this compactness argument. Regard the extended array at N as a point of the countable product $\prod_{p,m}[\widehat{\beta}_m/2, 2\widehat{\beta}_m]$. This product is compact metrizable. Extract one mesh subsequence along which every coordinate (p,m) converges to $a_{p,m}$. Then regard $(a_{p(k),m})_m$ along the requested volume sequence as points of $\prod_m[\widehat{\beta}_m/2, 2\widehat{\beta}_m]$, and extract one volume subsequence along which they converge to B_m for every m. For fixed m, the auxiliary extension agrees eventually with the original coupling because the mesh index tends to infinity. Thus the chosen refinement has the stated iterated limit for every level. Any later mesh and volume refinements preserve it.

These limiting coefficients retain the running slope. Write

$$t_m = \log(a_m^{-1}), \qquad \widehat{\beta}_m = \bar{\beta} + s t_m, \qquad s = 2b_0, \qquad b_0 = \frac{11 C_A}{48\pi^2},$$

where C_A is the adjoint Casimir fixed in the conventions. For every $\varepsilon > 0$ and all sufficiently large m, actual-history running gives

$$|\beta_{N,p,m} - \widehat{\beta}_m| \le \varepsilon t_m \qquad (N \ge m),$$

uniformly in p. Taking the ordered limit at fixed m preserves this closed inequality. Hence $(B_m - \widehat{\beta}_m)/t_m \to 0$ and $B_m/t_m \to 2b_0$.

Let $m(\mu)$ be the first level with $t_m \ge \log \mu$. The schedule mesh condition gives $t_{m+1}/t_m \to 1$. Since

$$t_{m(\mu)-1} < \log \mu \le t_{m(\mu)},$$

squeezing gives $t_{m(\mu)}/\log\mu \to 1$ as $\mu \to \infty$. The positive auxiliary history coupling $g(\mu) = B_{m(\mu)}^{-1/2}$ therefore satisfies

$$g(\mu)\sqrt{2b_0 \log\mu} \longrightarrow 1, \qquad g(\mu) \longrightarrow 0.$$

This concerns the coupling read from the represented history. Its identification with a separately defined position-space diagnostic requires the corresponding comparison theorem and is not made by this scalar compactness argument.

11.5 Common compact extraction and the field quotient

The physical field space is a quotient by the equation image, followed by identification through the complete spectator detector. Uniformly bounded projections permit its finite-dimensional coordinates to be retained on the common extraction. Kernel dimension stability uses closeness of idempotents: if $u \in \ker P$, $u = Qu$ and $\|P - Q\| < 1$, then $\|u\| = \|(Q - P)u\| \le \|Q - P\|\|u\|$, hence $u = 0$. Applying the same argument in the opposite direction gives equality of the two kernel dimensions. Inclusions and derivatives descend only after their compatibility with this equation image has been proved.

We describe explicitly the compactness step that puts the preceding estimates on one continuum table. It has two different inputs. At every fixed physical period p, each fixed finite list of insertions has eventual centered root bounds, with constants allowed to depend on p. Separately, on sufficiently large periods there are centered bounds uniform in those periods, eventually in the mesh at each period. The nonlinear insertion estimates give both inputs. The second is not inferred from the first.

Here is the bound for an ordinary correlation derived from these centered bounds. Fix n insertions and a compact separated set of physical positions, with separation margin $r > 0$. Choose an intermediate spacing $a_m < r$, and then a finer reporting level $k \ge m$ so that $6na_k < a_m$. The stencil of each insertion fits its reporting chamber at every sufficiently fine mesh, uniformly in the period: each fixed integer stencil offset times the fine spacing tends to zero, whereas $a_k > 0$ is fixed. Include the finitely many signed images and reflected insertions in the common centered bound A.

Choose the origin of the m-tiling so that all n insertion points have distance at least three reporting widths from its chamber boundaries. The inequality $6na_k < a_m$ leaves an allowed origin in each coordinate. Separation and $a_m < r$ place the insertion points in distinct chambers. Pull each chamber back to the reference chamber, with its prescribed signed motion and optional conjugation. The insertion locality and stencil fit give genuine chamber-supported functions, and each has rooted bound at most A. The reflection-positive chessboard inequality therefore gives

$$|\langle I_1(x_1)\cdots I_n(x_n)\rangle| \le A^n. \tag{CV1}$$

This argument applies to arbitrary separated tuples; it does not require every coordinate projection of the tuple to be separated.

For equicontinuity first work on the open set of tuples whose coordinate projections are pairwise distinct, using periodic coordinate differences at a fixed volume. Choose a compact thickening within that open set and a margin $4r$ for all its projected gaps. The centered translation-difference estimate, with a reporting spacing below r, bounds a displacement h of one slot in one coordinate by

$$|\langle I(x + he)\cdots\rangle - \langle I(x)\cdots\rangle| \le (A^n/r)|h|, \qquad |h| \le r.$$

For two nearby tuples change their $4n$ coordinates one at a time. Each intermediate tuple is obtained by choosing coordinates from one endpoint or the other, so its distance from either endpoint is at most their distance. It stays in the chosen thickening. Summing the $4n$ one-coordinate bounds

gives a local Lipschitz constant $4nA^n/r$. The same construction is available at each fixed period and, separately, uniformly in all sufficiently large periods.

Only countably many functions need be selected. At each degree choose a finite basis of the label space and take finite words in these basis vectors, all arities, and a countable compact exhaustion of their generic position sets. Multilinearity recovers arbitrary labels from these basis words after selection. Boundedness (CV1) and the equicontinuity estimate give, by the countable local Arzelà–Ascoli argument, a common mesh subsequence with a locally uniform limit M_p for every fixed p. Apply the large-period bounds to these limits: passing a boundedness or oscillation inequality through the mesh limit preserves it. A second countable Arzelà–Ascoli selection along the requested period sequence gives a table S, locally uniformly on generic configurations. All degrees and words use these same two selections.

For clarity, the countable compact selection itself introduces no hidden uniformity assumption. If the i-th coordinate lies eventually in a compact set K_i, choose one of its values in K_i and replace its exceptional earlier values by that value. The resulting sequence belongs to the countable compact product $\prod_i K_i$, which is compact metrizable in the cases used here. A convergent subsequence in that product converges in every coordinate. The replacements are eventually absent in each fixed coordinate, so the original sequence has the same coordinate limits.

This also retains additional countable numerical data $Y_i(N,p)$. For every active pair (i,p), first use its own compact fixed-period bound and select their mesh limits $z_i(p)$ simultaneously. For each i, a separate compact set K_i bounds $Y_i(N,p)$ at all sufficiently late periods and sufficiently fine meshes at each such period. Closedness of K_i puts $z_i(p)$ in K_i eventually in p. A common period subsequence then gives $z_i(p) \to y_i$ for every i. Fixed-period limits are kept at every active p, including periods absent from the selected period sequence. Apply this construction to the finite-dimensional coordinate projections and the coupling data used above.

Here is the construction and uniform bound for those finite-cutoff projections. At degree d, let $K_{N,p}$ be the image of the free equation ideal under the exact calibrated label map already constructed. This is a subspace of the single fixed coefficient space E_d, although it varies with N,p. Choose once a linear coordinate isomorphism $W : \mathbb{C}^{\dim E_d} \to E_d$, with Euclidean norm on its source, and let π_K be orthogonal projection onto $W^{-1}K$. Put

$$F_K = I - W\pi_K W^{-1}. \tag{CVP1}$$

It kills K, and $x - F_K x \in K$ for every x. Moreover

$$\|F_K\| \le \|I\| + \|W\|\|W^{-1}\|, \tag{CVP2}$$

independently of K, since $\|\pi_K\| \le 1$. Thus no conditioning estimate on a moving basis of $K_{N,p}$ is required.

When K is invariant under the signed label action V_s and coefficient conjugation, average F_K to

$$Q_K = \frac{1}{|\mathcal{S}|}\sum_{s\in\mathcal{S}} V_s F_K V_{s^{-1}}, \qquad P_K x = \tfrac{1}{2}\{Q_K x + \overline{Q_K \bar{x}}\}. \tag{CVP3}$$

Both maps still kill K and differ from the identity by a K-valued map. For the second statement, average the differences $x - V_s F_K V_{s^{-1}} x$, each of which belongs to K, and then use conjugation invariance. These two properties imply $P_K^2 = P_K$ and $\ker P_K = K$: apply P_K to $x - P_K x \in K$, and conversely use that same inclusion when $P_K x = 0$. Reindexing the finite group sum proves signed equivariance; the real average proves conjugation equivariance. Its norm is bounded by (CVP2) times the fixed finite average of $\|V_s\|\|V_{s^{-1}}\|$. The real average does not enlarge that bound, because conjugation is isometric in the frame coordinates. At the finitely many cutoffs where the calibrated

symmetry has not yet been established, use (CVP1) instead. This supplies a uniform bound at every cutoff and the required exact symmetries on the admitted tail.

The projection maps are therefore idempotent and uniformly bounded before extraction. Their limits are idempotent as well, at each fixed period and in the outer limit. Their kernel dimensions stabilize in both limits. Indeed, if P, Q are idempotents and $\|P - Q\| < 1$, the map $u \mapsto u - Qu$ sends $\ker P$ into $\ker Q$ injectively: if $u = Qu$ and $Pu = 0$, then $\|u\| = \|(Q - P)u\| \leq \|Q - P\|\|u\|$, forcing $u = 0$. Interchanging P, Q gives equal kernel dimensions. The physical detector estimates identify the limiting kernel with the exact equation-null space, rather than merely bounding its dimension. The transported insertion comparison shows that a label in this kernel has vanishing separated smeared correlations: its discrepancy has a uniform separated bound tending to zero, and the step-function integral argument gives the same vanishing for every compact separated test.

Conversely, the finite full-rank detector pairs separate the quotient. A label with zero continuum correlations against every spectator has zero detector readings, hence lies in the limiting kernel. Thus

$$\{\text{continuum-null labels of degree } \leq d\} = \ker P_d. \tag{CV2}$$

The quotient is finite dimensional, carries the real and signed-motion structures preserved by P_d, and is the same quotient used for the detector normalization and subsequent differentiation.

The derivative construction removes the coordinate-projection restriction on the raw table. Transport the finite differences to this quotient using their retained coefficient estimates. Their limits agree with the raw table on generic tuples. The local Taylor and next-difference argument proved earlier identifies their derivatives and extends them smoothly across coordinate ties, as long as no two full four-dimensional positions coincide. The generic set is dense there, so two continuous such extensions must agree. This gives the smooth table T with the original M_p as its fixed-period tables and the same separated smeared limits. No pointwise assertion at a coordinate tie is obtained from raw compactness alone.

In the present construction the full fixed-volume mesh limits were already established by the two-level Cauchy argument. The auxiliary mesh subsequence just used therefore converges to those same M_p. To see explicitly that this restores the required smeared order, fix a token word and one compactly supported separated test t. For each suitable period p, let $R_{N,p}(t)$ be its lattice smeared moment. The full-mesh result gives $R_{N,p}(t) \to r_p(t)$ as $N \to \infty$. The extracted table gives, along its mesh subsequence $\phi(N)$, a limit $D_p(t)$, followed by the stated volume limit of $D_p(t)$. Cofinality of ϕ also gives $R_{\phi(N),p}(t) \to r_p(t)$. Uniqueness of a complex limit therefore identifies $D_p(t) = r_p(t)$ at each suitable p. Substitute this equality into the retained volume-limit identity. The result has the full mesh order followed by precisely the same volume subsequence as the final table. This argument restores the smeared convergence assertion; it does not claim that every auxiliary structure originally extracted on a mesh subsequence has been extended to all meshes.

Finally, countably many further stages can be imposed without combining incompatible choices. If each successive order refines its predecessor, take its diagonal mesh and period sequences separately. For a nested sequence F_j of strictly increasing sequences, the diagonal $t_j = F_j(j)$ is strictly increasing: writing $F_{j+1} = F_j \circ g_j$, one has $g_j(j + 1) \geq j + 1$. For every fixed stage i, the tail of t lies in the range of F_i. Apply this to mesh and period orders. Every eventual bound and every retained limit from any fixed stage survives on the resulting diagonal order. This is why the hierarchy, quotient, numerical limits and later comparison services can all be attached to the same extracted table.

11.6 The calibrated equation image and its silent representatives

We specify the finite-cutoff change of coefficients used in the preceding projection construction. This also explains why its kernel directions are silent in continuum tests. Fix a degree, its coefficient frame, and a second sound local section V, such as the marked equation section. Write e_j for the coordinate vectors and $c \geq 1$ for the fixed cost of signed motions and coordinate decomposition in this frame. The local insertion transport constructs columns u_j satisfying

$$|e_j - u_j|_{\mathrm{coeff}} \leq (2c^2)^{-1}. \tag{CVR1}$$

The columns are combined by the signed and conjugate averaging already used for local sections, rather than by an unsymmetrized matrix. Explicitly, if $\chi_{s,j}(x)$ are the averaged coordinates of the inverse motion of x, then

$$Px = \sum_{s,j} \chi_{s,j}(x)\, su_j, \qquad x = \sum_{s,j} \chi_{s,j}(x)\, se_j, \qquad \sum_{s,j} |\chi_{s,j}(x)| \leq c|x|_{\mathrm{coeff}}. \tag{CVR2}$$

Here a conjugating motion conjugates the coordinate as well, so the combined map is complex linear. The fixed motion bound is $|sv|_{\mathrm{coeff}} \leq c|v|_{\mathrm{coeff}}$. Subtract the two sums and apply (CVR1):

$$|x - Px|_{\mathrm{coeff}} \leq \tfrac{1}{2}|x|_{\mathrm{coeff}}, \qquad |P^{-1}y|_{\mathrm{coeff}} \leq 2|y|_{\mathrm{coeff}}. \tag{CVR3}$$

Indeed the first inequality makes P injective; finite dimensionality makes it surjective, and applying the first inequality to $P^{-1}y$ gives the second. These bounds hold uniformly in the period on the eventual mesh tail.

For precision about the purchase of (CVR1), put $X = V.\mathrm{stencil} \cup T.\mathrm{stencil}$, where T is the original section. At stage N, write h_N for the support exponent, q_N for the number of columns, and $b_N \geq 0$ for the purchased derivative-credit bound of V. The transport estimate gives

$$L_{N-1}^{1/8}|e_j - u_j|_{\mathrm{coeff}} \leq 2(1+q_N)Y_N, \quad Y_N = 2 + e^{81h_N}(b_N + q_N + 2|X|). \tag{CVR4}$$

The delayed length requirement is precisely $2(1+q_N)Y_N \leq L_{N-1}^{1/8}/(2c^2)$, which proves (CVR1). It is bought together with $b_N \leq \widehat{\beta}_N$, not inferred from a limit after selecting the columns. The common error price can be taken to be

$$E_N = 4(1+q_N)Y_N + e^{81(h_N+2)} Z_N \zeta_{d,N}/u_N, \quad Z_N = (a_V + 2a_T) \sum_j \theta_N^{-f_j}, \tag{CVR5}$$

where a_V, a_T are the fixed coefficient costs of the two sections, $\theta_N > 0$ is the field threshold and f_j the field-number degree of coordinate j, $u_N > 0$ is the purchased insertion allowance, and $\zeta_{d,N} > 0$ its degree-dependent margin. The sum is the finite price of undoing the threshold dilation. Each factor in (CVR5), including u_N^{-1}, has the subexponential bound supplied by the common calendar. Finite sums and products preserve that bound by dividing a requested exponent among their finitely many factors.

At any fixed admitted reporting indices m, k, the rooted discrepancy between the insertion of e_j through V and that of u_j through T, including every signed or conjugate motion, is at most

$$C_{d,m,k} E_N e^{-s_N/8}, \qquad C_{d,m,k} = \exp\{(a_m/a_k)^4 \tau_k + 5\} \frac{a_k^{-d}}{\zeta_{d,k}/4} e^{s_k/8}. \tag{CVR6}$$

Here τ_k is the explicit vacuum treasury derived in the denominator estimates above. These are fixed finite constants after the reporting indices are chosen. Expand any subexponentially priced coefficient row by (CVR2). The finite number of terms and their subexponential coordinates, together

with (CVR5), absorb at most $e^{s_N/16}$ eventually. Thus the entire discrepancy has rooted bound tending to zero, with a remaining exponential factor $e^{-s_N/16}$. Its fixed stencil, gauge invariance and continuity make it eligible for the separated spectator product and smearing argument below.

Let $\mathcal{T}$ be the terminal coefficient isomorphism and B the marked embedding of free labels into terminal coefficients. The calibrated equation image is exactly

$$K_{N,p} = \mathcal{T} P \mathcal{T}^{-1} B(I_d). \tag{CVR7}$$

If $x \in K_{N,p}$, set $y = P^{-1}\mathcal{T}^{-1}x$. Then $\mathcal{T}y \in B(I_d)$. The actual insertion through V with coefficient y is therefore exactly zero against physically separated spectators on a sufficiently fine mesh, by the local marking identity. The uniform radius for that assertion was proved by the finite-basis maximum in the equation-lift construction. Formula (CVR3) preserves the subexponential price of y, and (CVR6) compares this exactly silent insertion with the original insertion of x. Consequently every subexponentially priced row in (CVR7) tends to zero against every separated test. This is the kernel cancellation used in (CV2) and (QD3); it follows from the actual calibrated map and a local identity, rather than from dimension counting alone.

11.7 Inclusions and derivatives on the same physical quotient

The differential system is constructed using all spectator tests on this quotient. Let E_d be the finite coefficient space and N_d the null subspace in (CV2). A test reading means: put the distinguished label at any chosen slot, put arbitrary other labels in all remaining slots, multiply the correlation by a smooth compactly supported test whose support lies in the separated configuration space, and integrate. Denote it by $\mathcal{L}_\psi(b)$, including the spectator data in ψ. The family of all these linear functionals has common kernel N_d. For a smooth table this follows from uniqueness of continuous kernels as distributions on the open separated set; a nonzero value has a neighborhood and a compact bump detecting it.

At finite mesh a_N, a derivative insertion is the actual forward difference

$$\delta_a I_b(o) = a_N^{-1}\{I_b(o + e_a) - I_b(o)\}. \tag{QD1}$$

In the finite torus sum, translating the distinguished summation variable gives the exact summation-by-parts identity

$$a_N^{4n} \sum_r \psi(a_N r)\, \langle \delta_a I_b(r_i)\, \text{spectators} \rangle = a_N^{4n} \sum_r \frac{\psi(a_N r - a_N e_{i,a}) - \psi(a_N r)}{a_N}\, \langle I_b(r_i)\, \text{spectators} \rangle. \tag{QD2}$$

Translation is a permutation of the finite torus, so this is true even at its seam. Once the physical period contains the compact support and its one-step thickening, both potentially nonzero terms have their actual cover coordinates. Taylor's theorem makes the test quotient converge uniformly to $-\partial_{i,a}\psi$, on one fixed separated compact. The separated correlation bound controls the remainder in the cell integral. Take the mesh limit at that fixed suitable period first, and then the retained volume limit. This proves the derivative test limit in the same mesh-then-volume order as the hierarchy; no uniform count of lattice points over all periods is needed.

The coefficient approximations used for both inclusion and (QD1) have subexponential top-coordinate cost: for every $\epsilon > 0$, their pullback to the terminal frame is eventually bounded by $\exp(\epsilon s_N)$, where s_N is the calendar scale. For bounded reference coefficients this follows by multiplying the backward column bounds: their logarithm is $\sum_{i<N} \log C_i$. The calendar gives $\log C_i \le \epsilon \log L_i/2$ after a fixed level, so the tail is at most $\epsilon(s_N - s_k)/2$; the finite head is absorbed

into $\epsilon s_N/2$. The relay estimates extend this cost statement to the actual inclusion and finite-difference approximations.

Here is why their errors vanish against every test. A local rooted discrepancy bound $y_{N,p} \to 0$, with reporting spacing chosen arbitrarily small after fixing its requested tolerance, gives a separated correlation bound $A|y_{N,p}|$. Choose that reporting spacing below the physical separation, place the spectator insertions in distinct cells, and apply the reflection-positive product estimate, leaving the discrepancy in its own cell. Its compactly smeared sum is the integral of a step function, bounded by this quantity times the test supremum and the volume of one fixed enlargement of its support. Hence it tends to zero. This also applies to an exactly silent separated remainder with zero bound.

One must control approximations before assuming their coefficients bounded. Write $P_{N,p}$ for the finite projection and $x_{N,p}$ for an approximating coefficient. Normalize its null remainder by

$$z_{N,p} = \frac{x_{N,p} - P_{N,p}x_{N,p}}{1 + \|P_{N,p}x_{N,p}\|}. \tag{QD3}$$

It lies in the finite projection kernel. The normalized projected part has norm at most one and therefore the subexponential cost just proved. Multiplication by the reciprocal denominator costs at most one, and subtraction preserves subexponential cost by spending $\epsilon/2$ on each summand. The relay's equation cancellation then gives $\mathcal{L}_\psi(z_{N,p}) \to 0$ for every test. This is precisely the normalized remainder required by the detector argument: it proves boundedness and recovery of the projected limit using the stable finite detector, without presuming boundedness of the unprojected approximant. The detector identifies its limiting responses with one vector in E_e/N_e.

Apply this recovery to an inclusion from horizon d to $e \geq d$, and to (QD1) from d to $d+1$, first on a fixed basis of E_d. Choose representatives of the recovered vectors and extend linearly. This gives linear maps $K_{d,e} : E_d \to E_e$ and $D_{d,a} : E_d \to E_{d+1}$ satisfying, simultaneously for all spectator tests,

$$\mathcal{L}_\psi(K_{d,e}b) = \mathcal{L}_\psi(b), \qquad \mathcal{L}_\psi(D_{d,a}b) = -\mathcal{L}_{\partial_{i,a}\psi}(b). \tag{QD4}$$

The same approximating coefficient family is used for every spectator; separate choices for different tests would not define a field.

If $b \in N_d$, the right sides of both equations (QD4) vanish for every test, including the differentiated test, which still has compact separated support. Thus $K_{d,e}b \in N_e$ and $D_{d,a}b \in N_{d+1}$. The quotient maps

$$\iota_{d,e}[b] = [K_{d,e}b], \qquad \partial_a[b] = [D_{d,a}b] \tag{QD5}$$

are consequently independent of the chosen representative. Their pointwise kernel identities follow from (QD4): inclusion gives equality of two continuous kernels, and integration by parts gives the ordinary derivative of the smooth kernel. Integration by parts is legitimate on the separated open set because the test support is a compact subset of that set. Extend the product of test and kernel by zero outside; it is smooth with compact support, so its total directional derivative integrates to zero.

Faithfulness against all spectators now proves every compatibility law. Inclusion is injective, since its zero class has exactly the same readings as its input. Inclusions compose and the identity inclusion is the identity, since both sides have identical readings. Differentiation commutes with inclusion, because it differentiates the same smooth kernel in either horizon. Two derivatives commute by the ordinary mixed-derivative theorem. The derivative of the unit is zero: the unit removes its insertion slot and the resulting kernel is independent of that coordinate. Permuting insertion slots transfers these identities from the distinguished first slot to every slot by the linear permutation chain rule.

Finally conjugation commutes with real directional differentiation. For a signed orthogonal motion s, its coordinate derivative transforms by

$$\partial_b(f \circ s) = \sum_a s_{ab}(\partial_a f) \circ s. \tag{QD6}$$

Apply the already established reality and signed covariance identities to the kernel, differentiate on the separated neighborhood, and use faithfulness again. It follows that the quotient inclusions commute with both symmetries and that the quotient derivatives obey (QD6). This establishes the differential structure on the actual null quotient used in reconstruction, with no extra representative-dependent operation.

12 Phase decomposition and Wightman reconstruction

The reconstruction has three levels. Finite reflection-positive correlations first give Hilbert vectors and bounded slab operators. The commutative ground-compression algebra then produces characters and a full-support phase measure. Finally tube continuation and tempered boundary limits construct fields on a common cyclic domain in each selected phase. The proof below keeps these levels separate: it does not assume a one-dimensional ground space before the phase argument establishes vacuum uniqueness.

12.1 Bounded slab insertions and their norm derivatives

Fix an axis a, width $d > 0$, and word $w = ((x_i, v_i))_{i=1}^n$ of distinct points inside $0 < x_i^a < d$. A buffer $0 < r \le 1$ requires pairwise maximum-coordinate distance at least r and $r \le x_i^a \le d - r$. For degrees at most D and coefficient norms at most R, we construct

$$\|U_d(w)\| \le C^n (n/r)^{An}, \tag{BS1}$$

with $C > 0, A \in \mathbb{N}$ depending only on D, R. The empty word has bound one. On a future plant the prescribed action inserts w before that plant translated by de_a. Boundedness precedes extension to the Hilbert completion.

The centered one-cell estimate gives a threshold F, exponent P, and $J > 0$, such that one insertion at level $k \ge F$ has price Ja_k^{-P}. Include all signed rotations and reflected conjugates by enlarging the coefficient ball by their finitely many operator norms on the finitely many raw degree spaces. These constants consequently work for every reflected copy.

The decreasing schedule satisfies $a_{k+1} \ge c_0 a_k^2$. At the first $m \ge F$ below $r/32$, we have $a_m \ge c_1 r^2$; decrease $c_1 > 0$ to include $m = F$. Choose fixed sufficiently small $s > 0$ and the first $k \ge F$ below $s(r/n)^2$. The same first-crossing argument gives

$$a_k \ge c(r/n)^4, \qquad 20na_k < a_m, \qquad k \ge m. \tag{BS2}$$

Indeed $20na_k < 20sr^2/n \le 20sr^2 < c_1 r^2$, and strict decrease gives $k \ge m$. Exterior insertions are paid at still finer scales, chosen after this main scale; their prices may depend on the exterior word.

First let d' be an integer multiple of a_m. Alternating insertion of w and its reflection produces, in the squared norm after q repetitions, $4q$ main copies and two exterior copies. The main positions repeat the original $2n$ position patterns modulo a_m. Thus cell packing needs only $20na_k < a_m$, independently of q. The buffer separates the main train from the translated exterior. The rooted product estimate and the fixed finite-word lattice limit give

$$\|\text{plant of the repeated word}\|^2 \le B_p[Ja_k^{-P}]^{4nq}, \tag{BS3}$$

with B_p independent of q. This is the train estimate in `Next/PHASE/Trains.travel_price` and Seats.replicas. By (BS2), its main price is

$$Q = [Ja_k^{-P}]^n \le [Jc^{-P}]^n (n/r)^{4Pn}.$$

We may use the latter expression for Q.

To remove the exterior prefactor, let X send algebraic future words to the reflection Hilbert space, V insert the main word, and $V^\sharp$ insert its opposite reflected word. Reflection and translation invariance give $\langle XVu, Xv\rangle = \langle Xu, XV^\sharp v\rangle$. Put $T = V^\sharp V$. Taking square roots in (BS3) and finite linear combinations gives $\|XT^q u\| \le B_u Q^{2q}$. Therefore $s_u = \sup_q Q^{-2q}\|XT^q u\|$ is finite. Moving powers of T across the pairing yields

$$(Q^{-2q}\|XT^q u\|)^2 \le \|Xu\| s_u.$$

Taking suprema gives $s_u^2 \le \|Xu\| s_u$, hence $s_u \le \|Xu\|$. Finally $\|XVu\|^2 = \langle Xu, XTu\rangle \le Q^2\|Xu\|^2$. Insertion thus annihilates null vectors, descends to the quotient, and extends uniquely with norm at most Q. This is OS/Train.producer.

For general d, use $d' = a_m\lfloor d/a_m\rfloor$. Since $a_m < r/32$, the same word lies in the narrower slab with buffer $r/2$. Compose its operator on the right with the heat contraction of width $d - d'$. On plants this is exactly the required insertion at width d, and its norm does not increase. The empty word is the heat contraction. All intermediate choices give the same operator because they agree on the dense plants. Replacing r by $r/2$ only enlarges C, proving (BS1) for one common family of operators. This is the padding and uniqueness argument of OS/Halves and Work/B01.

We next prove operator-norm regularity. Fix labels and width, and write $F(x) = U_d(w_x)$ on the open set of distinct slab configurations. For a slot-axis index $b = (i, j)$, replace its label by a representative of its physical j-derivative and call the resulting placement $F_b(x)$. Set

$$B(x)h = \sum_b h_b F_b(x). \tag{BS4}$$

Near a fixed configuration there is one positive buffer for the original word and these finitely many differentiated words. Equation (BS1) makes B locally bounded in operator norm. Between two plants, the placement is a separated correlation with a fixed exterior word. The scalar slot-derivative law gives its derivative as the matrix element of $B(x)$, also for finite linear combinations of plants.

On a convex position ball suppose $\|B(z) - B_0\| \le \eta$. The scalar mean-value estimate gives

$$|\langle u, [F(y) - F(x) - B_0(y - x)]v\rangle| \le \eta\|u\|\|v\|\|y - x\|.$$

Density gives the same estimate in operator norm. First set $B_0 = 0$: local boundedness makes every placement locally Lipschitz. Apply this also to each differentiated label array; all F_b, hence B, are norm continuous. Then set $B_0 = B(x)$ and shrink the ball. Its oscillation tends to zero, proving $DF(x) = B(x)$ in norm. Induction with differentiated labels gives smoothness of every order on the same position domain. This is the argument of `Work/B02/Analysis.norm_remainder` and `Work/B02/Slab.smooth_place`; it establishes the norm regularity used in the waves and compression-valued moments below.

12.2 Surjectivity of the reflection attachment

Surjectivity of attachment The reflection Hilbert space is formed from a positive pairing on future words. An attachment map can preserve that pairing and still have a proper closed range,

so its being an isometry is not the final step. Surjectivity below identifies a dense family of heated word vectors in that range and uses closedness to obtain the whole space. This identifies the reconstructed representation with the intended cyclic word space and is needed when the boundary fields are later placed on their common domain.

The attachment J is an isometry from the complete boundary Hilbert space, so its range is closed. First consider a regularized word vector integrated against the Fourier transform of a compact smooth test. The boundary tilt identity identifies it with J applied to the boundary word distribution tested against the correspondingly chilled test. Here is the Hilbert-space step in that identification. If A is the reflection-space regularized vector and B the proposed attached vector, the joint pairing formula gives $\langle A, A\rangle = \langle B, B\rangle$. Pair A with the regularized approximants tending to B. Their partial-boundary correlations converge by the boundary distribution construction, giving $\langle A, B\rangle = \langle B, B\rangle$. Therefore $\|A - B\|^2 = 0$ and $A = B$. The strict-time wave and independently heated word constructions below establish the boundary tilt and joint pairing identities used here.

Compact cutoffs approximate any Schwartz test after Fourier transform. The regularized word map is continuous, by its bound $C\|f\|_{L^1}$, so closedness of the range extends membership to every integrated regularized word vector. For a pointwise regularized word vector, pair with a vector orthogonal to the range. All its smeared pairings vanish. The pointwise pairing is continuous, so equality as a distribution implies that it vanishes everywhere. The pointwise vector belongs to the double orthogonal complement of the range, hence to the range itself.

A future Euclidean word with strictly ordered positive times is exactly such a regularized word: choose the regularization below every initial time and consecutive time gap, then recursively insert the local regularized fields and intervening translation and heat operators. This recursion reproduces the reflection-space plant of the word with its specified insertion representatives. Quotient-field representatives give the same plant because their table expectations agree against all future words. Hence every strictly time-ordered future plant lies in J's range.

For an arbitrary future word, permutation symmetry first sorts its times. Add εi to its i-th sorted time. For $\varepsilon > 0$ the times are strictly increasing and remain positive. The original configuration is separated; the perturbed configurations remain separated and converge to it as $\varepsilon \downarrow 0$. Smoothness of the separated table makes the associated plant vectors continuous in norm, as seen by expanding the squared distance into three table pairings. Since the range is closed, the unperturbed plant also lies in it. Future plants have dense linear span by the GNS construction. Therefore the closed range is the whole reflection Hilbert space, proving surjectivity.

12.3 Explicit tested boundary formula

Fix a field word v and let $\Psi(y, f)$ be its holomorphic wave smeared against a joint Schwartz test f at strictly increasing imaginary heights y. Let $a_i = i$ and define the continuous test operator $D = -i\partial_{\mathrm{dawn}(a)}$. The tested Cauchy–Riemann identity is

$$\frac{d}{dt}\Psi(y + ta, D^k f) = \Psi(y + ta, D^{k+1} f).$$

Assume the wave bound $|W(x+iy)| \le Ck^{-m}$ whenever the height gaps are at least k, $0 < k \le 1$. For any nondecreasing y and $t > 0$, y+ta has gaps at least t. Define

$$P_m(y, f; t) = \sum_{k=0}^{m} \frac{(-t)^k}{k!}\Psi(y + ta, D^k f).$$

Differentiating and cancelling adjacent terms yields $\partial_t P_m(y, f; t) = (-t)^m \Psi(y+ta, D^{m+1} f)/m!$. The wave bound and integration in x give

$$|\partial_t P_m(y, f; t)| \le \frac{C}{m!}\|D^{m+1} f\|_{L^1}, \qquad 0 < t \le 1.$$

Thus the expression

$$B(y, f) = \sum_{k=0}^{m} \frac{(-1)^k}{k!}\Psi(y + a, D^k f) - \int_0^1 \frac{(-t)^m}{m!}\Psi(y + ta, D^{m+1} f)\, dt.$$

is defined for every nondecreasing y. For strictly increasing y, the fundamental theorem of calculus gives $B(y, f) = \Psi(y, f)$. For y approaching any nondecreasing y_0, the finite terms are evaluated at interior heights and are continuous. The remainder integrands are continuous for every $t > 0$ and share the constant integrable bound just proved. Dominated convergence makes B continuous at y_0. Therefore $\Psi(y, f)$ tends to $B(y_0, f)$ along every strict-height approach, with no restriction on the ratios of the gaps.

Furthermore

$$|B(y, f)| \le C \sum_{k=0}^{m} \frac{\|D^k f\|_{L^1}}{k!} + \frac{C}{m!}\|D^{m+1} f\|_{L^1}.$$

This is a continuous Schwartz seminorm. Linearity in f and multilinearity in the labels pass through the interior limits and yield a tempered multilinear boundary distribution. Different choices of the admissible exponent m give the same distribution by uniqueness of these limits. The common character-independent C, m prove both boundary continuity in the character and passage of the character barycenter to the boundary. The construction of the waves and the tested Cauchy–Riemann and polynomial gap estimates used here is given in (SW1)–(SW7) and (CR1)–(CR2) below.

12.4 Character integration at strict imaginary times

For a fixed field word v of arity n, the heat-regularized field factors have bounds $Q_a\, p^{-d_a}$. Choosing the regularization $p = k/8$ for imaginary-time gaps at least k, with $0 < k \le 1$, gives a product bound

$$|W_\chi(v; x + iy)| \le [\prod_a Q_a(8)^{d_a}] k^{-\sum_a d_a},$$

uniformly in χ and the real positions x. The comparison of the phase regularized factor norm with the unselected factor norm supplies the same constants in every phase. The factors are constructed in (SW1)–(SW3); the phase norm comparison is (ST1)–(ST4), using the bounded slab construction (BS1)–(BS4).

At a strict imaginary-time configuration the phase values are continuous in χ and have the unselected value as their integral. For a joint Schwartz test f and fixed strict heights y, the preceding bound gives the integrable dominator $Ck^{-m}|f(x)|$. Dominated convergence therefore proves continuity of $\chi \mapsto \int W_\chi(v; x + iy) f(x)\mathrm{d}x$. The same bound, combined with the probability normalization of the phase measure, makes the double integral absolutely integrable. Fubini transfers the character integral inside the x-integral and yields the unselected strict smeared wave.

The source passes this identity to the Wightman boundary using a common finite-order Taylor representation along an interior height ray. Its polynomial part uses finitely many strict evaluations; its remainder has a common integrable bound on the ray parameter after enough test derivatives. Thus continuity in χ passes separately through that finite sum and the dominated remainder integral. The same representation bounds the smeared strict waves uniformly in χ as the heights approach the boundary. Dominated convergence in χ then gives the boundary barycenter identity. The tested boundary formula and strict-time wave estimates below provide this Taylor representation and its common polynomial-order bound.

12.5 Identification of the heat ground projection

Let E_s be the positive reflection-space contraction semigroup, and let Π be its strong $s \to \infty$ limit, the projection onto its fixed vectors. Its lifted joint momentum measure gives for every vector x

$$\langle x, E_s x\rangle = \int \exp(-sE(p))d\nu_x(p).$$

Since $E \geq 0$ almost everywhere, the integrands are bounded by one and converge to the indicator of $E = 0$. Dominated convergence gives the limit $\nu_x(\{E = 0\}) = \langle x, P(\{E = 0\})x\rangle$. The strong semigroup limit gives the same limit $\langle x, \Pi x\rangle$. Polarization of these quadratic pairings proves $\Pi = P(\{E = 0\})$, without an assumption that this projection has rank one.

For the isometric attachment J from the boundary Hilbert space to this reflection space, set $B_s = J^*E_sJ$. The attachment identifies the diagonal spectral measures, so $\langle x, B_s x\rangle$ is the Boltzmann integral for the boundary momentum measure. It also intertwines the zero-energy projections. Applying J^* to the reflection-space strong limit and using $J^*J = I$ gives $B_s x \to P_{\text{boundary}}(\{E = 0\})x$ for every x. The vacuum rank is established later by the locality and character argument, not used in this identification. Surjectivity of J, supplied separately by the attachment construction, then realizes the whole reflection heat evolution on the boundary Hilbert space.

12.6 Weak irreducibility from heat evolution

With the reconstructed Hilbert space and heat evolution, let Φ be the reconstructed fields on their common word domain, Ω the cyclic unit vacuum, $U(t)$ the time translations, and B a bounded operator satisfying the weak commutation identity

$$\langle \Phi(a, f)^* u, Bv\rangle = \langle u, B\Phi(a, f)v\rangle$$

for word-domain u,v, with the adjoint interpreted by its actual Hilbert-space adjoint restricted to that domain. No preservation of the domain by B is assumed. Repeated application moves a whole word across the pairing: $\langle w^*u, Bv\rangle = \langle u, Bwv\rangle$.

Fix a word vector $u = w\Omega$ and let $v = w^*\Omega$. The heat-space attachment supplies bounded continuous scalar functions on the closed upper half-plane, holomorphic in its interior, with real boundary values

$$F(t) = \langle u, U(t)B\Omega\rangle, G(t) = \langle B^*\Omega, U(t)v\rangle.$$

Translation covariance, invariance of Ω and the weak word identity give $F(t) = G(-t)$ for real t. Glue F above the axis to $G(-z)$ below it. The boundary equality makes the glued function continuous. Cauchy's theorem on rectangles cut along the real axis, with cancellation of the common boundary integrals, shows that the glued function is entire. It is bounded, so Liouville's theorem makes it constant. In particular

$$\langle u, B\Omega\rangle = \langle u, \exp(-sH)B\Omega\rangle, s \geq 0.$$

The constructed heat evolution converges strongly to the vacuum projection as s tends to infinity. Hence $\langle u, B\Omega\rangle = \langle u, \Omega\rangle\langle\Omega, B\Omega\rangle$ for every word vector u. Density gives $B\Omega = a\Omega$, $a = \langle\Omega, B\Omega\rangle$. Move each word across B once more to obtain $\langle u, Bw\Omega\rangle = a < u, w\Omega >$ for all word vectors u,wΩ. Density and boundedness imply $B = aI$. Conversely every scalar multiple of I satisfies the weak identity by the field adjoint relation. Thus the weak commutant is precisely the scalars. The heat-contraction, holomorphic-edge and vacuum-projection arguments of this chapter establish the properties used in this weak-commutant argument.

12.7 From low-energy measure bounds to uniform heat decay

Low-energy spectral bounds control heat decay without supplying a positive mass gap. The heat-ground projection and its compression algebra are the actual objects used in phase decomposition. Full support of the character measure subsequently upgrades certain continuous almost-everywhere identities to every character, after the required continuity and common bounds have been established.

The shell estimate uses positivity of a time average. Let μ, ν be the scalar spectral measures of a compact word combination and its adjoint word combination, both supported in the forward cone, with total mass sum at most M. Their characteristic-function difference $A(x,t) = \widehat{\mu}(x,t) - \widehat{\nu}(-x,-t)$ is the vacuum commutator pairing. Locality gives $A(x,t) = 0$ whenever $|x| > R$ and $|t| \leq 1$, with R uniform in the character: the words have fixed compact test supports and the label action is common across phases. Define

$$D(x) = \int_0^1 [A(x,t) - A(x,-t)]/(2i)dt,$$
$$\rho(E) = \int_0^1 \sin(Et)dt.$$

For $E > 0$, $\rho(E) = (1 - \cos E)/E \geq 0$, with $\rho(0) = 0$ and $|\rho| \leq 1$. For $0 \leq E \leq 1$, the inequality $\sin(Et) \geq 2Et/\pi$ gives $\rho(E) \geq E/\pi$. Fubini applies since both measures are finite and the integrands are bounded. It expresses D as the spatial Fourier transform of the sum of two positive measures: $\rho(E)\, d\mu$ pushed to spatial momentum, and $\rho(E)\, d\nu$ pushed to negative spatial momentum. Moreover, $|D(x)| \leq M$, and D vanishes outside the radius-R ball.

Fix $0 < u \leq 1$ and a smooth nonnegative spatial bump B_u equal to one on $|p| \leq u$ and supported in $|p| \leq 2u$, with $B_u \leq 1$. Its Fourier transform is bounded in absolute value by the volume of the radius-2u ball, hence by $C_3 u^3$. Since D is supported in a fixed ball, Fourier inversion and positivity imply

$$\int \rho(E) B_u(p_{\text{spatial}}) d\mu \leq K(R,M)u^3,$$

where the fixed positive Fourier normalization is absorbed into K. On the shell $u/2 < E \leq u$, the cone support gives $|p_{\text{spatial}}| \leq E \leq u$, and $\rho(E) \geq u/(2\pi)$. Thus

$$\mu(\{u/2 < E \leq u\}) \leq 2\pi K(R,M)u^2.$$

Decompose $0 < E \leq u$ into the countable dyadic shells with upper endpoints $u2^{-n}$. The corresponding masses sum to at most K' $u^2 \sum_n 4^{-n} \leq 2K'u^2$. This proves the low-energy bound used below. The uniform total masses follow from continuity of the word norms on the compact

character space; those norms are table expectations of adjoint words concatenated with words. The word-moment continuity, locality and spectral reconstruction in this chapter give these properties.

Let ν be a finite positive spectral measure supported in the forward cone, and let $E(p)$ be its time coordinate. Suppose $\nu(\{0 < E \le u\}) \le Cu^2$ for $u > 0$. For $s > 0$ the identity $e^{-sE} = \int_E^\infty se^{-st}\,dt$ on $E > 0$ and Tonelli's theorem give

$$\begin{aligned}&\int_{E>0} \exp(-sE)d\nu\\ &= s\int_0^\infty \exp(-st)\nu(\{0 < E \le t\})dt\\ &\le Cs\int_0^\infty t^2\exp(-st)dt = 2Cs^{-2}.\end{aligned}$$

The final equality follows by substituting $u = st$ and integrating $u^2\exp(-u)$ twice by parts. A bound initially supplied only for $0 < u \le 1$ extends to all u by increasing C to at least ν(total), since $u^2 \ge 1$ for $u \ge 1$. Applied to the scalar spectral measure of a word vector, this is the estimate for $\langle x, (\exp(-sH) - P_0)x\rangle$. It leaves the zero-energy atom separate throughout.

Continuity in the character is obtained from translation pairings. The characteristic function of the scalar spectral measure ν_χ of a fixed sum of words is the inner product of that vector with its translated vector. Both vectors are represented by fixed words with the same transported labels, so the continuum table's word-pair continuity gives continuity in χ at each translation. The total masses $\nu_\chi(total)$ have a common bound M. For a compact smooth momentum test g, Fourier inversion writes $\int$g $d\nu_\chi$ as a fixed Fourier normalization times the integral of the characteristic function against the Schwartz function $hat(g)$. The bound M$|hat(g)|$ is integrable. Dominated convergence therefore proves continuity of $\int$g $d\nu_\chi$.

For the heat multiplier $\exp(-sE)$, use a smooth bump equal to one on the Euclidean ball of radius 2N and supported in the ball of radius 3N. On the forward cone, outside the smaller ball $E \ge N$. Thus the error in every phase is at most M $\exp(-sN)$, tending to zero uniformly in χ. The heat pairings are consequently uniform limits of the continuous compact-test pairings, and are continuous. These arguments supply the heat continuity and shell-to-heat steps used in the removal of exceptional characters. The preceding time-average argument supplies the shell bound; the compression-valued word construction below supplies continuity of the word pairings.

12.8 Removal of exceptional characters in vacuum uniqueness

From almost every phase to every phase The measure-theoretic argument first gives the relevant property outside a null set of characters. To obtain it at every character, one also needs continuity of the fixed probe readings and full support of the phase measure. A nonzero continuous violation would persist on a nonempty open set, which full support gives positive measure. The common countable probe family makes the exceptional set compatible before this continuity argument is applied. Almost-everywhere statements alone would not establish the asserted all-phase conclusion.

The phase measure has full support: every nonempty open subset of the compact character space has positive measure. Consequently a closed set containing almost every character is the whole space; otherwise its nonempty open complement would have both positive measure and measure zero.

For a finite sum of compactly supported words w let x_χ be its vector in phase χ and Ω_χ its vacuum. The construction supplies a common $M \ge 0$ and, for almost every χ, a unique heat ground state together with the inverse-square estimate for the zero-energy-subtracted heat pairing. In those

phases the zero-energy projection P_χ is the vacuum projection. Indeed the heat attachment sends its range into the one-dimensional heat ground space; normalization and self-adjointness determine its scalar coefficient as $\langle \Omega_\chi, x\rangle$. Thus for each $s \geq 1$ almost every χ obeys

$$0 \leq \langle x_\chi, \exp(-sH_\chi)x_\chi\rangle - |\langle \Omega_\chi, x_\chi\rangle|^2 \leq Ms^{-2}. \tag{4}$$

For fixed s this pairing is continuous in χ, as is the vacuum overlap. The set where (4) holds is therefore closed. Full support extends (4) to every χ, with the same M. This step uses continuity of heat pairings, without first assuming continuity of the ground projections themselves.

Fix χ and let s tend to infinity. The spectral theorem gives strong convergence $\exp(-sH_\chi) \to P_\chi$ for nonnegative H_χ. Equation (4) yields $\langle x_\chi, P_\chi x_\chi\rangle = |\langle \Omega_\chi, x_\chi\rangle|^2$. Since P_χ is an orthogonal projection fixing Ω_χ, expansion of the norm gives

$$\begin{aligned} &\|P_\chi x_\chi - \langle \Omega_\chi, x_\chi\rangle \Omega_\chi\|^2 \\ &= \langle x_\chi, P_\chi x_\chi\rangle - |\langle \Omega_\chi, x_\chi\rangle|^2 = 0. \end{aligned}$$

Thus P_χ acts as the vacuum projection on these compact word vectors. Approximate each Schwartz test in a word by compact cutoffs. Continuity of the joint smeared word map gives convergence of the word vectors, and boundedness of the two projections transfers their equality to every Schwartz word vector. Cyclicity and boundedness then extend it to the whole Hilbert space. This proves rank-one zero energy in every character, supplying the vacuum-projection heat limit used in the weak-commutant argument. The low-energy estimate, character probability construction, heat-pairing continuity and attachment arguments in this chapter provide the ingredients used here.

12.9 Propagation of analytic limits on a complex slice

Here is the analytic limit argument needed when continuing strict Euclidean values to real-time slices. Let $U \subset \mathbb{C}$ be open and connected, let F_a be eventually holomorphic on U, and suppose that at each $x \in U$ some ball $B(x, r)$ has a common bound $|F_a| \leq A$ on a tail of the indexing filter. The tail may depend on the ball. If a holomorphic f is the pointwise limit on a subset accumulating at an interior point of U, then $F_a(y) \to f(y)$ at every $y \in U$.

First derive the requisite local equicontinuity. Choose $s > 0$ with $B(x, s) \subset U$, put $t = \min(r, s)/2$, and $D = \max(1, A/t)$. For $z \in B(x, t)$, the closed disc of radius t about z lies in $B(x, r) \cap U$. The Cauchy derivative estimate gives $|F_a'(z)| \leq A/t \leq D$ on the common tail. Integration along the segment from x to $y \in B(x, t)$ yields

$$|F_a(y) - F_a(x)| \leq D|y - x|.$$

If F_a has pointwise limit g throughout U, passing to the limit gives the identical local Lipschitz estimate for g. Given $\epsilon > 0$, restrict further to $|y - x| < \epsilon/(4D)$ and use $|F_a(x) - g(x)| < \epsilon/4$. The triangle inequality proves local uniform convergence to g; hence g is holomorphic.

To remove the assumption of an everywhere pointwise limit, pass to an arbitrary ultrafilter finer than the original filter. Each scalar family $F_a(z)$ is eventually in a compact closed disc, so this ultrafilter has a limit $g(z)$ for every $z \in U$. The preceding argument makes g holomorphic. At every point of the prescribed accumulation set its value equals f, by uniqueness of scalar limits. The identity theorem gives $g = f$ throughout U. Thus every finer ultrafilter converges to $f(y)$, proving convergence along the original filter. This argument applies equally to the boundary-height filter and to limits of phase characters, and uses only local bounds on tails.

12.10 Bounded slab operators in every character chamber

The operator estimates needed in reconstruction hold at every character of the compression algebra. Their proof does not require discarding exceptional characters. We give the power argument, including the step that removes a character-dependent constant.

Let $X : V \to \mathcal{H}$ map the algebraic span of future words into its reflection-positive Hilbert completion. Let $A, B : V \to V$ satisfy

$$\langle XAv, Xw\rangle = \langle Xv, XBw\rangle.$$

Set $T = BA$. Suppose $D > 0$ and, for every $v \in V$, there is a finite constant C_v such that

$$\|XT^n v\| \le C_v D^{2n} \qquad (n \ge 0). \tag{ST1}$$

The constant need not be uniform in v. The adjoint relation implies symmetry of T, and then of each power, relative to this Gram form. Put

$$d_n = \|XT^n v\| D^{-2n}, \qquad s = \sup_{n\ge 0} d_n.$$

The supremum is finite by (ST1). Symmetry and Cauchy–Schwarz give

$$d_n^2 = D^{-4n}|\langle Xv, XT^{2n}v\rangle| \le \|Xv\| d_{2n} \le \|Xv\| s.$$

Taking the supremum shows $s^2 \le \|Xv\| s$, hence $s \le \|Xv\|$. Finally,

$$\|XAv\|^2 = |\langle Xv, XTv\rangle| \le \|Xv\|\|XTv\| \le D^2 \|Xv\|^2. \tag{ST2}$$

Thus A descends through the null space and extends to a bounded operator of norm at most D on the Hilbert completion.

Apply this to insertion of a slab word and insertion of its reflected word. Translation invariance of the hierarchy gives the displayed adjoint relation. Starting with a future word v, denote by v_n the word obtained by adjoining n successive reflected/unreflected slab pairs. It suffices to have

$$|\operatorname{Gram}(v_n, v_n)| \le Q_v D^{4n}. \tag{ST3}$$

For a scalar multiple av, the Hilbert norm bound is $|a|\sqrt{Q_v} D^{2n}$. A finite linear combination has (ST1) by the triangle inequality, with constant equal to the corresponding finite sum. Hence a diagonal estimate on individual words really does produce a bounded operator on their entire dense span. This proves the slab operator construction without assuming boundedness beforehand.

To obtain (ST3) in a character chamber, let U be the slab operator in the unselected chamber and let X_v be its ground-space word map. Concatenation says

$$X_{\mathrm{enter}(w,v)} = UX_v,$$

while reflection gives U^* for the reflected slab. Therefore

$$X_{v_n} = (U^*U)^n X_v, \qquad \|X_{v_n}\| \le \|X_v\|\|U\|^{2n}.$$

Let χ be the chosen continuous character of the closed compression algebra. The character chamber receives the scalar Gram values of that algebra, and the compressed two-sided word is $X_v^* X_v$. Consequently

$$|\operatorname{Gram}_\chi(v, v)| = |\chi(X_v^* X_v)| \le \|\chi\|\|X_v\|^2.$$

For any $D > \|U\|$, substitution of v_n proves (ST3) with $Q_v = \|\chi\|\|X_v\|^2$. The power argument removes this prefactor, producing a character-chamber slab operator of norm at most D.

All the operators obtained for different $D > \|U\|$ agree on the dense future-word span, where their action is prescribed by insertion. They are therefore the same bounded operator. Letting $D \downarrow \|U\|$ proves the exact estimate

$$\|U_\chi\| \le \|U\|. \tag{ST4}$$

This also covers $\|U\| = 0$. Every unselected arbitrary-copy estimate thus transfers with its original constants to every character, uniformly over the character space.

For each character construct its reflection-positive chamber, these bounded slab maps, and the derivative calculus inherited by the character hierarchy. The transferred arbitrary-copy estimates give the slab laws required to form its contour and strict waves. Apply the field reconstruction to those data to obtain its Hilbert space and fields, with the common conjugation and rotation labels. At the previously selected character, choose the already constructed chamber, slab maps, contour and fields themselves. This is legitimate because they satisfy the same requirements; consequently no comparison between independently chosen copies is needed there. The resulting family supplies the common phasewise operator bounds used in the strict-wave polynomial estimates and in passage to the tested Wightman boundaries. The remaining contour and field construction is developed below; vacuum uniqueness is a separate conclusion and is not assumed in this step.

12.11 Centered insertions and the strict-time waves

Write $E(u)$ for the chamber's heat contraction, $F(\zeta)$ for its holomorphic contraction semigroup on $\Re\zeta > 0$, and $U(x)$ for spatial translation. Thus $F(u) = E(u)$ for $u > 0$; translations commute with the heat flow and fix the unit vacuum Ω. For a field label v and $p > 0$, define its centered, buffered insertion by the bounded slab operator

$$B_p(v) = \mathrm{place}\big(2p, [(0, p; v)]\big). \tag{SW1}$$

Initially v belongs to one finite field space. Its matrix element between any two future-word vectors is the Euclidean hierarchy with that insertion between the reflected left spectators and the right spectators translated forward by $2p$. The spectators and insertion have distinct positions. Multilinearity of the hierarchy in the insertion therefore gives linearity of B_p. If two finite-level labels represent the same element of the common field union, the derivative-plan inclusion identity makes these matrix elements equal. Equality on the dense future-word span then makes the bounded operators equal. Hence (SW1) defines a linear map on the common field union, rather than a map depending on a representative.

The arbitrary-copy slab bound applied to this one-insertion word gives, for each fixed union label v, constants $Q_v > 0$ and $d_v \in \mathbb{N}$ with

$$\|B_p(v)\| \le Q_v p^{-d_v} \qquad (0 < p < 1). \tag{SW2}$$

The word has both boundary buffers equal to p, which is why that estimate applies at this width. Width zero is unused and may be assigned the zero map.

For $w \in \mathbb{C}$ and $s \ge 0$, define

$$R(w, s, x) = F(-iw - s)U(x), \qquad \Im w > s.$$

Its norm is at most one, because $\Re(-iw - s) = \Im w - s > 0$. Let $z_1, \dots, z_n$ have strictly increasing imaginary parts, let $x_i \in \mathbb{R}^3$, and take

$$0 < p < 1, \qquad 4p < \Im z_j - \Im z_i \quad (i < j). \tag{SW3}$$

The buffered matrix element is

$$M_p(z,x;v) = \langle \Omega,\, B_p(v_1)R(z_2-z_1,2p,x_2-x_1)B_p(v_2)\cdots R(z_n-z_{n-1},2p,x_n-x_{n-1})B_p(v_n)\Omega\rangle\,. \tag{SW4}$$

For $n=0$ it is $\langle\Omega,\Omega\rangle$. All factors are defined under (SW3), and this expression is multilinear in the n union labels.

There are two exact identities behind the construction. Slab stretching and concatenation give

$$B_{p+u}(v) = E(u)B_p(v)E(u). \tag{SW5}$$

When the larger width is admissible, the semigroup identity and commutation with spatial translation give

$$R(w,2p,x) = E(u)R(w,2(p+u),x)E(u). \tag{SW6}$$

Insert (SW5) into the larger-width version of (SW4). Each intervening pair of heat contractions combines with its larger-width propagator according to (SW6); the two exterior contractions disappear against Ω. Thus $M_{p+u}=M_p$. Comparing the smaller and larger of any two admissible widths proves complete width independence.

At purely imaginary times $z_i = iy_i$, the same product is exactly the original Euclidean hierarchy at the points (x_i,y_i). To see this directly, build its ket recursively from the final insertion backwards. The centered insertion adds the next plant at time p; the positive heat interval $y_{i+1}-y_i-2p$ and spatial translation move the remaining word into its prescribed relative position. Slab insertion and chamber translation act on planted word vectors by those very operations. Taking the vacuum inner product gives the expectation of the complete word, translated by a common vector. Translation invariance removes this vector, and the quotient-label identity identifies the result with the hierarchy on the common union. This proves the Euclidean identity for every admissible p.

Define $F_n(z,x;v) = M_p(z,x;v)$ on the strict-time region. An admissible p exists because there are finitely many positive imaginary gaps; the choice is irrelevant by width independence. A single such p remains admissible on a neighborhood of any fixed strict-time point. On that neighborhood all propagators in (SW4) vary holomorphically in the complex time variables and continuously in the spatial variables. Bounded composition therefore proves the same properties for F_n. In particular this argument requires no continuity of the arbitrarily chosen width as a function of z.

Repeated contraction bounds and Cauchy–Schwarz show

$$|F_n(z,x;v)| \le \|\Omega\|^2 \prod_{i=1}^{n} \|B_p(v_i)\|.$$

If $0<k\le 1$ and every imaginary gap is at least k, choose $p=k/8$, which satisfies (SW3). From (SW2),

$$|F_n(z,x;v)| \le C_v k^{-m_v}, \qquad m_v=\sum_i d_{v_i}, \quad C_v = \|\Omega\|^2 \prod_i (Q_{v_i} 8^{d_{v_i}}). \tag{SW7}$$

The bound is uniform in all real times and spatial positions. At arities zero and one the pairwise gap condition is empty; the same choice of p is still allowed. No uniformity in arity or in the labels is asserted.

12.12 Tested Cauchy–Riemann equations

For a real spacetime configuration X, let

$$f_{v,y}(X) = F_n\big((t_i(X)+iy_i)_i,(x_i(X))_i;v\big), \qquad \Psi_v(y;\varphi) = \int f_{v,y}(X)\varphi(X)\,dX,$$

where $y_1 < \cdots < y_n$ and φ is a joint Schwartz test. A neighborhood of any such y has one common positive lower gap k. This follows by decreasing the finitely many strictly positive gaps slightly; equivalently subtract a sufficiently small multiple of the increasing height vector. Equation (SW7) therefore bounds $f_{v,u}$, uniformly in X and nearby u, by a fixed constant. The integral defining Ψ_v is absolutely convergent, and dominated convergence proves continuity in y.

For a height direction a, write e_a for the spacetime direction whose i-th time coordinate is a_i and whose spatial coordinates vanish. The one-complex-variable slice

$$h_{X,u}(\zeta) = f_{v,u+(\Im\zeta)a}(X + (\Re\zeta)e_a)$$

is holomorphic wherever its heights remain strictly ordered. A single disc $|\zeta| \le r$ is admissible for every X and all u in a sufficiently small height neighborhood. The same common bound C holds on these discs. Cauchy's estimate gives

$$|\partial_{e_a} f_{v,u}(X)| \le C/r. \tag{CR1}$$

The radius is independent of the noncompact real configuration X; this is essential for integration. Holomorphy of the slice also gives

$$\frac{d}{ds} f_{v,y+sa}(X) = i\,\partial_{e_a} f_{v,y+sa}(X).$$

Bound (CR1) permits differentiation under the joint test integral. Integration by parts along the real direction e_a then yields

$$\frac{d}{ds}\Psi_v(y + sa; \varphi) = -i\,\Psi_v(y + sa; \partial_{e_a}\varphi). \tag{CR2}$$

Both the original integrand and the derivative integrand are absolutely integrable by the Schwartz decay and the uniform bounds above. These hypotheses justify integration by parts directly on the whole real space. Thus (CR2) holds for arbitrary joint tests and arbitrary height directions staying in the strict cone, not just for tensor tests or a single approach ray. Together with (SW7), it supplies the tested Taylor and tempered boundary arguments used earlier.

12.13 The frequency identity and the time spectral cone

Use the Fourier normalization of the construction, denoting its Schwartz transform by $\mathscr{H}$; it includes the fixed coordinate rescaling that removes the 2π from the following derivative identity. If $d = e_y$ is the pure-time spacetime vector belonging to a height vector y, that identity is

$$-i\partial_d(\mathscr{H} g) = -\mathscr{H}(\langle d, \cdot\rangle g). \tag{FT1}$$

For compactly supported smooth g, multiplication by $e^{s\langle d,q\rangle}$ defines a Schwartz test g_s, with the same compact support, for every real s. Moreover

$$\frac{d}{ds} g_s = \langle d, \cdot\rangle g_s$$

in the Schwartz topology. Indeed the divided difference of the scalar exponential is smooth at zero. Its products with g, and every derivative of those products, are continuous jointly in the parameter and the point on a common compact support. Uniform convergence of every weighted derivative follows from compactness; outside that support all those derivatives vanish. This proves the asserted Schwartz derivative, including at an arbitrary parameter value.

Let W_v be the tempered boundary functional of the strict waves. For strictly increasing y, consider

$$A(s) = \Psi_v(sy; \mathscr{H} g_s), \qquad s > 0.$$

The tested Cauchy–Riemann identity and (FT1) show $A'(s) = 0$: the derivative due to the heights is $\Psi_v(sy; -i\partial_d \mathscr{H} g_s)$, and the derivative of the test is its negative. This product rule is valid also near the boundary. To justify moving tests there, the tested Taylor bound supplies one continuous Schwartz seminorm p with

$$|\Psi_v(u; \varphi)| \le p(\varphi)$$

for every strict height u in the boundary approach region. Thus a test difference tending to zero in Schwartz topology contributes a quantity tending to zero uniformly in u. On the interior, the ordinary locally dominated derivative argument applies. Constancy of A on $(0, \infty)$ and the boundary limit at $s \downarrow 0$ give

$$\widehat{W}_v(g) := W_v(\mathscr{H} g) = \Psi_v\Big(y; \mathscr{H}\big(e^{\langle e_y, q\rangle} g(q)\big)\Big). \tag{FT2}$$

This equality uses absolute frequency coordinates; it is not an assertion only in relative coordinates.

Now let y be weakly increasing and let g have compact support in $\{q : \langle e_y, q\rangle < -1\}$. Fix a strictly increasing height vector r with unit gaps, and apply (FT2) at $Y(t) = ty + r$, $t \ge 0$. These heights retain gaps at least one. The tests

$$g_t(q) = e^{\langle e_{Y(t)}, q\rangle} g(q)$$

tend to zero in Schwartz topology. For completeness, each derivative of order j is bounded in a weighted Schwartz seminorm by a finite sum of terms

$$\binom{j}{\ell} e^{-t} (t\|e_y\|)^{\ell}\, p_{i,j-\ell} \qquad (0 \le \ell \le j),$$

where the constants include the fixed factor $e^{\langle e_r, q\rangle} g(q)$. Every term tends to zero. The unit-gap version of (SW7) bounds the right side of (FT2) by a fixed continuous Schwartz seminorm of $\mathscr{H} g_t$, so it tends to zero. Its left side is independent of t, and hence $\widehat{W}_v(g) = 0$.

If a frequency point q has $\langle e_y, q\rangle < 0$ for some weakly increasing y, rescale y so that this pairing equals -2. An open neighborhood of q then has pairing less than -1, and the preceding compact-test argument makes the distribution vanish there. It follows that

$$\operatorname{supp} \widehat{W}_v \subset \bigcap_{y_1 \le \cdots \le y_n} \{q : \langle e_y, q\rangle \ge 0\}. \tag{FT3}$$

The weak ordering, including constant height vectors of either sign, is essential: it enforces total time-frequency conservation as well as the ordered partial-energy inequalities. The passage from this time cone to the relativistic spectrum also uses the rotation and covariance arguments developed next.

12.14 Removing the outer angular flux and obtaining all Euclidean rotations

Continuous Euclidean rotations are obtained from shell Ward identities and control of the outer angular flux. They are not inserted as an extra invariance assumption. The corresponding label representation must descend to the physical quotient and agree with the finite-degree free comparison. Only after this step can the analytic tube identities be converted into the real boost action used in reconstruction.

The lattice symmetries give only signed coordinate permutations. Continuous rotations of the physical quotient require the conserved stress and its shell maps constructed below. Write p_α for the translation shell map and $Y(K)$ for the angular shell map, where K is a real skew 4-by-4 matrix. Before their normalization, the angular Ward expression is

$$\mathcal{A}_K(v;x)=\sum_\ell \{S(v_1,\dots,Y(K)v_\ell,\dots;x)+(Kx_\ell)^\alpha S(v_1,\dots,p_\alpha v_\ell,\dots;x)\}. \tag{RT1}$$

Its vanishing cannot be inferred by simply enlarging a spherical cutoff: the correlations have subpower infrared bounds, rather than the decay that such an argument would require. The source removes the outer flux through tested spectral estimates. We give the details of that passage here.

The tests in this argument have the following precise meaning. Fix a nonnegative even smooth compactly supported function c, equal to one on $[-1,1]$, and let

$$c_w(p)=w^{-1}c(p/w), \qquad w>0.$$

The source uses this central frequency profile and real Schwartz profiles with the same normalization. Its j-th moment is $\int p^j f(p)\,dp$; substitution gives the moment of f_w as w^j times that of f. In the Fourier convention used here,

$$\partial_x^j \widehat{f}(0)=(-i)^j\int p^j f(p)\,dp.$$

Differentiation under the integral is justified by the integrable Schwartz majorant $|p|^j|f(p)|$. If moments below m vanish, repeated integration of the derivative along a segment bounds $|\widehat{f}(x)|\le C|x|^m$ on a fixed ball. This is the origin-jet cancellation used in the compensated profiles.

For four one-dimensional tests f_a, the tested stress pair is the product-test extension of

$$R_{ab,cd}(f)=\int \langle J_{ab}(x)J_{cd}(0)\rangle \prod_i f_i(x_i)\,d^4x. \tag{RT0a}$$

When a factor vanishes sufficiently far in Taylor order at zero this is the ordinary integral. For arbitrary Schwartz factors, subtract the finite product Taylor polynomial near zero; the polynomial short-distance bound makes the remainder integrable. This defines the extension without presuming rotation invariance. Pair exchange, tensor symmetry and signed covariance follow by the corresponding changes of variables in the remainder integral.

Its tested conservation defect is

$$\Delta_{b,cd}(f)=\sum_i R_{ib,cd}(f_0,\dots,f_i',\dots,f_3).$$

Let $g_j(x)=x^j c(x)/j!$, and let $f^{(a)}$ have factor g_1 at a and g_0 elsewhere. Set

$$w_{ab,cd}=\Delta_{b,cd}(f^{(a)}), \qquad R^\circ_{ab,cd}(f)=R_{ab,cd}(f)-w_{ab,cd}\prod_i f_i(0). \tag{RT0b}$$

The subtracted term is supported at the single contact $x=0$. Its differentiated value on $f^{(a)}$ is exactly $w_{ab,cd}$, because $g_j^{(\ell)}(0)=1_{\{j=\ell\}}$. The resulting defect therefore vanishes on all first-order product jets. It vanishes on the order-zero and order-two jets by total inversion: the three remaining tensor indices give an odd sign whereas an even total test order gives an even sign. These finite contact subtractions preserve the pair and signed symmetries; reflection in a single axis kills a coefficient unless its four tensor indices occur in pairs. Thus the compensated conservation defect

has no jets of total degree below three. This precise contact remainder, rather than an assumed exact conservation law at coincident points, enters the spectral band estimates.

For a fixed exterior word on coordinate line l, put

$$F_{db}(y) = \langle J_{db}(y) \prod_{\ell} B_{\ell}(x_{\ell}) \rangle.$$

Realize that exterior word as a finite word vector and pair it with the one-stress vector. The hierarchy's permutation and conjugation laws identify this pairing with F_{db} whenever the points are distinct. This is why the spectral bounds apply to an arbitrary exterior word, not only to the stress two-point function. The actual outer term is

$$\int (\partial_d \phi_{\lambda}^{ab,l})(y) F_{db}(y)\, d^4y, \qquad d \neq a, \tag{RT0c}$$

for the compensated angular primitive $\phi_{\lambda}^{ab,l}$. Its frequency widths are λ^{-100} in axes a, b, λ^{-150} on the remaining line axis l, and λ^{-300} otherwise; it carries the primitive amplitude λ^{100}. The band subtractions described next give the actual estimate $C\lambda^{-1}$ for (RT0c), including its integrability. The finite excluded insertion set has Lebesgue measure zero and does not change this integral. Compact shell conservation identifies the difference between the summed local charges and this outer integral. Its vanishing therefore proves (RT1).

For each coordinate i, reflection positivity supplies its own positive joint spectral measures ν_i^A, supported on $E = p_i \geq 0$, for finite linear combinations A of reflected words. They include atoms at $E = 0$. Mixed readings are obtained by polarization and obey the weighted Cauchy–Schwarz inequality

$$|\langle B, f(P)A \rangle| \leq \frac{1}{2} \left(t \int X\, d\nu_i^B + t^{-1} \int Z\, d\nu_i^A \right) \quad \text{if } |f|^2 \leq XZ, \quad t > 0. \tag{RT2}$$

This follows by applying $2uv \leq tu^2 + t^{-1}v^2$ to the two spectral vectors; approximation by simple spectral functions proves it for the displayed measurable weights. No identification of spectral measures for different axes is needed.

All transverse moments needed below are finite after positive time damping. To prove this, take the characteristic function of the finite damped measure and apply k pairs of forward and backward differences in a transverse direction j. Its integral multiplier, up to the cancelling sign of the difference pair, is

$$\left| \frac{e^{ihp_j} - 1}{h} \right|^{2k}.$$

The fundamental theorem of calculus bounds the difference expression by the corresponding $2k$-th derivative of the smooth separated correlation. Fatou's lemma as $h \to 0$ gives the moment of $|p_j|^{2k}$. The same argument for moved word clusters uses their actual positive time gaps. For a cluster collapsed in the i direction this yields, for some finite D, C depending on the word and k,

$$\int e^{-tE} |p_j|^{2k}\, d\nu_i^A \leq Ct^{-D}, \qquad 0 < t \leq 1. \tag{RT3}$$

Thus moment orders can be chosen before shrinking the spectral bands.

Here is the interpolation that controls low energy. For distinct nodes $z_j = (j+2)^2$, let L_j be their Lagrange basis polynomials, $0 \leq j < m$. Polynomial interpolation gives

$$\frac{1}{x+y} - \sum_{j<m} \frac{L_j(x)}{z_j + y} = \frac{1}{x+y} \prod_{j<m} \frac{z_j - x}{z_j + y}, \qquad y \geq 0, \quad x + y \neq 0. \tag{RT4}$$

For completeness, multiply the left side by $x + y$. It is a polynomial of degree at most m, zero at every z_j, and equal to one at $x = -y$. These $m + 1$ values determine it and give the product on the right. Rescale $x = (p/a)^2, y = (E/a)^2$. Multiplying by the even Laplace numerator $2E$ gives the exact response after subtracting the finite high-band interpolant. For moment-controlled profiles the subtracted moments are exactly those furnished by (RT4). The residual response is nonnegative for $E \geq 0$, and the choice of the central profile gives

$$\mathcal{R}_{m,a}(E) \geq \frac{c_m}{E + a}, \qquad 0 \leq E \leq 1, \quad 0 < a \ll 1, \tag{RT5}$$

with $c_m > 0$. At zero this is the continuous value of the full response, including the profile's zero-energy contribution; one does not divide a bare factor $2E$ by E and discard an atom. The numerator factors in (RT4) stay of one sign because the high nodes lie outside the primary profile band. Their fixed ratios give c_m; their denominators give the stated reciprocal bound.

The reason this positive response is small is a sign identity for the same tested position correlation. Suppose A has reflection parities ϵ_i, ϵ_j of opposite sign on axes i, j. Write

$$P_i(h; f) = \int \mathcal{L}_i(h)(p_i) \prod_{l \in Q \setminus \{i\}} f_l(p_l)\, d\nu_i^A.$$

Here f_l are nonnegative even Schwartz frequency profiles and $\mathcal{L}_i$ is the real even Laplace response of the inverse Fourier profile in direction i. The two spectral formulas for that one tested correlation give $P_i = -P_j$. If some transverse axes are absent, insert profiles equal to one on larger and larger intervals; bounded dominated convergence gives the same identity for that subset Q. Split each active profile into its Lagrange interpolant and residual. Both residuals have nonnegative response. Consequently the primary residual on axis i is bounded by the sum of the two cross terms containing an interpolant. The interpolants are finite sums of profiles in a high band.

For clarity the choice of orders in this repeated operation is important. Put $\lambda = e^N$, and scale each frequency profile as $f_l^{(N)}(p) = \lambda^{-w_l} f_l(p/\lambda^{w_l})$, $w_l \leq 0$. A bound $C\lambda^\gamma$ will be called a bound of exponent γ in this paragraph. If a high band has exponent $a' > 0$, a transverse $2k$-th moment from (RT3) bounds its far remainder with exponent at most

$$-2ka' + K|Q \setminus \{i, j\}| + \max(K, dD), \tag{RT6}$$

where $-K \leq w_l$ and d is the earlier chosen band exponent. Given any $\eta > 0$, choose k so that the negative term exceeds the other fixed losses, and then choose $d > 0$ so small that $dD < \eta$. At each switch the coordinate just treated is removed from the list still requiring a bound. If its parity agrees with the next required coordinate, insert an opposite-parity coordinate first. A finite induction on that list terminates; the terminal case has only nonshrinking transverse widths and is bounded by (RT3). Choose moment orders and admissible band exponents backwards through this finite list. This proves an exponent η, for every $\eta > 0$, for the residual response. Equation (RT5) therefore gives

$$\int_{0 \leq p_i \leq 1} \frac{\prod_{l \in Q \setminus \{i\}} f_l^{(N)}(p_l)}{p_i + \lambda^q}\, d\nu_i^A \leq C_\eta \lambda^\eta, \qquad q < 0, \tag{RT7}$$

when the collapsed cluster has both parities on Q.

A word on a coordinate line l is collapsed on the other axes but can have a genuine time gap on l. Subtract there a shifted positive-time profile. Its response is positive, dominates a fixed multiple of $(E + \lambda^q)^{-1}$ on $0 \leq E \leq 1$, and its damped transverse moments cost at most λ^{-q}, by (RT3) at

the actual gap. The same switch pays this loss by increasing k in (RT6). Thus (RT7) holds on $i = l$. On $i \neq l$, bound the additional l-profile by its supremum $C\lambda^{-w_l}$. The resulting bound has exponent $\eta - w_l$. This proves the line estimate even when the time-axis cluster is not collapsed.

For a diagonal stress reading a further estimate is necessary. Subtract its finite position Taylor jet with a cutoff equal to one at zero. The remainder vanishes to the required order: repeated integration along the segment from zero bounds its k-th derivative by $C|x|^m$ when derivatives $k, \dots, k+m-1$ vanish. Tensor products of these subtractions preserve reflection parities. Conservation, followed by integration by parts on these tests, gives the three tested frequency expressions

$$B = p_i^2 A + R_1, \qquad C = p_i^2 B + R_2. \tag{RT8}$$

On profiles of width at most $W = \lambda^a$, $a \leq 0$, the remainders have respectively bounds $O(W^4)$ and $O(W^6)$, with an arbitrarily small exponent loss. On profiles with the prescribed zero moments the remainder is exactly zero: it is a finite contact Taylor jet and each of its moments is zero.

Apply (RT4) at widths $W, 2W$ to A, B, C. After division by W^{2k} at levels $k = 0, 1, 2$, denote their six band readings by a, b, v, V, r, R, and their residual differences by f, g, h. The positive spectral responses have signs $f \geq 0, g \leq 0, h \geq 0$. Equations (RT8) and the two-band subtraction give small bounds for

$$v - a, \quad V - 4b, \quad r - v, \quad R - 4V, \quad f - (a - b), \quad g - (v - V), \quad h - (r - R).$$

These imply a bound for f with the same exponent. Indeed if the sum of the seven absolute errors is e, these equalities within error e imply $a - b \geq -e$, $a - 4b \leq 3e$, $a - 16b \geq 5(-e)$; eliminating a, b bounds f by $100e$. In the exact case they force $a = b = f = 0$; the numerical factor safely bounds the finite error elimination. The constant-moment term of each interpolant cancels between W and $2W$; every other term has at least a factor W^2. Combining this bound with the positive residual lower bound gives

$$\int_{0 \leq E \leq 1} \frac{EW^4}{(E+W)^4} \prod_{j \neq i} f_j^{(N)}(p_j)\, d\nu_i^{J_{ii}} \leq C_\eta \lambda^{4a+\eta}. \tag{RT9}$$

This deduction uses positivity of spectral measures and of the indicated scalar responses. It does not differentiate a positive matrix or posit positivity of its derivatives.

Now approximate the plane Killing field $x_a e_b - x_b e_a$ by compactly supported smooth vector fields. Choose frequency widths λ^{-100} on axes a, b, λ^{-150} on the remaining line direction if present, and λ^{-300} on the other axes. An odd primitive in direction a, multiplied by λ^{100}, and its partner with a, b interchanged give the two components. Their derivatives in the a, b pair agree exactly. Since $J_{ab} = J_{ba}$, these two terms in $\sum_{i,j} \partial_i F_j J_{ij}$ cancel before any integral is split. The remaining derivatives are supported outside the insertion neighborhoods, where no contact term is present.

Split each such tested spectral integral into $E \leq 1$ with transverse frequencies inside the box widened by λ, and its complement. On the complement, the Schwartz bound of order k gives an arbitrary factor λ^{-k} in a transverse direction. For $E > 1$ the actual exterior time gap gives e^{-E/λ^w}; for every fixed R, M this is at most $\lambda^{-M} e^{-RE}$, eventually. The finite damped moments from (RT3) then pay the whole complement. Inside the box use (RT2) with the two reciprocal weights from (RT7) for shear components, and (RT7),(RT9) for a diagonal component. Split the latter once more at $E = \lambda^{-3000}$. Below that threshold (including zero), the large gain in the line estimate pays a crude stress bound; above it (RT9) applies. With the stated widths these estimates, including the primitive amplitude λ^{100}, leave a bound $C\lambda^{-1}$. Thus the exterior flux tends to zero.

Here are explicit definitions of the responses and contact expressions in this argument. For a real Schwartz frequency profile f, write $D_s f(p) = s^{-1} f(p/s)$, $m_q(f) = \int p^q f(p)\, dp$, and $\widehat{f}(t) =$

$\int e^{-ipt} f(p)\,dp$, with the Fourier constants used in (RT2)–(RT4). Choose an even nonnegative smooth u of integral one supported in $1 \le |v| \le 2$, and set $u_k(v) = v^{-2k}u(v)$, which is smooth because its support misses zero. Write $L_j(X) = \sum_k l_{jk}X^k$ for the node polynomial in (RT4). The actual finite interpolant is

$$B_{m,s}f = \sum_{j<m} \sum_{k} l_{jk} m_{2k}(f) s^{-2k} D_s D_{j+2} u_k, \qquad R_{m,s}f = f - B_{m,s}f. \tag{RT11}$$

Integrating against a continuous scalar function and changing variables in each term turns its kernel into $\int u(v) \sum_j L_j((p/(sv))^2) g(sv(j+2))\,dv$. Putting $g(p) = (p^2+E^2)^{-1}$ and using (RT4) gives

$$\mathcal{R}_{m,s,f}(E) = 2E \int u(v) \int \frac{f(p)}{p^2+E^2} \prod_{j<m} \frac{(sv(j+2))^2 - p^2}{(sv(j+2))^2 + E^2}\,dp\,dv \quad (E>0). \tag{RT12}$$

It is the real part of $\int e^{-|t|E} \widehat{R_{m,s}f}(t)\,dt$, not a rational expression assigned independently at $E=0$. The latter integral gives its continuous value there, namely pref $f(0)$, since the interpolant vanishes near zero. The odd-parity response uses $\int \operatorname{sgn}(t) e^{-|t|E} \widehat{f}(t)\,dt = -2i \int pf(p)/(p^2+E^2)\,dp$ for $E>0$. Its subtraction is obtained by applying the even interpolant to $pf(p)$ and dividing the result by p; division is on support away from zero. Even and odd residuals have the required first $2m$ moments zero.

For $f \ge 0$ supported in $[-s,s]$, every product factor in (RT12) lies between zero and one, because $|v| \ge 1$ and $j+2 \ge 2$. Each is at most $(2(j+2))^2(s/E)^2$; if $s \ge 1$, it is at least $\frac{1}{2}(1+E)^{-2}$. The lower estimate follows from $(sv(j+2))^2 \ge 4s^2$ and $p^2 \le s^2$. For the central profile $f = D_w\chi$, where $\chi = 1$ on $[-1,1]$, integration just over $0 < p < wE/(E+w)$ gives

$$2E \int \frac{D_w\chi(p)}{p^2+E^2}\,dp \ge (E+w)^{-1}.$$

Consequently

$$\mathcal{R}_{m,s,D_w\chi}(E) \ge 2^{-m}(E+w)^{-1}(1+E)^{-2m}. \tag{RT13}$$

Continuity proves the same inequality at zero; on $0 \le E \le 1$, (RT5) follows with $c_m = 2^{-3m}$. The fixed subtraction band $s \ge 1$ and the shrinking primary width w are distinct. At high energy the product upper bound gives, for $q \le 2m$, a bound $C_{m,q}(w^{-1}+s^q)(1+E)^{-q}$. These are the precise positive majorants used in the parity switch. For an arbitrary profile of either parity, a nonnegative even compact majorant and the odd division construction bound its subtracted response by a constant times an even response of this form.

To specify the contacts in (RT8), let $R_{ab,cd}$ denote the compensated product-test distribution (RT0b), let $T_{ab,cd}$ be its value after Fourier transforming all four test slots, and let $Q_{bcd} = i\,\text{Fourier}(\sum_j \partial_j R_{jb,cd})$. Writing P_j for multiplication of frequency slot j by p_j, the source identity is $Q_{bcd} = \sum_j P_j T_{jb,cd}$. Thus the quantities called A, B, C in (RT8) are

$$A = T_{ii,ii}, \qquad B = \sum_{j,a\neq i} P_j P_a T_{ai,ji}, \qquad C = \sum_{j,a,b,c\neq i} P_j P_a P_b P_c T_{cj,ab}. \tag{RT14}$$

Their contact remainders are exactly

$$R_1 = \sum_{j\neq i} P_j Q_{iji} - P_i Q_{iii}, \qquad R_2 = \sum_{j,a,b\neq i} P_j P_a P_b Q_{jab} - P_i \sum_{j,a\neq i} P_j P_a Q_{aji}. \tag{RT15}$$

Expand $Q = \sum PT$, use the exchange symmetries of the two tensor pairs, and commute the scalar multipliers P_j; this gives $B = P_i^2 A + R_1$ and $C = P_i^2 B + R_2$ exactly. The compensated Q has no

contact jet of total order below three by (RT0b). On widths $\le \lambda^a$ this costs $3a$ with arbitrarily small exponent slack. Multiplication by one frequency costs a and by three costs $3a$; hence (RT15) costs $4a$ and $6a$, respectively. This proves the powers asserted after (RT8). Vanishing on profiles with the required moments follows directly by evaluating the finite contact polynomial on those moments.

The exterior-flux powers can also be made explicit. Put $E = p_d$, $H_N(p) = \prod_{j\neq d} D_{\lambda^{1+w_j}}\chi(p_j)$, and

$$D_q(p) = \mathbf{1}_{E\le 1}\frac{H_N(p)}{E+\lambda^q}, \qquad Z_z(p) = \mathbf{1}_{E\le 1}\frac{E\lambda^{4z}}{(E+\lambda^z)^4}H_N(p). \tag{RT16}$$

The exterior-word integral of D_{-3000} has exponent $u \le 150$: $u = 1$ if $d = l$, and $u = -(1+w_l)+1$ otherwise, by the line estimate. For a shear stress, the stress integral of D_{w_d} has exponent one. The squared response is at most $C^2\lambda^{2(3+w_d)}D_{-3000}D_{w_d}$, by $e^{-E/\lambda^{w_d}} \le \lambda^{w_d}/(E+\lambda^{w_d})$. Weighted Cauchy–Schwarz therefore gives exponent $3 + w_d + (u+1)/2$. If $d = l$, then $u = 1$ and $w_d \le -150$, giving at most -146. If $d \neq l, a, b$, then $w_d = -300$ and $u \le 150$, giving at most -221.5. These exhaust the shear case $d \neq a, b$; both bounds are stronger than exponent -101.

If $d = b$, split instead at $E = \lambda^{-3000}$. On the upper part use Z_{-99}. Its stress integral has exponent $1 + 4(-99) = -395$, by (RT9) with slack one, and the squared response is at most $C^2\lambda^6 D_{-3000}Z_{-99}$. The resulting exponent is at most $3+(150-395)/2 < -101$. On the lower part use the crude stress weight D_0, whose integral has exponent at most $3\cdot 300 = 900$. The squared response is at most $C^2\lambda^{2(3-1500)}D_{-3000}D_0$, giving exponent at most $3-1500+(150+900)/2 < -101$. These pointwise estimates are valid on the nonnegative spectrum and include $E = 0$. The complement of the low box has arbitrary negative exponent by the damped moments and Schwartz tails discussed after (RT9). Thus the spectral reading has exponent -101. Multiplying by the angular primitive's λ^{100} leaves exponent -1, as claimed in (RT0c). Decomposition into the finitely many reflection parities preserves the bound; parity-mismatched terms integrate to zero by reflecting their odd axis.

The compact cutoff flux identity is therefore (RT1) equal to zero for every separated tuple on a coordinate line. Translation along that line changes Kx_ℓ by the fixed vector Ke_l and leaves the correlation unchanged. Subtracting the two angular identities gives the corresponding sum of the translation shell maps equal to zero. Empty tuples give empty sums throughout.

12.15 The order of choices in the spectral sign walk

We record the quantitative induction that supplies the small mixed-band bound in (RT6). Let Q be the finite set of axes being tested, with two signs assigned so that each axis has an opposite-sign partner. Let $U \subseteq Q$ be the axes whose width exponent is still allowed to be negative, and let i be the active axis. A finite route alternates signs and removes the currently visited axis from U. Such a route exists by induction on $|U|$: if the active axis is in U, remove it and move to an opposite partner. If it is not, choose an axis still in U; reach it directly if its sign is opposite, and otherwise take one opposite-sign detour. This reduces the nonempty set after at most two such moves.

The induction statement is uniform in a width vector w with $-K \le w_j \le a$, $w_i \le d$, and $w_j \ge 0$ outside U, for every $0 \le d \le a$. It asserts a residual budget $\delta > 0$ after sufficiently high moment order, for nonnegative even compact profiles and all fixed support factors. If U is empty, the direct spectral moment bound of order zero has a finite price exponent D. Choose $a > 0$ with $aD < \delta$. The nonnegative widths then pay its residual by that budget.

At an induction step from axis i to an opposite-sign axis j, first obtain the future bandwidth $a' > 0$ from the induction hypothesis for $U \setminus \{i\}$. Put $p = K|Q \setminus \{i,j\}|$ and $z = \max(K, \delta)$. Choose an integer

$$k > (p+z)/a'. \tag{RSW1}$$

Only now purchase the moment estimate of order k, with its finite price exponent D, and choose $0 < a < a'$ with $aD < \delta$. The positive profile majorant constructed by the finite interpolant bounds the baked profile and keeps the remainder nonnegative on the spectral half-line. The far part of the two-axis comparison has exponent at most

$$-2ka' + p + \max(-w_i, dD) \le -2ka' + p + z < 0. \tag{RSW2}$$

The near part is paid by the future hypothesis after replacing the active profile by that majorant and its exponent by the nonnegative d. The two residuals have nonnegative spectral integrands. The signed comparison inequality from (RT5) therefore bounds their sum by the two paid positive quantities; no spectrum is transferred from one axis to another. Moment orders and thresholds can be enlarged to their maximum. Induction along the finite route gives the bound for arbitrary initial nonpositive widths, with one sufficiently small positive a. This explains why the future bandwidth is chosen before the present moment order and why the contact-price exponent is chosen before the present final bandwidth. It supplies the precise quantifiers behind the mixed-band estimate used to derive the collinear Ward identities.

12.16 Normalized two-band diagonal estimate

For clarity, the six normalized readings used in (RT8)–(RT9) have the following common normalization. Let $W = \lambda^a$, $a \le 0$, and let $q_k^{(pW)}$ denote the order-k profile from (RT4) at width pW, with all other profiles fixed. Put

$$a_0 = \operatorname{Re} A(q_0^{(W)}), \quad b_0 = \operatorname{Re} A(q_0^{(2W)}),$$
$$v = W^{-2} \operatorname{Re} B(q_1^{(W)}), \quad V = W^{-2} \operatorname{Re} B(q_1^{(2W)}), \quad r = W^{-4} \operatorname{Re} C(q_2^{(W)}), \quad R = W^{-4} \operatorname{Re} C(q_2^{(2W)}). \tag{DB1}$$

The normalization on the second band is still by W, not by $2W$. Multiplication of the order-$k+1$ profile by the frequency square gives $(pW)^2$ times its order-k predecessor. Consequently (RT8) gives small errors for $v - a_0, V - 4b_0, r - v, R - 4V$. The first contact error has unnormalized order W^4, the second order W^6; their respective normalizations leave order W^2, including the arbitrary small exponent slack. The zero-moment contact subtraction has the same bound: its constant interpolation coefficient cancels between the two bands, and every nonconstant even coefficient has at least a factor W^2.

Let f, g, h be the real spectral gap readings at these three orders, normalized by $1, W^{-2}, W^{-4}$. Their signs are $f \ge 0, g \le 0, h \ge 0$. The difference between a gap profile and its two-band difference is exactly the finite interpolation debris just bounded. Hence the additional errors $f - (a_0 - b_0), g - (v - V), h - (r - R)$ have the same small budget. Let K be the sum of the absolute values of these seven errors. From $g \le 0$ we get $a_0 - 4b_0 \le 3K$; from $h \ge 0$ we get $a_0 - 16b_0 \ge -8K$. Therefore $12b_0 \le 11K$, and

$$0 \le f \le a_0 - b_0 + K \le 3b_0 + 4K \le 7K \le 100K. \tag{DB2}$$

Finite sums and fixed scalar multiples preserve the budget exponent, so (DB2) proves the small diagonal gap bound. This is the algebraic step in (RT9); it needs positivity of the three response signs and the contact bounds, and makes no positivity claim for the six individual band readings.

12.17 Ward kernel transfer and the rotation action

We explain how these identities extend to general tuples and to a group action. Differentiating a shell integral is allowed because its fixed compact annulus remains separated under small movements of the distinguished point. Differentiation acts both on the correlation and on the centered

angular coefficient. The latter derivative is the relevant constant translation coefficient. Hence every position derivative of a line Ward test is again a finite sum of line Ward tests for successor-degree labels. The all-order scaling comparison (WS7) transports these tests to the free quotient, where the translation and angular actions are explicit derivatives of polynomials. Finitely many separated reading maps suffice at each finite degree, by the detector argument in (WS4)–(WS5). Their matrices converge to injective free matrices. This transfers the free rank statements for the spaces of possible internal defects and forces any such defect to vanish. Thus the translation shell map is one positive scalar c times the already constructed derivative, and the angular rows, divided by c, satisfy the Ward identity at every separated tuple.

The finite-dimensional rank arguments just used have two different domains. For translation rows, the domain at cap d is the real subspace of $\prod_{i=1}^{4} \mathrm{Hom}_{\mathbb{C}}(\mathcal{F}_d, \mathcal{F}_{d+1})$ consisting of rows that kill the unit, preserve each lower filtered image, commute with conjugation, and transform as a vector under signed permutations. A reading is a position derivative at a fixed separated collinear tuple of the Ward sum $\sum_\ell S(v_1, \dots, T_i v_\ell, \dots, v_n)$. The free readings have common kernel precisely the real line spanned by the derivative row. Select finitely many such readings by minimizing the dimension of their common kernel. Adjoining a further reading cannot lower that minimal dimension, so this finite kernel already equals the full kernel. Pullback by U_k turns the physical readings into matrices on this one fixed real space; their convergence is (WS7). Eventual rank and the two known rows in their kernels then give the translation normalization stated above.

For the internal angular defect, the domain is instead the real subspace

$$\mathcal{I}_d \subset \mathrm{Hom}_{\mathbb{R}}(\mathfrak{so}(4), \mathrm{Hom}_{\mathbb{R}}(\mathfrak{so}(4), \mathrm{End}_{\mathbb{C}}\, \mathcal{F}_d))\,. \tag{RT17}$$

Its elements $T(K, L)$ are alternating in K, L, kill the unit, commute with conjugation, preserve every lower filtered image, and intertwine simultaneous conjugation of K, L by orientation-preserving signed permutations with their action on labels. The reading for a separated tuple, labels v_ℓ, and K, L, is $\sum_\ell S(v_1, \dots, T(K, L)v_\ell, \dots, v_n)$. The free internal Ward separation says the common kernel of these readings is zero at the required sufficiently large caps; lower caps follow by inclusion. Select finitely many readings by the same decreasing-kernel-dimension argument. Their joint map is injective. Transport by $T \mapsto U_k T U_k^{-1}$, retaining both Lie algebra slots, identifies the actual and free domains; reality, signed covariance, unit preservation and lower-image preservation follow from (WS1). On fixed free basis inputs the actual reading matrices converge to the free ones by order-zero joint comparison. Linear independence of finitely many matrix columns is open, so they are eventually injective. This permits testing the unknown, scale-dependent pulled-back defect without assuming it is bounded. Its readings vanish by the Ward identity; injectivity makes the defect itself zero.

The free internal separation used here has a concrete angular argument. Identify fields with their Wick polynomials as in (FM3)–(FM7). For a fixed pair K, L, the assumed zero Ward sums say that the corresponding polynomial endomorphism T is silent in every moment; in particular its two-point pairings obey

$$\langle Tp(x), q(0)\rangle + \langle p(x), Tq(0)\rangle = 0. \tag{RT18}$$

It preserves the degree-four subspace by the lower-image condition. The explicit curvature contraction (FM2), expanded on its degree-four invariant polynomials, gives at equal radii $x, y \neq 0$ the identity

$$\langle p(x), q(0)\rangle = \langle (\rho(H_y H_x)p)(y), q(0)\rangle, \tag{RT19}$$

where $H_z = I - 2zz^{\mathsf{T}}/|z|^2$ is reflection in the hyperplane normal to z, and ρ acts on all tensor legs. This follows by substituting the orthogonal factor $I - 2zz^{\mathsf{T}}/|z|^2$ in the degree-four Wick pairings;

the radial factors agree, and contractions use the invariant color inner product. The explicit low-degree evaluations and the extension to every filtered cap are (FS6)–(FS14), with (WE1)–(WE4) and (HC1)–(HC4) supplying their concrete evaluation and collision steps. Apply (RT18) at x and y, use (RT19) in both terms, and subtract. Positive definiteness of the reflected polynomial pairing at $x = -re_i$ forces T to commute on degree four with $\rho(H_x H_y)$. Taking $y = e_i + e_j$, with the matching radius, makes this product a signed quarter-turn in the ij-plane.

Signed covariance of the two Lie algebra slots now makes $T(K, L)p$, for each degree-four p, invariant under simultaneous quarter-turn conjugation of K, L. An alternating bilinear map with this invariance is zero. To check this without representation theory, use the plane basis E_{ij} of $\mathfrak{so}(4)$. Equal arguments vanish by alternation. For two planes sharing one axis, the square of a quarter-turn in the remaining two axes preserves one plane and negates the other. For disjoint planes, two quarter-turn substitutions exchange the arguments with the sign which, combined with alternation, negates the value. Every coefficient is therefore its own negative. This proves that the internal endomorphism kills the whole degree-four subspace, including the free stress.

The local angular Ward identity for the free stress then propagates this to higher degrees: insert the stress in the silent moment sum, apply the compact angular shell integral, and use that its T-term is zero. The remaining identity is the moment pairing of the commutator of T with the free angular generator. Nondegeneracy of the Wick polynomial tests, proved by the reflected Gram argument following (FM7), makes this commutator zero. Finite-dimensional ODE uniqueness exponentiates that commutation to each quarter-turn at the higher cap. Signed covariance and the same plane-basis calculation again force $T = 0$. This proves the separation of $\mathcal{I}_d$ for every $d \geq 4$ used in (RT17), rather than assuming an additional internal symmetry classification.

For the translation detector there are two further points behind its free one-dimensional kernel. A separated target tuple can be joined to a tuple on any prescribed coordinate line by a segment staying in configuration space: order the line coordinates by the target coordinates on that axis, break ties by distinct arbitrarily small increments, and interpolate. Strict inequalities prevent a collision along the segment, and tied target coordinates remain separated until the endpoint. All position derivatives of the Ward sum vanish at the line tuple when all the scene readings vanish. Analyticity of the explicit Wick kernels gives zero in its neighborhood and then along this segment. Thus scene vanishing is the full silent-moment identity. The color-vector, reality and filtration restrictions of (RT17)'s translation analogue then give precisely the derivative line: the free Wick coefficient calculation identifies its degree-four charge as a real scalar translation charge; subtract that scalar derivative. The remaining silent row kills the free stress. The radial shell identity makes its commutator with the engineering-weight derivation a ghost. The fixed exterior-cap separation (FS13) kills that commutator, and the inversion-parity calculation (FS14) kills the row on every higher cap. The complete charge and polynomial-separation calculations are (FS1)–(FS14), built on the hierarchy (FM1)–(FM7) and the invariant spaces (FC1)–(FC15). The zero scalar case is included.

Apply two orbital derivatives to that identity and subtract in opposite orders. The vector fields $x \mapsto Kx$ have bracket with the opposite sign to the matrix bracket, while the substitutions in a multilinear correlation have the same opposite sign. The resulting defect is $[Y(K), Y(L)] - cY([K, L])$, tested in the internal defect space just eliminated. It is zero. This proves the Lie relations on the actual quotient.

There is a useful additional detail in fitting this representation to the free one: the deformed matrices need not be normal. Put $D = c^{-1}Y$, and work at one finite degree in the space of linear maps between the two quotients. Its relative generators are $R_K(T) = D(K)T - TY_{\mathrm{free}}(K)$. Their sum of squares commutes with every R_K: expand the commutator, use the plane matrix brackets, and pair the two opposite ordered-plane terms. On the free side this sum is negative semidefinite

in the invariant inner product. A real polynomial in it projects onto the zero eigenspace, by interpolation on its finite real spectrum. Evaluate that same polynomial on the nearby deformed sum of squares. The resulting matrices P_k converge to the free projection P, commute with all the deformed generators, and obey $P_k^2 - P_k \to 0$; they are not yet projections.

Correct them by iteration of $\phi(Z) = 3Z^2 - 2Z^3$. If $e(Z) = Z^2 - Z$, direct polynomial multiplication gives

$$\phi(Z) - Z = (1 - 2Z)e(Z), \qquad e(\phi(Z)) = (4e(Z) - 3)e(Z)^2. \tag{RT10}$$

For bounded Z and sufficiently small initial defect, the defect decreases at least geometrically by $1/2$, and the increments are bounded by a fixed constant times that defect. The iterates are Cauchy and converge to an exact idempotent Q_k, with $Q_k - P_k \to 0$. Polynomial iteration preserves every commutation and every closed real symmetry constraint.

It remains to show that the relative generators vanish on its range. Each plane generator is the bracket of two other plane generators. Since Q_k commutes with them and is idempotent, the restricted matrices $R_i Q_k$ satisfy those same perfect bracket identities. They tend to zero because the free generators kill P. If M is the largest of their finitely many norms, the bracket identities give $M \le 2M^2$. Eventually $M < 1/2$, hence $M = 0$. Apply Q_k to the comparison map corresponding to the identity: its limit is the identity, so the result is invertible for large k and is an exact intertwiner. Reality and signed permutations preserve it, because signed conjugation permutes the plane generators with signs whose squares are one. This proves the required compatibility without assuming a spectral theorem for the deformed matrices.

Transport the free orthogonal representation by this intertwiner. Its exponential derivative is $c^{-1}Y(K)$. Along $x_\ell(t) = e^{tK}x_\ell$ and the simultaneously moved labels, the ordinary product rule and the Ward identity show that the derivative of the correlation is zero. The product rule is valid jointly in positions and labels: expansion in fixed finite bases is a finite sum of products of smooth position functions and polynomial coefficient functions. Thus the correlation is constant along every plane rotation. Plane rotations align any unit vector with a coordinate axis; every hyperplane reflection is a conjugate of a coordinate reflection by such rotations. Reflections generate the orthogonal group, so covariance holds for all of $O(4)$.

All actions preserve the inclusions. This is first true for their generators by the quotient derivative and Ward identities; the two sides of an inclusion intertwining equation then solve the same finite-dimensional linear differential equation along every exponential, with the same initial value. Uniqueness proves it on the exponentials and the group generators. Finally average the normalized stress $c^{-1}J$, with its two tensor indices moved as well, over normalized Haar measure on compact $O(4)$. Conservation, reality and each shell identity are finite-dimensional linear identities, so they pass through the integral. This produces a symmetric conserved covariant stress with the normalized translation and angular charges on the same hierarchy. The next subsection uses precisely this Euclidean action.

12.18 From Euclidean rotations to the tested boost identity

Let Ω_a be the Euclidean skew generator in the time–spatial-a plane, with vector field $x_a\partial_t - t\partial_a$. Write L_a for its action on a finite field space, and put $K_a = -iL_a$ for the corresponding Lorentz boost generator on labels. Differentiating Euclidean covariance at a configuration of distinct points gives

$$D_{(\Omega_a X_i)_i} S_v(X) = -\sum_i S_{v[i \leftarrow L_a v_i]}(X). \tag{BC1}$$

Indeed move the points by $e^{s\Omega_a}$, transfer this rotation to inverse rotation of all labels, and differentiate at zero. Each label belongs to a finite-dimensional field space, so the multilinear scalar reading

is differentiable by its finite basis expansion. Smoothness of the Euclidean hierarchy on distinct configurations justifies the positional derivative.

The extension of (BC1) to strict complex times is made after joint testing; no spatial differentiability of the strict waves is assumed. First take a compactly supported joint test φ. For fixed ordered y, introduce the complex scaling parameter z by

$$\zeta_i(z, X) = z\, t_i(X) + i y_i.$$

Let U be the set of z for which these times are strictly ordered in imaginary part for every $X \in \operatorname{supp}\varphi$. The set U is open by compactness of the support and continuity of the finitely many gaps. It is convex, since each gap inequality is affine in $\Im z$. It contains the entire real axis, in particular zero and one.

For $z = iw$, the waves in this integral are Euclidean values at the folded points

$$Q_w(X)_i = (x_i(X),\, y_i + w t_i(X)).$$

Consider the real vector field on joint configuration space

$$V_w(X) = \sum_i \left(x_{i,a} \partial_{t_i} - w(y_i + w t_i) \partial_{x_{i,a}} \right).$$

Its divergence is zero. Its image under the linear part of Q_w is exactly w times the Euclidean rotation vector field at $Q_w(X)$. This identity avoids division by w and remains valid at $w = 0$. Equation (BC1), the chain rule, and integration by parts applied to $S_v(Q_w(X))\varphi(X)$ give

$$\int \left[\begin{array}{l} F_v(\zeta(iw, X), x(X)) \displaystyle\sum_i \{ x_{i,a} \partial_{t_i} + iw\, \zeta_i(iw, X) \partial_{x_{i,a}} \} \varphi(X) \\ \quad - iw \displaystyle\sum_i F_{v[i \leftarrow K_a v_i]}(\zeta(iw, X), x(X)) \varphi(X) \end{array} \right] dX = 0. \tag{BC2}$$

The product being differentiated is smooth near the compact support: its folded points have strictly ordered heights and hence distinct positions. It extends by zero smoothly away from that support. The zero divergence makes its integrated directional derivative vanish.

Replace iw in the integrand of (BC2) by arbitrary $z \in U$. The integral is holomorphic on U. To justify this, regard the integrand as a continuous function on the compact support. Its dependence on z is holomorphic at every point, and it is jointly continuous; locally the Cauchy formula is uniform on that compact set. Integration is a bounded linear functional on the continuous-function space and therefore preserves holomorphy. Equation (BC2) makes this holomorphic integral vanish at imaginary parameters accumulating at zero. The identity theorem on the connected set U makes it vanish at $z = 1$. Thus

$$\sum_i \Psi_{v[i \leftarrow K_a v_i]}(y; \varphi) = \Psi_v(y; \mathcal{B}_{a,y}\varphi), \quad \mathcal{B}_{a,y} = \sum_i \left(x_{i,a} \partial_{t_i} + (t_i + i y_i) \partial_{x_{i,a}} \right). \tag{BC3}$$

For a general Schwartz test, use compact cutoffs converging to it in Schwartz topology. The polynomial-coefficient differential operator $\mathcal{B}_{a,y}$ is continuous on that space, and the fixed-height tested functional is continuous. Passing to the limit proves (BC3) for all joint Schwartz tests.

12.19 Strong motion of tests and finite boosts

The Poincaré action on joint tests is inverse placement,

$$(g \cdot \varphi)(X) = \varphi(g^{-1} X).$$

It is continuous in the Schwartz topology. Here is the bound that accounts for this assertion. For an invertible linear map A, translation b, and seminorm $p_{i,k}(\varphi) = \sup_X \|X\|^i \|D^k \varphi(X)\|$, the chain rule and $\|X\| \le (\|A^{-1}\| + \|b\|)(1 + \|A(X-b)\|)$ give

$$p_{i,k}(\varphi(A(\cdot - b))) \le (\|A^{-1}\| + \|b\|)^i \|A\|^k \, 2^i \max_{\substack{r \le i \\ s \le k}} p_{r,s}(\varphi). \tag{BC4}$$

The derivatives vary continuously on compact sets as A, b vary, and (BC4), with higher weighted seminorms of the fixed test, controls the complement. This proves convergence in every Schwartz seminorm.

For a boost g_s in direction a, pointwise differentiation of inverse placement gives

$$\frac{d}{ds}(g_s \cdot \varphi)\Big|_{s=0} = -\mathcal{B}_{a,0}\varphi. \tag{BC5}$$

This is also a derivative in Schwartz topology. In fact its pointwise derivative at arbitrary s is $g_s \cdot (-\mathcal{B}_{a,0}\varphi)$, a strongly continuous Schwartz family by (BC4). The fundamental theorem of calculus after evaluation at X writes the difference between the difference quotient and this derivative as

$$\int_0^1 \big[D(us) - D(0)\big](X)\, du.$$

The same equality holds after each spatial derivative. Taking each weighted supremum bounds its seminorm by $\sup_{0 \le u \le 1} p_{i,k}(D(us) - D(0))$, which tends to zero by continuity on the compact parameter interval. This proves (BC5) strongly without integrating in the Schwartz space itself.

Pass (BC3) to the boundary. Its extra height terms are the finite sum

$$i \sum_i y_i \, \Psi_v(y; \partial_{x_{i,a}} \varphi);$$

each tends to zero, because the tested boundary limit exists for that fixed derivative test and $y_i \to 0$. Therefore

$$\sum_i W_{v[i \leftarrow K_a v_i]}(\varphi) = W_v(\mathcal{B}_{a,0}\varphi). \tag{BC6}$$

Choose one finite field space containing all labels, using the directed inclusions. Consider

$$H(s) = W_{(\rho(g_s)v_i)_i}(g_s \cdot \varphi).$$

Differentiate at zero. The label derivative is the sum on the left of (BC6), and the test derivative is its negative by (BC5). Their cancellation gives $H'(0) = 0$. This product rule needs no topology on the entire field union: expand the multilinear distribution in a basis of the one finite field space, differentiate the finitely many scalar coordinate products and the pairings with fixed basis labels, and sum. At an arbitrary parameter value repeat the same argument with the already moved labels and test. The group law gives $H'(s) = 0$ everywhere, hence

$$W_{(\rho(g_s)v_i)_i}(g_s \cdot \varphi) = W_v(\varphi). \tag{BC7}$$

Real translations and spatial rotations follow first at strict heights from their Euclidean identities by one-complex-parameter uniqueness, and then at the boundary by change of variables, whose Jacobian has modulus one. The transformations satisfying the boundary covariance identity form a subgroup: composition follows from the representation law and test pushforward law, and inverses follow by applying the identity to inversely transformed labels and tests. Spatial rotations and coordinate boosts generate the specified Lorentz group, so (BC7) gives full covariance once that generation statement is established.

12.20 Dual cone and a smooth frequency cutoff

The absolute-frequency cone has a useful elementary dual description. Write $p = (p_1, \dots, p_n)$, with total momentum zero, and $P_j = \sum_{i=j}^{n} p_i$. Its defining remaining condition is that each proper suffix momentum belongs to the closed forward cone. If y has all successive differences in the corresponding closed dual forward cone, summation by parts gives

$$\langle p, y\rangle = \sum_{j=2}^{n} \langle P_j, y_j - y_{j-1}\rangle \geq 0. \tag{DC1}$$

The endpoint terms vanish because the total momentum and the empty suffix are zero. Self-duality of the forward cone proves the inequality term by term. The identity also covers the empty tuple.

There is a quantitative interior form. If y is in the open tube-height cone, choose $r > 0$ such that $B(y, r)$ is contained in that cone. For $\|u - y\| < r/2$ and nonzero p in the frequency cone, the point

$$u - \frac{r}{2\|p\|} p$$

lies in $B(y, r)$. Apply (DC1) there to obtain

$$\langle p, u\rangle \geq (r/2)\|p\|. \tag{DC2}$$

For $p = 0$ the estimate is immediate. This gives one positive exponential damping margin on a whole neighborhood of any interior height.

A smooth cutoff near the frequency cone is constructed explicitly. Convolve the indicator of its open distance-two neighborhood with a nonnegative, unit-integral smooth bump supported in the ball of radius $1/2$. Denote the result by η. It equals one on the distance-one neighborhood of the cone: every point sampled by the convolution then has distance less than $3/2$ from the cone. Its support is contained in the distance-$5/2$ neighborhood, and its closed support lies in the distance-three neighborhood. All derivatives are globally bounded, since derivatives can be put on the compactly supported bump and the indicator has absolute value at most one:

$$\|D^k \eta\|_\infty \leq \|D^k b\|_{L^1}.$$

The smoothing and support assertions use the actual Euclidean metric on the full absolute-frequency space, including its dimension-zero case.

For a closed tube height y and a frequency x in this closed support, choose a cone point z within distance three of x. Equation (DC1) and Cauchy–Schwarz imply

$$-\langle y, x\rangle \leq \langle y, z - x\rangle \leq 3\|y\|.$$

Hence on $\|y\| \leq R$,

$$\|D_x^\ell e^{-\langle y,x\rangle}\| \leq e^{3R} R^\ell.$$

Leibniz's formula and the cutoff derivative bounds show that $\eta(x) e^{-\langle y,x\rangle}$ has uniformly controlled derivatives for bounded closed heights. More generally the polynomial-growth cutoff estimates used in the source give

$$\|D_x^k(\eta e^{-\langle y,\cdot\rangle})(x)\| \leq C_{k,R}(1 + \|x\|)^{N_k}. \tag{DC3}$$

Outside the closed support all its derivatives vanish. Thus these multipliers act continuously on Schwartz tests, with bounds uniform over bounded closed heights; interior heights additionally have the exponential margin (DC2). These are the cutoff and damping estimates needed for the subsequent tube construction.

12.21 Holomorphic tube functions from the frequency distribution

The tube is an open complex-time domain with strictly ordered imaginary heights. Positive heights provide the damping needed for the frequency distributions and Hilbert vectors. Removing them is a tested tempered boundary limit; it is not substitution of a pointwise value at zero height. The Gram argument below shows that independently heated word vectors give one positive boundary form on the common word space.

Here the radian Fourier transforms on $\mathbb{R}^N$, $N = 4n$, are explicitly

$$\mathscr{H}g(q) = \int e^{-i\langle x,q\rangle} g(x)\,dx, \qquad \mathscr{P}f(q) = (2\pi)^{-N} \int e^{i\langle q,x\rangle} f(x)\,dx.$$

They are inverse continuous maps of the Schwartz space. For a closed tube height y, set

$$D_y g(q) = \eta(q) e^{-\langle y,q\rangle} g(q).$$

The estimates (DC3) and Leibniz's rule bound each seminorm of $D_y g$, uniformly for y in bounded subsets of the closed height cone, by finitely many seminorms of g. Joint smoothness of the multiplier gives convergence of every jet on compact frequency sets as y varies. The same bounds with one extra polynomial weight control the complements. It follows that $y \mapsto D_y g$ is continuous in the Schwartz topology even at closed-cone boundary heights.

For a full complex configuration $Z = X + iY$ in the open tube, define the lens

$$\lambda_Z(q) = \eta(q) e^{i\langle Z,q\rangle} = \eta(q) e^{i\langle X,q\rangle - \langle Y,q\rangle}.$$

It is a genuine Schwartz test. More precisely, (DC2), together with the distance-three support bound, gives constants $c > 0, A$ such that

$$-\langle Y, q\rangle \le A - c\|q\| \quad (q \in \operatorname{supp}\eta); \qquad A = 3(\|Y\| + c) \text{ is sufficient.}$$

Each frequency derivative of the exponential contributes at most a power of its linear functional norm. Derivatives of η have polynomial growth. Leibniz's formula and

$$(1+t)^k e^{A-at} \le k! a^{-k} e^{A+a} \qquad (t \ge 0,\ a > 0)$$

therefore prove, for every k, j, a finite bound

$$\sup_q (1 + \|q\|)^k e^{(c/2)\|q\|} \|D_q^j \lambda_Z(q)\| < \infty. \tag{TL1}$$

The exponential margin retained in (TL1) justifies complex differentiation of the lens as a Schwartz-valued function.

For completeness, if L is a small real-linear, complex-valued frequency functional, with $\|L\| \le \min(a, 1)$, the scalar exponential remainder and its derivatives satisfy

$$\|D^j(e^{L(q)} - 1 - L(q))\| \le \|L\|^2 (1 + \|q\|)^2 e^{a\|q\|}. \tag{TL2}$$

For $j = 0$ this is the exponential-series remainder. For $j = 1$ differentiate once and bound $e^L - 1$. For $j \ge 2$, derivatives of the affine terms vanish and the bound follows from $\|L\|^j e^{\Re L(q)}$, using $\|L\| \le 1$. Multiply (TL2) by the lens and use (TL1) and Leibniz's rule. Every Schwartz seminorm of $\lambda_{Z+H} - \lambda_Z - i\langle H, \cdot\rangle \lambda_Z$ is at most $C_{k,j}\|H\|^2$. Dividing by $\|H\|$ proves the complex Fréchet derivative

$$d\lambda_Z(H)(q) = i\langle H, q\rangle \lambda_Z(q) \tag{TL3}$$

in the actual Schwartz topology, jointly in all complex coordinates.

For the tempered frequency functional $\widehat{W}_v$, define

$$\mathcal{F}_v(Z) = (2\pi)^{-N}\widehat{W}_v(\lambda_Z). \tag{TL4}$$

Continuity and linearity of $\widehat{W}_v$ transfer (TL3) to complex differentiability of $\mathcal{F}_v$. Finite-dimensional complex analysis makes it holomorphic on the full open tube. At this point no spectral support assertion is needed for holomorphy itself; the support condition will make the result independent of the cutoff.

The identification with a damped distribution requires an integral argument in Schwartz topology. At a fixed interior Y, put $g(q) = \eta(q)e^{-\langle Y,q\rangle}$, a Schwartz test by (TL1), and $g_X(q) = e^{i\langle q,X\rangle}g(q)$. Its Fourier transform is the translate $(\mathscr{H}g)(\,\cdot - X)$. For any Schwartz function h,

$$p_{i,j}\big(h(\,\cdot - X)\big) \le 2^i\left[p_{i,j}(h) + \|X\|^i p_{0,j}(h)\right]. \tag{TL5}$$

This follows from $\|q\| \le \|q-X\| + \|X\|$, splitting according to which summand is larger. Continuity of a tempered distribution bounds it by finitely many Schwartz seminorms. Thus $X \mapsto \widehat{W}_v(g_X)$ is continuous and has at most polynomial growth. Its product with any joint Schwartz test $\varphi(X)$ is absolutely integrable.

For compactly supported φ, integrate $g_X\varphi(X)$ over its compact support. Uniform control of every jet permits the continuous functional $\widehat{W}_v$ to pass through this integral. Fourier inversion gives

$$\widehat{W}_v(D_Y(\mathscr{P}\varphi)) = \int \mathcal{F}_v(X+iY)\varphi(X)\,dX. \tag{TL6}$$

For arbitrary Schwartz φ, use its smooth compact cutoffs. The left side converges by continuity on Schwartz space; the right side converges by (TL5), polynomial growth, and dominated convergence. This proves (TL6) for full joint tests, with an absolutely convergent right side. It identifies the holomorphic emission with the actual damped boundary distribution at every fixed interior height, without a fictitious norm on Schwartz space or an exchange of unbounded integrals.

12.22 Identification of the tube continuation and its covariance

Spectral support now removes the auxiliary cutoff. A distribution supported in the closed frequency cone annihilates a Schwartz test that vanishes on a neighborhood of that cone. For compactly supported g, the two tests

$$\eta(q)e^{-\langle Y,q\rangle}g(q), \qquad e^{-\langle Y,q\rangle}g(q)$$

agree on the distance-one neighborhood of the cone. Their difference has closed support disjoint from it; hence the supported distribution gives them the same value. This argument uses the compact support of g when defining the second test. It does not assert that an unrestricted exponential is a Schwartz multiplier.

At a fixed strictly ordered temporal height y, the strict restriction $f_{v,y}$ is bounded and continuous by (SW7). Its integral against a joint test therefore defines a tempered distribution u_y. In (FT2) use the compact test $e^{-\langle e_y,q\rangle}g(q)$. Its exponential lift cancels exactly. The preceding cutoff removal then gives

$$\widehat{W}_v(D_{e_y}g) = \widehat{u}_y(g)$$

for all compactly supported smooth g. Both sides are continuous on Schwartz space, so smooth compact cutoffs extend equality to every Schwartz test. Fourier inversion and (TL6) show

$$\int \mathcal{F}_v(X+ie_y)\varphi(X)\,dX = \int f_{v,y}(X)\varphi(X)\,dX.$$

Both functions are continuous. Equality against all compact smooth tests gives almost-everywhere equality, and full support of Lebesgue measure upgrades it to pointwise equality:

$$\mathcal{F}_v(X + ie_y) = f_{v,y}(X). \tag{TI1}$$

Thus the constructed full-tube function continues the very strict waves previously constructed.

Let r be a real Lorentz transformation in the specified group, and $r^\vee$ its dual for the Euclidean frequency pairing, so that $\langle rX, r^\vee q\rangle = \langle X, q\rangle$. The tube-height cone is invariant under r, and the frequency cone is invariant under $r^\vee$. Compare the two Schwartz lenses

$$\lambda_Z(q), \qquad \lambda_{rZ}(r^\vee q).$$

Their exponential factors agree everywhere. Their cutoff factors need not agree globally, but both equal one near every cone point p: use the intersection of the unit ball about p and the inverse image of the unit ball about $r^\vee p$. Consequently their difference has closed support disjoint from the cone and is annihilated by $\widehat{W}_v$. Fourier transformation of boundary covariance, followed by this lens comparison, gives

$$\mathcal{F}_{(\rho(r)v_i)_i}(rZ) = \mathcal{F}_v(Z). \tag{TI2}$$

This proves real Lorentz covariance inside the tube without requiring an invariant cutoff. The same local support argument compares any two permitted cutoffs and proves that their emissions agree.

12.23 Analytic germs at an adjacent spacelike pair

Fix adjacent argument slots $i, j = i + 1$, and write σ for their transposition. Choose height profiles $y^{(m)} \to 0$ with exactly this adjacent gap flattened: $y_i^{(m)} = y_j^{(m)}$, and all other ordered gaps strictly positive. One explicit choice is $y_k^{(m)} = a_k/(m+1)$, where $a_k = k$ except $a_j = i$. At two arguments this is the zero profile. Both the original and transposed imaginary configurations lie in the closed tube-height cone, and their complex configurations are in the closures of the respective tubes.

At a real configuration with spacelike difference $v = X_j - X_i$, set $T = \|\mathbf{v}\| > |v^0|$, and choose the spatial unit vector $e = \mathbf{v}/T$. The infinitesimal boost

$$C_e v = (v^0 e,\ e \cdot \mathbf{v})$$

is strictly future timelike, since its spatial norm is $|v^0|$ and its time component is T. Let S_e project onto the spatial-e and time plane. The complex boost is

$$A_e(z)Z = Z + \sinh z\, C_e Z + (\cosh z - 1) S_e Z.$$

For $z = ib$, its imaginary difference at any adjacent pair is

$$Y + \sin b\, V + (\cos b - 1) W, \tag{LG1}$$

where Y is the original imaginary difference, $V = C_e$ applied to the real difference, and $W = S_e Y$.

For every unflattened adjacent gap, Y is already strictly timelike, so (LG1) remains so in a neighborhood of $b = 0$ and of the given configuration. At the flattened pair, V is strictly timelike. Use the exact identity

$$\sin b\, V + (\cos b - 1) W = \sin b \left(V - \frac{\sin b}{1 + \cos b} W \right).$$

The vector in parentheses is strictly timelike throughout a neighborhood of $(b, Z) = (0, Z_0)$, and $\sin b > 0$ for sufficiently small $b > 0$. Adding it with that positive coefficient to a closed forward-cone vector Y enters the open cone. Since there are finitely many adjacent gaps, one product neighborhood and one small positive angle work for all of them, uniformly down to angle zero. Apply the same calculation to the transposed configuration with the opposite imaginary boost; at the flattened pair the two reversals of sign cancel. Thus the original tube and the transposed tube each have a small imaginary-boost path into their interiors, with this path property holding for nearby points of the corresponding tube, not just at the boundary point itself.

The labels also have an analytic boost action. The following elementary Banach-algebra argument supplies it. For a continuous one-parameter group $u(t)$, put $V(t) = \int_0^t u(s)\,ds$. Then $V'(t) = u(t)$ and $V(t)/t \to 1$. Choose nonzero a so that $g = V(a)/a$ is invertible. The group law gives

$$u(t)V(a) = V(t + a) - V(t), \qquad u(t) = a^{-1}[V(t + a) - V(t)]g^{-1}.$$

The right side proves differentiability. If $B = u'(0)$, differentiating the group law gives $u'(t) = u(t)B$, whence $\frac{d}{dt}[u(t)e^{-tB}] = 0$. Evaluation at zero yields $u(t) = e^{tB}$. For each finite-dimensional field space this gives the entire complex extension $u(z) = e^{zB}$ of its boost representation.

Fix the finite field space containing the finitely many labels. For a boundary configuration Z_0, take the small imaginary-boost endpoint it just constructed and define near Z_0

$$G(Z) = \mathcal{F}_{(u(it)v_k)_k}(A_e(it)Z). \tag{LG2}$$

The endpoint lies in the tube, so (LG2) is holomorphic on a full neighborhood of Z_0. It agrees with the original emission on nearby points already in the original tube. To prove this agreement, for such a point Z consider

$$Q_Z(z) = \mathcal{F}_{(u(z)v_k)_k}(A_e(z)Z).$$

This is holomorphic in a neighborhood of the entire imaginary-boost path. For moving labels, expand in the finite field-space basis; each coefficient product is holomorphic and each fixed-label emission is holomorphic. At real z, (TI2) makes $Q_Z(z)$ equal to the original emission at Z. The one-variable identity theorem propagates this constant value along the connected path from zero to it. This proves the claimed agreement and hence the analytic germ.

For the transposed tube repeat the argument with the opposite path and with labels permuted by σ. The permutation map is complex linear and commutes with real boosts, so the same argument applies. We have thus constructed both holomorphic germs at every spacelike adjacent pair, with the common spectator-height sequence tending to zero. The next argument identifies the two germs and passes their equality to the boundary exchange relation.

12.24 Identifying the germs and realizing their partial boundaries

First consider a real configuration whose time coordinates are zero, with the adjacent pair spacelike, and the spectator profile y chosen above. The associated Euclidean configuration $q = ((\mathbf{x}_k, y_k))_k$ has no collisions: every pair other than the flattened pair has different heights, and the flattened pair has different spatial positions because it is spacelike.

Use the totally real Wick slice

$$M(q) = (\mathbf{q}_k,\, iq_k^0)_k.$$

Its image spans the full complex configuration space over $\mathbb{C}$. Indeed each complex configuration is $M(u)+iM(v)$, by taking spatial real and imaginary parts and, in the time coordinates, respectively the imaginary part and the negative real part.

Let G_+ and G_- be the two holomorphic germs previously constructed at $M(q)$. On the open Euclidean side with strictly ordered heights, $G_+ \circ M$ equals the original smooth Euclidean hierarchy S_v. On the different open side where the transposed heights are strictly ordered, $G_- \circ M$ equals the same function: the strict-wave identification gives the permuted Euclidean hierarchy there, and permutation symmetry identifies it with S_v. Both open sides accumulate at q, since the tied heights can be separated in either direction.

For each order m, differentiate on the relevant open side and approach q. Continuity of the m-th derivatives of the analytic pullbacks and of the smooth hierarchy proves

$$D^m(G_+ \circ M)(q) = D^m S_v(q) = D^m(G_- \circ M)(q).$$

The chain rule identifies these real derivatives with restrictions of the complex m-linear derivatives of $G_\pm$ to arguments in the image of M. Complex multilinearity and the spanning property extend equality to all complex arguments. Thus all complex derivatives agree. The convergent analytic Taylor series give

$$G_+ = G_- \quad \text{on a neighborhood of } M(q). \tag{GE1}$$

Notice that equality was not assumed on a common open part of the two Euclidean sides; equality of their limiting jets is what proves (GE1).

We also need to know what the germs represent as distributions. The following local argument supplies that fact. Suppose E is an interior tube function, G is a holomorphic germ through $X+iY$, and the interior slices have a weak distributional limit Q as their heights approach Y. Choose a small ball of real configurations about X and a small height neighborhood such that G is defined and bounded in their product, and agrees with E at the interior points there. For a joint compact smooth test φ supported in that ball, dominated convergence with the majorant $C|\varphi|$ gives

$$Q(\varphi) = \int G(X' + iY)\varphi(X')\, dX'. \tag{GE2}$$

The height approach is nonempty because the base point belongs to the closure of the tube. Holomorphy gives continuity of G, and the already established damped-distribution boundary theorem supplies the weak limit. Thus (GE2) identifies that particular partial boundary; it does not merely define a new distribution.

Apply this argument separately to the original and transposed tubes. In the second application reindex all real integration variables by the adjacent transposition. This map preserves Lebesgue measure, is its own inverse, and sends the test to its corresponding pullback; it also permutes the imaginary heights. Therefore the transposed partial boundary is represented locally by the second germ in precisely the same original variables. This proves the local representation of both partial distributions by their respective germs. Combined with (GE1), it proves their local equality at the zero-real-time spacelike configurations. The extension from these configurations to general spacelike pairs and the final zero-height limit are the next steps.

12.25 Independently heated word vectors and their Gram limit

The positivity argument uses actual vectors in the reflection-positive chamber. For a positive ordered height profile y, choose a width p smaller than its initial height and its gaps. The buffered shot vector is

$$\xi_{p,y,v}(X) = R(z_1, p, x_1)B_p(v_1)R(z_2 - z_1, 2p, x_2 - x_1)B_p(v_2) \cdots R(z_n - z_{n-1}, 2p, x_n - x_{n-1})B_p(v_n)\Omega,$$

where $z_i = t_i(X) + iy_i$. At arity zero it is Ω. Recursive contraction estimates give

$$\|\xi_{p,y,v}(X)\| \le \|\Omega\| \prod_i \|B_p(v_i)\|. \tag{HV1}$$

It is strongly continuous in X: the heat factors vary in operator norm while their parameters remain in the open half-plane; the spatial translations are strongly continuous; induction handles their application to the recursively varying ket.

If the real times vanish, the shot is exactly the planted future word at the Euclidean positions (x_i, y_i). This follows by the insertion and translation recursion used for (SW4), now with the initial reference width zero. Its inner product with a second shot, of independent width q and positive profile z, is the Euclidean hierarchy of the first word reflected and reversed, followed by the second. The first block of heights is $(-y_n, \dots, -y_1)$, the second is z; the first labels are $Jv_n, \dots, Jv_1$, with J the induced adjoint on union labels.

The same identity holds for arbitrary real times. Multiply the first array of real times by $\bar{c}$, and the second by c, in a complex parameter c. Antilinearity of the first Hilbert slot makes the scalar inner product holomorphic in c. There is an open horizontal strip containing both zero and one where all initial-height and gap inequalities for these fixed widths hold. The hierarchy side is holomorphic on that strip as well. At purely imaginary c both shots are Euclidean planted words, with the corresponding shifted heights, so the preceding Gram identity holds there. The identity theorem gives it at $c = 1$. Thus

$$\langle \xi_{p,y,v}(X), \xi_{q,z,w}(Y)\rangle = f_{(Jv_n,\dots,Jv_1,w_1,\dots,w_m),\,(-y_n,\dots,-y_1,z_1,\dots,z_m)}(X_n, \dots, X_1, Y_1, \dots, Y_m). \tag{HV2}$$

In particular no equality of the two widths was needed. The real coordinates in the reversed first block are reversed without a further geometric reflection; the imaginary heights have already changed sign.

For a smeared word $b = ((v_i, \varphi_i))_{i=1}^n$, let $\varphi_b(X) = \prod_i \varphi_i(X_i)$, take $y_i(t) = it$ with slots numbered from one, and take $p(t) = \min(t, 1)/8$. Define

$$\eta_b(t) = \int \varphi_b(X) \xi_{p(t),y(t),v}(X)\, dX, \qquad t > 0.$$

This is an absolutely integrable Hilbert-valued integral by (HV1) and Schwartz integrability. Translating every real argument by a common a translates the shot by the chamber's real-time translation unitary. For the empty word this uses invariance of Ω under both the spatial group and the real-time edge group: invariance under heat first extends to the holomorphic half-plane by analytic uniqueness, and then to its real-time edge. For nonempty words the initial propagator changes by that edge translation, whereas all relative propagators are unchanged. Change of variables in the integral therefore gives

$$\eta_{a\cdot b}(t) = U(a)\eta_b(t). \tag{HV3}$$

Apply (HV2) under the two integrals. Their absolute integrability follows from the same fixed-height bounds. The map $(X, Y) \mapsto (X_n, \dots, X_1, Y_1, \dots, Y_m)$ is a real linear isometry and preserves Lebesgue measure. The product test becomes $\overline{\varphi_b(X)}\varphi_c(Y)$, exactly the test of the reversed adjoint word followed by c. Hence the all-height boundary theorem gives the independent two-parameter limit

$$\lim_{\substack{t\downarrow 0\\ u\downarrow 0}} \langle \eta_b(t), \eta_c(u)\rangle = W(b^* c). \tag{HV4}$$

The joined profile is strictly ordered for all $t, u > 0$ and tends to zero as the pair tends to zero, which justifies using that boundary theorem. The independence of the parameters is stronger than a diagonal limit and is what will make the shot vectors Cauchy.

12.26 Density of product tests

Finite sums of product Schwartz tests are dense in the joint Schwartz space. Here is the concrete approximation used. First take a compactly supported joint test φ, and choose a product cutoff $\chi(X) = \prod_i \chi_i(X_i)$ equal to one on its support. Approximate its Fourier transform by compactly supported Fourier cutoffs. Inverting and multiplying by χ gives Schwartz convergence to φ. Each approximant is an integral over a compact frequency set of the tests

$$\chi(X)e^{i\langle q,X\rangle} = \prod_i \left[\chi_i(X_i)e^{i\langle q_i,X_i\rangle}\right].$$

This is a continuous family in Schwartz topology: all derivatives vary jointly on a common compact spatial support, giving uniform control of every seminorm. Quadrature on the compact frequency set approximates its integral in each seminorm by finite sums of these product tests. Thus each approximant is in the closure of the product-test span, and so is φ. Finally smooth compact spatial cutoffs converge to any Schwartz test, proving density in general. At arity zero the product test is the constant one on the one-point configuration space and already spans that space. This density permits subsequent joint-test conclusions to be recovered from word-vector identities.

12.27 Spacelike exchange on joint tests and on words

Two germs anchored to the same tube at a point of its closure agree. To prove this, take a ball on which both extensions are holomorphic and agree with the tube function in that ball's intersection with the tube. This intersection contains a point because the base point belongs to the tube closure; it is open because the tube is open. Analytic uniqueness on the connected ball makes the extensions equal there.

This also proves propagation along a continuous path of base points in the closures of two tubes. At each base point take the two germ extensions on a ball. At nearby base points their germs agree with the germs selected there, by the preceding uniqueness statement. Whether the two extensions agree is therefore locally constant along the path: agreement propagates across the ball by analytic uniqueness, in either direction. On a connected parameter interval it is constant.

Given any real configuration X with the designated pair spacelike, contract all its real time coordinates to zero:

$$X_k(s) = (\mathbf{x}_k, st_k), \qquad 0 \le s \le 1.$$

The pair remains spacelike, since $s|t_j - t_i| \le |t_j - t_i| < \|\mathbf{x}_j - \mathbf{x}_i\|$. With the fixed spectator heights, the path $X(s) + ie_y$ belongs to both tube closures. Both germs exist throughout it. Their equality at $s = 0$, proved in (GE1), therefore propagates to $s = 1$. Thus the two local extensions agree near every spacelike configuration.

Equation (GE2) now identifies the two partial distributions locally on a sufficiently small real ball. Their difference vanishes there on every compact smooth test. Consequently its distributional support misses every point of the spacelike-pair open set, and it vanishes on all Schwartz tests whose closed support is contained in that set. Apply this to every profile $y^{(m)}$. The damped distribution at $e_{y^{(m)}}$ converges, on each fixed test, to the undamped boundary distribution: the closed-height multiplier is continuous in Schwartz topology, and its zero-height cutoff is one near the frequency support. The same is true after the adjacent transposition. Since the two partial distributions agree for every m, uniqueness of their scalar limits gives

$$W_v(\varphi) = W_{v\circ\sigma}(\varphi \circ \sigma) \quad \text{if } \operatorname{supp}\varphi \subset \{X : X_j - X_i \text{ is spacelike}\}. \tag{LC1}$$

This is the joint-test exchange identity with all spectator variables present.

For product tests $\varphi = \prod_k \varphi_k$, its closed support is contained in the inverse image of $\operatorname{supp}\varphi_k$ under every coordinate projection. If the supports of φ_i, φ_j are mutually spacelike, that joint support lies in the set in (LC1). Permuting the joint product permutes exactly those two test factors. Place these factors between any two spectator word lists to obtain

$$W\big(a\,(v,\varphi_i)(w,\varphi_j)\,b\big) = W\big(a\,(w,\varphi_j)(v,\varphi_i)\,b\big). \tag{LC2}$$

No support restriction is placed on the spectator tests. This is the word relation that will give locality on the common field domain.

12.28 The label quotient remains faithful at the boundary

Suppose a union label x is a boundary marker: every word expectation containing its smearing, with arbitrary words before and after, is zero. Then every joint distribution with that label in any prescribed slot is zero. It vanishes first on product tests, since a word can be split just before that slot, and then on all joint tests by product-test density and continuity.

Represent x in a finite quotient field space by a raw label w. The full-tube emission for any distribution containing x is zero, so (TI1) makes its strictly time-ordered Euclidean reading zero. This remains true after any permutation of the arguments, because the marker may occupy any slot. For a general Euclidean configuration of distinct points, sort its time coordinates. They are weakly ordered, possibly with ties. Approach the sorted heights by strictly ordered heights while keeping spatial coordinates fixed. The configurations remain distinct near the given one. Smoothness of the hierarchy there makes their readings converge to the given reading, which is therefore zero. Undoing the sorting permutation proves vanishing at every distinct configuration.

The Euclidean null-space criterion now says that w belongs to its finite-level null subspace. The selected hierarchy has exactly the retained null subspaces of the original hierarchy, so its class in the original quotient is zero. Hence $x = 0$ in the common field union. This proves faithfulness of the boundary label quotient; it is not an additional field assumption.

12.29 Boundary positivity and the common field domain

The common domain is the complex span of finite smeared field words applied to the vacuum. Field multiplication, adjoints and spacetime covariance are asserted on that domain. Its density follows from the cyclic construction, and invariance follows from adjoining another smeared field or moving every test by the spacetime action. No extension to a common bounded-operator algebra is claimed.

For each word b, (HV4) with $c = b$ makes $\eta_b(t)$ Cauchy as $t \downarrow 0$. Indeed expand

$$\|\eta_b(t) - \eta_b(u)\|^2 = \langle \eta_b(t), \eta_b(t)\rangle - \langle \eta_b(t), \eta_b(u)\rangle - \langle \eta_b(u), \eta_b(t)\rangle + \langle \eta_b(u), \eta_b(u)\rangle.$$

Every term has the same limit by the independent two-parameter convergence, so the expression tends to zero. Completeness gives a vector η_b. Taking independent limits in (HV4) yields

$$\langle \eta_b, \eta_c\rangle = W(b^*c), \qquad \eta_\varnothing = \Omega, \qquad \eta_{a\cdot b} = U(a)\eta_b. \tag{FR1}$$

Consequently the kernel $K(b,c) = W(b^*c)$ is Hermitian and positive: for any finite coefficients,

$$\sum_{\alpha,\beta} \overline{z_\alpha} z_\beta K(b_\alpha, b_\beta) = \Big\| \sum_\alpha z_\alpha \eta_{b_\alpha} \Big\|^2 \geq 0.$$

The empty expectation is one because the chamber vacuum has norm one.

Form the vector space of finite formal sums of smeared words, with sesquilinear form K, linear in its second argument. Divide by its zero-norm subspace and complete. Denote the resulting Hilbert space by $\mathcal{H}_W$, the word classes by $[b]$, and their algebraic span by $\mathcal{D}$. This span is dense, and $\Omega_W = [\varnothing]$ is a unit vector. The common operator domain will be exactly $\mathcal{D}$.

For a letter $\ell = (x, \varphi)$, left insertion on formal sums is $L_\ell[b] = [\ell b]$. The word involution gives

$$\langle L_\ell p, q\rangle = \langle p, L_{\ell^*} q\rangle, \qquad \ell^* = (Jx, \overline{\varphi}). \tag{FR2}$$

This identity also proves insertion preserves the null space. If $[p] = 0$, choose a formal sum q representing the vector $L_\ell p$ in $\mathcal{D}$; then (FR2) makes $\|L_\ell p\|^2 = \langle [p], L_{\ell^*}[q]\rangle = 0$. Thus

$$\Phi(x, \varphi)[b] = [(x, \varphi)b] \tag{FR3}$$

is a well-defined linear operator $\mathcal{D} \to \mathcal{D}$. It is bilinear in x, φ. To prove this directly, pair the proposed additive or scalar relation with every word vector. The corresponding word expectation is linear in the inserted label and test, since the joint distribution is multilinear in its labels and the product test is multilinear in its factors. All pairings vanish, and density proves the relation. There is no boundedness claim for the fields themselves.

Equation (FR2) is the adjoint relation on the common domain. Repeated application of (FR3) gives the word vectors, and their vacuum expectations are the prescribed W's. Cyclicity follows because $\mathcal{D}$ is exactly their span. The smearing map $\varphi \mapsto \Phi(x, \varphi)\psi$ is norm-continuous for each $\psi \in \mathcal{D}$. For a word vector, continuity of the diagonal and mixed joint-distribution pairings implies continuity of the norm difference by its squared-norm expansion; finite linear combinations give the assertion for ψ. The jointly continuous two-smearing Gram pairing used here follows from continuity of product tests in Schwartz topology and continuity of each fixed-label joint distribution.

The word exchange identity (LC2), paired against all word probes, gives

$$[\Phi(x, \varphi), \Phi(y, \psi)]\big|_{\mathcal{D}} = 0$$

when the two test supports are mutually spacelike. Finally if $\Phi(x, \varphi) = 0$ on $\mathcal{D}$ for every φ, all expectations with x in any slot vanish: pair its zero insertion against the reversed adjoint left spectator word. Involutivity of J removes that double reversal. The boundary marker criterion already proved gives $x = 0$. Thus the field representation retains the full prescribed label quotient.

12.30 Unitary spacetime action on the reconstructed space

For a Poincaré transformation g, transform each letter by its finite-dimensional label action and inverse placement of its test. Boundary covariance preserves $K(b, c)$; the adjoint on letters commutes with this transformation. Thus

$$\mathcal{U}(g)[b] = [g \cdot b]$$

extends isometrically to $\mathcal{H}_W$. Its inverse is $\mathcal{U}(g^{-1})$, and the word transformation law gives the representation law on the dense span and hence everywhere.

Its strong continuity can be proved without a topology on the full label union. Pin the finitely many labels in one finite field space and expand the multilinear distribution in a fixed basis there. Each moved expectation is a finite sum of products of continuous label coordinates and continuous fixed-label distribution pairings with moved tests. Strong continuity of test motion was proved in (BC4). Hence all mixed word-vector pairings vary continuously. The squared-norm identity gives norm continuity of $\mathcal{U}(g)[b]$, and then of finite sums. To extend to any vector, approximate by such a

sum within $\epsilon/3$; isometry controls the two approximation errors uniformly in g. This proves strong continuity on the Hilbert space.

The domain $\mathcal{D}$ is invariant, the empty word fixes the vacuum, and insertion transforms by

$$\mathcal{U}(g)\Phi(x,\varphi)\mathcal{U}(g)^{-1} = \Phi(\rho(g)x, g\cdot\varphi) \quad \text{on } \mathcal{D}. \tag{FR4}$$

The isometric attachment to the original reflection-positive chamber is also concrete: map $[b]$ to η_b in (FR1), extend by continuity, and use equality of all Gram pairings to see that it is an isometry. It sends the vacuum to Ω and intertwines translations by (FR1). The spectral and vacuum conclusions can therefore be transferred through this attachment by the arguments below.

12.31 Coordinate boosts, the spectral cone, and attachment

We supply the group-generation fact used above. For a strictly future timelike vector x, its time component T is positive and every spatial component x_a satisfies $|x_a/T| < 1$. The coordinate boost of rapidity

$$s = -\operatorname{artanh}(x_a/T)$$

sets that component to zero, since $T\sinh s + x_a\cosh s = 0$. It preserves all other spatial components and keeps the vector in the strict future cone. Three successive coordinate boosts therefore bring any future unit timelike vector to the unit time vector e_0. For any proper orthochronous Lorentz transformation g, apply these boosts to ge_0, obtaining $kge_0 = e_0$.

A Lorentz transformation fixing e_0 preserves the time coordinate, by invariance of the Lorentz pairing. Since it also preserves the Minkowski quadratic form, it preserves the Euclidean norm. It is therefore an orthogonal spatial turn, with the original determinant. Thus kg is such a spatial turn and $g = k^{-1}(kg)$ is a finite product of coordinate boosts and spatial turns. This proves the generation statement used to obtain full boundary covariance.

For precision, the absolute-frequency cone is initially defined as the intersection, over all Lorentz transformations r, of the inverse images under $r^\vee$ of the time cone in (FT3). Boundary covariance and the Fourier change-of-variables formula show that the Fourier distribution is supported in every one of these inverse images: pull a test disjoint from one inverse image through $r^\vee$, use the time-cone support for the transformed labels, and transform back. Hence it is supported in their intersection.

This intersection has the suffix-momentum description used in (DC1). For any suffix, test the time cone against the weakly increasing step height that is zero before that suffix and one on it. Its time component is nonnegative. Doing this after every Lorentz transformation makes the suffix vector future causal: a vector with nonnegative transformed time component under every proper orthochronous Lorentz transformation is in the closed forward cone. To see the converse characterization of a single vector, if its time is negative the identity already excludes it; if its spatial norm exceeds its nonnegative time, a boost toward its spatial direction makes its time negative. Constant negative as well as positive heights force total time momentum zero. The total suffix is future causal and has zero time, so its spatial part is zero too. Conversely, zero total momentum and future-causal proper suffixes imply all the time-cone inequalities by telescoping and imply them after every Lorentz transformation. This proves the description, with no change to relative coordinates.

For the field Hilbert space let P be the joint projection-valued measure of its strongly continuous translation representation, with convention

$$\langle v, \mathcal{U}(a)v\rangle = \int e^{i\langle p,a\rangle}\, d\mu_v(p), \qquad \mu_v(B) = \langle v, P(B)v\rangle.$$

The translation attachment in (FR1) is an isometry I into the reflection-positive chamber and intertwines translations. Equality of these scalar characteristic functions and uniqueness of Fourier transforms of finite measures give equality of the diagonal measures μ_v with the chamber's diagonal translation measures at Iv. The latter have nonnegative time support, as established by the semigroup argument below.

Lorentz covariance of the unitary representation gives

$$\mu_{\mathcal{U}(r)v} = (r^{\vee})_* \mu_v. \tag{SP1}$$

This follows by conjugating translations and comparing characteristic functions, using preservation of Hilbert inner products. If p is in the support of μ_v, then $r^{\vee}p$ is in the support of the right side of (SP1), and hence has nonnegative time component for every r. The characterization just proved makes p future causal. Thus every diagonal measure is supported in the closed forward cone, and polarization gives

$$P(\overline{V}_+^{\,c}) = 0. \tag{SP2}$$

This is the spectral condition for the actual field translation representation, not merely for the hierarchy distributions.

12.32 Reflection-positive heat contractions

The chamber heat flow is obtained from future-word translation, rather than postulated. On finite sums p of future words let Xp be the reflection-positive word vector and $T_t p$ its forward translate through $t \geq 0$. Translation and reflection of the Euclidean hierarchy give

$$\langle XT_t p, XT_u q \rangle = \langle Xp, XT_{t+u} q \rangle. \tag{OS1}$$

Its fixed-word correlation bound implies $\sup_{t \geq 0} \|XT_t p\| < \infty$ for every finite sum: first bound each individual translated word norm by the square root of its reflected two-word correlation bound, then sum the finitely many coefficient-weighted bounds.

Put $c = \sup_{t \geq 0} \|XT_t p\|$. Equation (OS1) and Cauchy–Schwarz imply

$$\|XT_t p\|^2 = |\langle Xp, XT_{2t} p \rangle| \leq \|Xp\| c.$$

Taking the supremum and using $c \geq 0$ yields $c \leq \|Xp\|$. Thus forward translation preserves null vectors and extends to a contraction $E(t)$ on the completed reflection-positive Hilbert space. Equation (OS1) makes $E(t)$ self-adjoint. Translation parallel to the reflection plane preserves the Gram form, extends to an isometry, and its opposite translation is its inverse. These give the commuting spatial unitary group.

The semigroup and commutation laws hold on planted words and hence everywhere. Strong continuity follows first on planted words: the squared norm of the difference of two translated vectors is a sum of continuous Euclidean correlation functions. Finite sums inherit continuity; their density and the uniform contraction bounds extend it to the Hilbert space by approximation. The same argument proves strong continuity of the spatial group. This constructs the self-adjoint contraction semigroup and commuting spatial translations required throughout the wave construction.

12.33 Holomorphic heat powers and the real-time edge

Let $E(t)$ be the self-adjoint contraction semigroup just constructed. Each $E(t)$ is positive, since $E(t) = E(t/2)^2$, so its spectrum is contained in $[0, 1]$. Put $A = E(1)$. For $\Re z > 0$, define

$$F(z) = A^z$$

by continuous functional calculus, with $0^z = 0$. The function $\lambda \mapsto \lambda^z$ is continuous on $[0,1]$, including at zero. On compact subsets of the right half-plane it varies continuously in uniform norm; point evaluations are holomorphic in z. Equivalently, uniform estimates for $\lambda^{\Re z}|\log\lambda|^k$ on smaller half-planes justify its derivatives. Thus F is holomorphic in operator norm there. Functional calculus gives

$$\|F(z)\| \le 1, \quad F(z+w) = F(z)F(w), \quad F(z)^* = F(\bar{z}). \tag{HE1}$$

It also commutes with the spatial translations.

These powers coincide with the supplied heat flow. For positive integers p, n, let $B = E(1/n)$, so $A = B^n$ and $E(p/n) = B^p$. Functional calculus on the nonnegative spectrum and the scalar equality $(\lambda^n)^{p/n} = \lambda^p$ give $F(p/n) = E(p/n)$. Approximate any positive t by positive rational numbers; operator-norm continuity of F and strong continuity of E, applied to each vector, give $F(t) = E(t)$.

Every $E(t)$ is injective. If $E(t)v = 0$, the semigroup law gives $E(s)v = 0$ for $s \ge t$. The holomorphic vector function $F(z)v$ consequently vanishes on a real interval in the right half-plane, hence everywhere there by analytic uniqueness. Strong continuity at zero then gives $v = 0$. Self-adjointness implies dense range: the orthogonal complement of the range is its kernel.

For real s, (HE1) gives

$$\|F(r - is)v\| = \|E(r)v\| \qquad (r > 0),$$

because both squared norms are the quadratic form of $E(2r)$. Define an isometry on the dense range of $E(1)$ by

$$V(s)E(1)v = F(1 - is)v.$$

It extends uniquely to the Hilbert space. Checking on that dense range with (HE1) gives

$$V(s)F(z) = F(z - is) = F(z)V(s), \qquad V(s)V(t) = V(s+t), \qquad V(0) = 1.$$

Therefore $V(s)$ is unitary, with inverse $V(-s)$. It commutes with spatial translations. Its strong continuity holds on the range of $E(1)$ by operator-norm holomorphy of F, and on all vectors by density and the isometric bound. Together these operators form the chamber's strongly continuous real spacetime translation group. The sign is $F(r - is)$, consistent with the spectral convention $e^{i\langle p,a\rangle}$.

12.34 The Gaussian half-plane support argument

We prove the precise support fact used for the chamber diagonal spectral measures. Suppose a finite positive measure μ, a continuous real function ω, and a bounded function h, continuous on the closed upper half-plane and holomorphic in its interior, satisfy

$$h(t) = \int e^{it\omega(q)}\, d\mu(q), \qquad |h(z)| \le C.$$

For $b \in \mathbb{R}$, $\epsilon > 0$, multiply by $e^{-ibz - \epsilon z^2}$. Cauchy's theorem on rectangles between imaginary heights zero and $s \ge 0$, followed by sending their vertical sides to infinity, gives

$$\left| \int_{\mathbb{R}} h(t) e^{-ibt - \epsilon t^2}\, dt \right| \le C\sqrt{\pi/\epsilon}\, e^{bs + \epsilon s^2}. \tag{PA1}$$

The vertical side integrals vanish because their norms are bounded by a constant times $e^{-\epsilon T^2}$; the horizontal integrals are absolutely integrable Gaussian majorants. Fubini and the scalar Gaussian Fourier integral also give

$$\int_{\mathbb{R}} h(t) e^{-ibt-\epsilon t^2}\, dt = \sqrt{\pi/\epsilon} \int e^{-(\omega(q)-b)^2/(4\epsilon)}\, d\mu(q). \tag{PA2}$$

If a support point q_0 had $b = \omega(q_0) < 0$, the open set $U = \{q : (\omega(q) - b)^2 < b^2/4\}$ would have finite positive measure a. Choose $s = -b/(2\epsilon) > 0$ in (PA1). Equations (PA1)–(PA2) and positivity imply

$$a\, e^{-b^2/(16\epsilon)} \le \int e^{-(\omega-b)^2/(4\epsilon)}\, d\mu \le C e^{-b^2/(4\epsilon)}.$$

Thus $0 < a \le C e^{-3b^2/(16\epsilon)}$, impossible as $\epsilon \downarrow 0$. Hence $\operatorname{supp}\mu \subset \{\omega \ge 0\}$.

Apply this to the chamber's diagonal translation measure with $\omega(q) = q^0$. Its required upper-half-plane extension is

$$h(z) = \langle v, V(\Re z) E(\Im z) v\rangle, \qquad \Im z \ge 0.$$

It is continuous to the edge, holomorphic inside because it equals $\langle v, F(-iz)v\rangle$, and bounded by $\|v\|^2$. At real z it is precisely the characteristic function of the time marginal of the joint translation measure. The support argument proves nonnegative time energy. Combined with (SP1), it completes the forward-cone spectral conclusion (SP2).

12.35 The ground compression algebra and its probability space

The phase space is the spectrum of the commutative ground compression algebra. Its probability measure is obtained from the vacuum state on that algebra. Character evaluation gives the phase moments; their barycenter returns the unselected hierarchy. The countable common family of probes and the uniform phase bounds are what allow positivity, regularity and vacuum uniqueness to be asserted simultaneously.

Let $\mathcal{H}_0$ be the heat-fixed subspace of the unselected chamber and P_0 its orthogonal projection. Compress each slab operator to that subspace:

$$m(w, d) = P_0 \operatorname{place}(d, w)\big|_{\mathcal{H}_0}.$$

Take the norm-closed unital star algebra generated by these operators. Once its commutativity and countable norm density are established below, it is a separable commutative C^*-algebra $\mathcal{A}$. Its character space Θ is compact and metrizable. To see metrizability explicitly, take a dense sequence $a_k \in \mathcal{A}$. The map $\chi \mapsto (\chi(a_k))_k$ is continuous and injective: equality on the dense sequence extends to all of $\mathcal{A}$ by continuity of each character. A continuous injection from the compact character space into a countable product of complex planes is an embedding.

The vacuum vector state has the representation

$$\langle \Omega_0, a\Omega_0\rangle = \int_\Theta \chi(a)\, d\nu(\chi). \tag{PH1}$$

Here the Gelfand representation identifies $\mathcal{A}$ with continuous functions on Θ, and the representation of a positive functional by a finite measure supplies ν. It is a probability measure since $\|\Omega_0\| = 1$.

It has full support. The vacuum is cyclic for $\mathcal{A}$ on $\mathcal{H}_0$, as follows from the compression identity below. Commutativity then makes it separating: if $a\Omega_0 = 0$, then $a(b\Omega_0) = b(a\Omega_0) = 0$ on a dense set, so $a = 0$. If a nonempty open subset of Θ had measure zero, choose a nonzero continuous bump

supported there. Its squared integral is zero, so the represented operator kills Ω_0, contradicting this separating property. Thus every nonempty open set has positive measure.

The compressed operator of a distinct word is independent of how it is placed in a slab. Increasing the slab width adds heat evolution on its right, which fixes the input ground vector. Increasing its initial time adds heat evolution on the left, which disappears under P_0. Spatial translation disappears on both sides because ground vectors are spatially invariant. Positive or negative common time translations reduce to these identities by comparing the two placements. Denote the resulting algebra element by $M(w)$. A component hierarchy S_χ is specified by

$$S_\chi(w) = \chi(M(w)). \tag{PH2}$$

For a future word b, let $X_b : \mathcal{H}_0 \to \mathcal{H}$ be its slab insertion with ground input. This is independent of the trailing slab width, and

$$X_b\Omega_0 = [b], \qquad M(b)\Omega_0 = P_0[b], \qquad M(b^*c) = X_b^* X_c. \tag{PH3}$$

The last equality is slab concatenation and the reflected adjoint identity. Since the future-word vectors span densely, their ground projections do too; the middle equality gives cyclicity of Ω_0 for $\mathcal{A}$. Equations (PH1)–(PH2) also give the barycenter of every distinct Euclidean word:

$$\int_\Theta S_\chi(w)\, d\nu(\chi) = S(w). \tag{PH4}$$

Choose any common slab placement to prove this equality, apply (PH1), and use the vacuum slab expectation and translation invariance. The source constructs component hierarchies satisfying (PH2); their existence and regularity must use these same compression elements, rather than arbitrary disintegrated point values.

12.36 Almost-everywhere one-dimensional heat-fixed space

For a future word b in a component chamber, set $a_\chi = P_{0,\chi}[b]_\chi$. The heat limit gives

$$\|a_\chi\|^2 = \lim_{s\to\infty} \|E_\chi(s)[b]_\chi\|^2.$$

Use integer s. Each expression before the limit is a continuous character reading of the reflected translated word. It is bounded by its value at $s = 0$, a continuous integrable function on the compact character space, because heat is contractive. Dominated convergence and (PH4) therefore give

$$\int_\Theta \|a_\chi\|^2\, d\nu = \|P_0[b]\|^2.$$

On the other hand, multiplicativity of a character and (PH1) imply

$$\int_\Theta |S_\chi(b)|^2\, d\nu = \|M(b)\Omega_0\|^2 = \|P_0[b]\|^2.$$

The vacuum pairing satisfies $\langle \Omega_{0,\chi}, a_\chi\rangle = S_\chi(b)$. Since the component vacuum has norm one,

$$\|a_\chi - S_\chi(b)\Omega_{0,\chi}\|^2 = \|a_\chi\|^2 - |S_\chi(b)|^2 \geq 0.$$

Its integral is zero by the preceding identities, so it vanishes almost everywhere for this word.

Choose one countable sequence of future words whose linear spans of planted vectors are dense in every component chamber; the source obtains such a sequence from the countable probe construction, rather than assuming separability of arbitrary label sets. Intersect the countably many

full-measure sets just obtained. For each character in that intersection, the projections of all those word vectors lie in the vacuum line. Their projected span is dense in the heat-fixed space, because the ground projection is continuous and onto. The vacuum line is finite dimensional and closed, so the heat-fixed space equals it. This proves the almost-everywhere vacuum property and supplies an admissible selected character. The earlier full-support continuity argument upgrades the corresponding final reconstructed uniqueness statement to every character.

12.37 Spatial invariance of the heat ground space

Write $P_0 = \lim_{t\to\infty} E(t)$ strongly. The existence of this limit and its identification with the heat-fixed projection follow from the positive contraction spectral calculus: on the spectrum of $E(1)$, the integer powers λ^n decrease to $1_{\{1\}}$, and strong dominated convergence applies; general times follow from the semigroup law and contraction bounds. Its range is exactly the vectors fixed by all $E(t)$.

For any two distinct Euclidean words p, q, place their reflected and untranslated copies in the positive half-space. Their correlation with separation x, for sufficiently large time component, has the form

$$S(p\,\tau_x q) = \langle v, U(\mathbf{x})E(x^0 - r - s)w\rangle \tag{GI1}$$

for fixed future-word vectors v, w and fixed offsets r, s. This is an exact translation and reflection identity. Along any sequence with $x^0 \to \infty$, adding a fixed purely temporal displacement ue_0 changes (GI1) by a quantity tending to zero: both heat vectors converge to $P_0 w$, and the intervening, possibly varying spatial unitaries preserve the norm of their difference.

The same limiting insensitivity holds when adding any fixed b with $b^0 > 0$, along the pure-time separation te_0. Choose a Euclidean orthogonal transformation sending b to $\|b\|e_0$ whose transformed e_0 has positive time component. Such a transformation is obtained by reflection in the bisecting hyperplane between b and $\|b\|e_0$; preservation of their inner product gives the positive time component. Rotate both clusters and their displacements. The previous purely temporal comparison applies because the rotated large separation still has time tending to infinity. Euclidean covariance then gives

$$S(p\,\tau_{te_0+b}q) - S(p\,\tau_{te_0}q) \longrightarrow 0. \tag{GI2}$$

Given a spatial displacement a, apply (GI2) to $b = a + e_0$, and subtract the purely temporal comparison for adding e_0 to the sequence $a + te_0$. This proves that adding a has no effect on the limiting correlation. Taking heat limits in (GI1) yields

$$\langle v, U(a)P_0 w\rangle = \langle v, P_0 w\rangle$$

for all future-word vectors v, w. Density implies $U(a)P_0 = P_0$, so every ground vector is spatially invariant. This argument does not assume spatial invariance while comparing the varying large separations.

12.38 Long cluster products and commutativity

Let $w_1, \dots, w_r$ be distinct-point cluster words, and let their displacements $d_k(N)$ have pairwise ordered time separations tending to $+\infty$. Then

$$S\left(\prod_{k=1}^{r} \tau_{d_k(N)} w_k\right) \longrightarrow \langle \Omega_0, M(w_1)\cdots M(w_r)\Omega_0\rangle. \tag{GI3}$$

To prove this, first place each word in the positive half-space by a fixed shift, compensated in its displacement. Induct from the last cluster backwards. Between successive cluster insertions there is a heat duration tending to infinity and an arbitrary spatial translation. If vectors $v_N \to v$ and durations $t_N \to \infty$, contractivity and the just-proved spatial invariance give

$$U(a_N)E(t_N)v_N \longrightarrow P_0 v.$$

Indeed its difference from $P_0 v$ has norm at most $\|v_N - v\| + \|E(t_N)v - P_0 v\|$. Apply the fixed bounded insertion operator and continue the induction. Vacuum pairing gives (GI3). The large time separations also ensure that the combined word has distinct points eventually.

The four-cluster construction proves the needed exchange of two compressed factors. In the plane of Euclidean time and one spatial coordinate choose displacement slopes

$$(0,0), \quad (1,2), \quad (2,1), \quad (3,3)$$

for clusters a, b, c, d, respectively, and multiply all slopes by N. Their time order is a, b, c, d. Interchange those two Euclidean axes: the time order is now a, c, b, d. At finite N, permutation symmetry and Euclidean covariance identify the two combined expectations. Taking (GI3) gives equality between the first compressed product and the rotated second product. Remove the simultaneous rotation using another application of (GI3), placing the four clusters along a direction with both axis components positive, such as $e_0 + e_1$; its time order remains the same before and after the interchange. Therefore

$$\langle \Omega_0, M(a)M(b)M(c)M(d)\Omega_0 \rangle = \langle \Omega_0, M(a)M(c)M(b)M(d)\Omega_0 \rangle. \tag{GI4}$$

Choose the exterior words in (GI4) as the reversed reflected word of a future probe and as another future probe. Equation (PH3) and the adjoint relation turn (GI4) into equality of matrix elements of $M(b)M(c)$ and $M(c)M(b)$ between projected future-word vectors. Those vectors span densely in $\mathcal{H}_0$, so the two bounded operators are equal. Compressed generators are closed under adjoints: reflection in the slab sends one generator to its adjoint. Thus all generators and their adjoints commute. Algebraic star polynomials in them commute, and norm continuity of multiplication extends this equality to their norm closure. This proves commutativity of the actual compression algebra used in (PH1).

The resulting spatial ground invariance, cluster limit and commutative compression algebra will now supply one compact phase space and one countable family of future probes for all its characters.

12.39 One countable family of future probes for every phase

The countable family used in heat-vacuum selection can be chosen before any character chamber is given. For each arity n and tuple of finite label degrees $d = (d_1, \dots, d_n)$, use the parameter space consisting of distinct positions in the positive half-space and raw coefficient vectors in the finite-dimensional spaces of those degrees. This is a separable metrizable space; choose a countable dense subset. There are only countably many n, d. The union of the realized future words from these subsets, together with the empty word, is therefore a single countable list.

For any smooth symmetric hierarchy and any of its reflection-positive realizations, the planted-vector map on each such parameter space is norm-continuous. Its Gram pairing at two parameters is the hierarchy evaluated on their reflected and joined configurations. These configurations remain distinct because the two blocks are on opposite sides of the reflection plane. To see joint continuity in the coefficients, expand in fixed bases of the finite label spaces. The reading is a finite sum of products of continuous coordinate coefficients and smooth fixed-label hierarchy readings. The

adjoint coefficients in the reflected block are continuous as well. Thus both the diagonal Gram value and the mixed pairing with a fixed planted vector are continuous; expanding the squared norm difference proves norm continuity.

Every future word has some arity, degree tuple and parameter in this list of parameter spaces. Its vector is therefore a norm limit of vectors from the corresponding countable dense subset. Since all future-word vectors have dense linear span, the chosen common countable list has dense linear span in every such chamber. This proves exactly the simultaneous density needed when intersecting the countably many full-measure conditions in phase selection.

12.40 Character readings, regularity, and positive quadratic forms

The source constructs the component moment at a distinct configuration X by first forming the compression-algebra-valued multilinear moment $m_d(X; v)$, and then applying the character:

$$S_{\chi d}(X; v) = \chi(m_d(X; v)). \tag{CH1}$$

Outside the distinct-point set it uses the specified zero extension. The algebra-valued moment is the same whole-word compression $M(w)$ used in (PH2), not a product of separate point insertions. Its permutation, unit insertion, common translation, adjoint and quotient-label identities therefore transfer through the character. Once individual Euclidean rotation and reflection invariance of these compression elements is established, those identities transfer too.

Smoothness in X is also a consequence of the actual algebra-valued moment. A continuous character on the closed compression subspace extends to a bounded linear functional on the ambient bounded-operator space by the norm-preserving extension theorem. Compose that extension with the smooth ambient operator-valued moment. On the compression subspace it agrees with the original character, proving smoothness of (CH1) on distinct configurations. This argument avoids imposing an extra differentiability structure on the character algebra.

Reflection positivity follows by applying χ to one positive quadratic expression. For finitely many future words b_i and coefficients c_i, set $F = \sum_i c_i X_{b_i}$. Equation (PH3) gives the algebra element

$$a = \sum_{i,j} \overline{c_i} c_j M(b_i^* b_j), \qquad a\big|_{\mathcal{H}_0} = F^* F.$$

It is positive in the ambient bounded-operator algebra. Spectral permanence for a closed unital C^*-subalgebra makes its spectrum in the compression algebra nonnegative as well. Since $\chi(a)$ belongs to that spectrum, it is real and nonnegative. Expanding $\chi(a)$ by linearity gives the component reflection-positive Gram sum. This proves positivity at each character, with no exceptional set.

For derivative estimates, suppose positional differentiation is represented by replacing a label v_i by its prescribed derivative label $\partial_a v_i$. At distinct configurations it then holds for iterated derivatives as well, by differentiating the local identity:

$$\partial_{X_i^a} D^j S_{\chi,v} = D^j S_{\chi, v[i \leftarrow \partial_a v_i]}.$$

If each fixed-label scalar reading has a separation bound Cr^{-A}, induction gives such a bound for every fixed derivative order. Specifically the operator norm of a linear map on joint position space is at most the sum of its norms on the coordinate unit vectors. At the induction step take the maximum of the finitely many exponents supplied by the replacement labels and the sum of their constants. Since $0 < r \le 1$, all terms are bounded by that common exponent. Quotient compatibility transfers the result to the raw component labels. All constants remain dependent on the fixed arity, labels and derivative order; no unproved uniformity in those data is introduced.

The next sections construct the compression-valued moments, prove norm separability and individual Euclidean invariance, and transfer the derivative and shell identities. These supply the premises used in the character formulas above.

12.41 Constructing the compression-valued multilinear moments

Fix a distinct configuration and choose one translation and one strict slab containing it. The same translation and slab contain all sufficiently nearby configurations. When varying a label at a fixed position, they also work for every coefficient vector of that label degree, since the slab condition depends on positions alone.

For future spectator words p, q, the matrix element of this slab operator is the hierarchy reading of the reflected left word, the inserted slab word, and the translated right word. All these positions are distinct. Move a distinguished insertion to the front by permutation symmetry. Its reading is then an exterior-field functional, with those actual two spectator words as exterior data. By definition of the all-spectator null quotient it depends only on the projected field label and is linear in that label.

Equality of every such matrix element implies equality of bounded slab operators, because the future-word vectors have dense span on both sides. Thus the slab operator itself depends linearly on the projected distinguished label. Compression is a bounded linear operation, so the same is true for its element of $\mathcal{A}$. Move any chosen slot to the front to obtain this result at every slot. It proves the multilinearity of $m_d(X; v)$. Replacing labels one at a time proves that equal projected tuples give equal compression elements. Off the distinct-point set the zero convention preserves these algebraic identities.

The identical matrix-element argument proves permutation symmetry and deletion of a unit label before compression. The empty compressed word is the identity because its slab operator is heat evolution and acts identically on ground inputs. Reflection of the slab gives the adjoint of its compressed word. Common-translation invariance was proved by placement independence. These establish the structural identities used in (CH1) without assuming commutativity or applying a character.

Smoothness in position uses the locally fixed translation and slab. On that neighborhood the compression-valued moment is exactly

$$X \longmapsto P_0 \ \mathrm{place}(d, \tau_s w(X))\big|_{\mathcal{H}_0}.$$

The slab regularity estimate gives smoothness in operator norm at fixed d; the position translation is smooth and the compression map is bounded linear. The resulting smooth function agrees locally with the canonical compression moment by placement independence, proving its smoothness on the entire distinct-point set.

12.42 Norm separability and common phase bounds

For each arity and degree tuple, fix basis labels in every slot. The image of the distinct configuration space under its compression moment is separable in operator norm, since that space is separable and the moment map is norm-continuous. The union of these images over all basis tuples, degrees and arities is still separable. Multilinear expansion puts every compressed word in its complex linear span. Finite products of a separable set form a separable set, by continuity of multiplication and countable union over the product length. Include adjoints and finite complex linear combinations, and take the norm closure. These operations preserve separability and generate exactly $\mathcal{A}$. This supplies the countable norm-dense sequence used for metrizability of its character space.

Characters have the uniform bound $|\chi(a)| \le \|a\|$. A word with separation r, $0 < r \le 1$, can be translated into a slab with one unit of padding on both sides; its mutual separation remains r. Compression has norm at most that of the slab operator. Applying the unselected arbitrary-copy slab estimate therefore bounds the entire character reading with precisely the same copy constants, uniformly over characters and all receiving hierarchies with the same specified word values. For fixed arity n and bounded degrees and coefficient norms, the source's radius estimate is of the form

$$|S_{\chi d}(X; v)| \le K^n n^{Bn} r^{-Bn},$$

with its empty-word convention understood. This gives the fixed-label scalar separation bound required for the derivative induction above. Thus the component bounds come from whole-word operator norms, not from separate uncontrolled character norms.

12.43 The reflection square from bounded lattice approximants

For a fixed generic grid word w with real coefficient labels, let w^θ be its geometrically time-reflected word with the prescribed reflected labels, and set

$$D = M(w) - M(w^\theta), \qquad a = \Re\langle \Omega_0, D^2 \Omega_0 \rangle.$$

The grid construction supplies real measurable approximants $f_i^{(k)}$ to widely translated copies of this reflection difference, with $|f_i^{(k)}| \le R_k$, and pair correlation errors at most $\epsilon_k \to 0$. Fix the approximation index k, a finite number $q+1$ of copies, and their displacement parameter before taking the ordered mesh/volume limit. The four terms of each difference-product converge to the corresponding continuum two-cluster readings. Next send the displacement parameter to infinity. Equation (GI3) makes each off-diagonal comparison tend to $a + \epsilon_k$.

At finite lattice parameters positivity of an ordinary real square gives

$$0 \le \mathbb{E}\left(\sum_{i=0}^{q} f_i^{(k)} \right)^2 .$$

Bound the diagonal terms by R_k^2, and the two symmetric off-diagonal terms by the prescribed pair comparisons. Passing through the limits in the stated order yields

$$0 \le (q+1)\big(R_k^2 + q(a + \epsilon_k)\big). \tag{RS1}$$

Only finitely many pair conditions are needed at each q, so their eventual tails can be intersected before the limit. If $a + \epsilon_k < 0$, choosing q large contradicts (RS1). Thus $a + \epsilon_k \ge 0$ for every k, and $a \ge 0$.

For real labels, the adjoint and permutation identities give $M(w)^* = M(w^\theta)$, so $D^* = -D$. It follows that

$$\|D\Omega_0\|^2 = -\Re\langle \Omega_0, D^2 \Omega_0 \rangle = -a.$$

Both sides' nonnegativity and $a \ge 0$ force $D\Omega_0 = 0$. The separating vacuum property of the commutative compression algebra then gives $D = 0$.

To extend the identity beyond the generic grids, choose increasingly fine grids from the common funded approximation sequence and approximate a fixed generic configuration by their canonical grid addresses. The grid points are eventually generic and converge to it. Norm continuity of both compression moments gives equality there. Generic configurations have full measure in the distinct-point region; continuity and full support of Lebesgue measure extend equality throughout that open region. Finally expand the multilinear compression moment in real coordinate basis labels, each fixed by coefficient conjugation. Equality on those basis tuples proves equality for arbitrary complex labels. Thus individual geometric time reflection fixes each compressed observable.

12.44 Individual Euclidean invariance of compressed words

Long two-cluster limits first give simultaneous rotation invariance of pair stages

$$\langle \Omega_0, M(p)M(q)\Omega_0 \rangle.$$

If a Euclidean orthogonal map r sends e_0 to a vector with positive time component, use clusters escaping along e_0 in (GI3). If that component is zero, use instead the direction $e_0 + r^{-1}e_0$, which has positive time component both before and after rotation. If it is negative, first compose r with time reflection; individual time-reflection invariance removes the added reflection from each compressed factor afterwards. These three cases prove simultaneous pair invariance for every Euclidean orthogonal map.

Conjugate the individual time reflection by an arbitrary orthogonal map r. Rotate a pair stage by r^{-1}, reflect its second compressed factor only, and rotate both factors back. The first factor returns to itself and the second has been transformed by $r\theta r^{-1}$. Hence its pair readings with every first word agree with their originals. Test with the adjoint word of the second factor and of its transformed copy. The Gram identity then makes the norm of their vacuum-vector difference zero. The separating vacuum makes the compressed elements equal.

Every hyperplane reflection is a conjugate of θ, except the identity case of a zero normal vector. Orthogonal transformations in finite dimension are finite products of hyperplane reflections. The individual invariance just proved is closed under multiplication by the label representation law. Therefore

$$M(r \cdot w) = M(w) \qquad (r \in O(4)).$$

This completes the Euclidean rotation and reflection identities used for every component hierarchy in (CH1). The reflection-square argument uses its bounded Wilson approximants before taking limits, and the rotation argument uses the same compressed whole words throughout.

12.45 Inheritance of derivatives and stress-window identities

Fix a character component $T = S_\chi$. Extend χ, as a bounded complex-linear functional only, from the closed compression algebra to all operators on the ground space, then compose with compression. This gives a single bounded linear functional ϕ on chamber operators such that

$$S \operatorname{useAt} T(X; v) = \phi(\operatorname{place}(D, \tau_s w(X; v))) \tag{HW1}$$

whenever the translated distinct configuration lies strictly in the slab $(0, D)$. No positivity or multiplicativity is used for ϕ outside the original algebra.

One common s, D can contain any given finite distinct configuration together with any specified compact set of potential insertion points: take lower and upper bounds for their time coordinates and add a strict margin. This placement persists on a neighborhood of the finite configuration. Inclusion identities for finite field labels hold for the slab operator itself, by pairing with arbitrary future words on both sides and using the original all-spectator inclusion identity. Applying ϕ proves inclusion compatibility in T.

Likewise norm differentiability of the fixed-width slab insertion gives

$$\partial_{X_i^a} \operatorname{place}(D, \tau_s w(X; v)) = \operatorname{place}(D, \tau_s w(X; v[i \leftarrow \partial_a v_i])).$$

The common translation has derivative the identity on position space. Differentiate (HW1) through the bounded linear functional, keeping that width and translation fixed throughout its neighborhood. This proves the component derivative identity used in the preceding separation estimates.

For stress identities, let χ_0 be a compact smooth window equal to a constant near the distinguished insertion and zero near every spectator insertion. Pack its entire closed support in the same slab. For a continuous scalar weight u, form the operator integral

$$\mathcal{I}_{a,u} = \int (\partial_a \chi_0)(y)\, u(y)\ \mathrm{place}(D, \tau_s w(y, X))\, dy. \tag{HW2}$$

This is a norm-integrable operator-valued function. Its support is compact. On the support of its scalar factor, y cannot collide with the distinguished point, because the window is locally constant there; it cannot collide with any spectator, because the window vanishes near each of them. The insertion is therefore norm-smooth locally on that support. Off it the complete integrand is locally zero, even if the bare insertion at a collision was defined arbitrarily. These observations prove continuity and compact support of the full integrand.

To prove the required identity for (HW2) as an operator, pair it with two arbitrary future-word vectors. The added exterior spectators lie strictly below or above the whole chosen slab, and hence outside the entire window support. The original stress-window identity therefore applies with these genuine additional spectators. Passing the bounded matrix-element functional through the integrable operator expression proves the operator identity by density. Pass ϕ through the same integral and use (HW1); this gives exactly the numerical identity in T.

In particular, if T_{ab} denotes the prescribed degree-four stress label and $Y_b B$ the translation derivative of a distinguished label B, the identity is

$$-\sum_a \int (\partial_a \chi_0)(y)\, T(T_{ab}(y), B(X_0), v(X_{\mathrm{sp}}))\, dy = T((Y_b B)(X_0), v(X_{\mathrm{sp}})). \tag{HW3}$$

For a skew Euclidean generator g, use $u_b(y) = [g(y - X_0)]_b$ in (HW2) and sum over a, b. The corresponding identity is

$$-\sum_{a,b} \int (\partial_a \chi_0)(y)[g(y - X_0)]_b\, T(T_{ab}(y), B(X_0), v(X_{\mathrm{sp}}))\, dy = T((Z_g B)(X_0), v(X_{\mathrm{sp}})). \tag{HW4}$$

The geometric window and the same stress and generator labels are retained. Thus both translation and rotation stress laws transfer at every character.

12.46 Smeared reflection positivity of the components

The temporal point Gram positivity already proved transfers to any coordinate reflection plane by a signed coordinate permutation taking that axis to time. This permutation intertwines the two word reflections, preserves the positive side, and preserves the component readings by their real-spin identity. It therefore transports the complete finite Gram sum.

Point positivity also implies the required positivity for joint smearing tests. Fix an axis and construct the Gram Hilbert space from its positive future-word kernel. For a compact joint test $f(X)$ supported on positive, mutually distinct configurations, let

$$\zeta_f(X) = f(X)[w(X)]$$

there, and define it as zero elsewhere. Its diagonal and mixed Gram pairings are the corresponding smooth smeared hierarchy kernels. They are continuous, including across the cutoff support boundary. The squared-norm criterion proves strong continuity of ζ_f; its compact support then gives an integrable Hilbert-valued function. Fubini and the Gram formula yield

$$\iint K_{\mathrm{smear}}(f, g; X, Y)\, dX\, dY = \left\langle \int \zeta_f(X)\, dX, \int \zeta_g(Y)\, dY \right\rangle.$$

Every finite smeared reflection-positive sum is consequently the squared norm of a finite linear combination of these integrated vectors. This proves nonnegativity for joint tests directly from point positivity, without assuming a full hierarchy in order to prove one of its clauses.

12.47 Assembling the bounded reflection witnesses

The real bounded witnesses in (RS1) are obtained from the same local observable approximants used in the continuum comparison. For a real label observable O, suppose a local complex approximant O' has $|O'| \le A$ and rooted discrepancy at most ϵ. Replace it by $\Re O'$. The observable O is real for the prescribed real label, and

$$\Re O' - O = \tfrac{1}{2}(O' - O) + \tfrac{1}{2}\overline{(O' - O)}.$$

The rooted seminorm is subadditive, homogeneous and invariant under conjugation. Its discrepancy is therefore still at most ϵ. Local dependence and measurability are preserved, and $|\Re O'| \le A$. This supplies genuinely real approximants, without requiring the cutoff observables themselves to be uniformly bounded.

For a fixed finite list of degree-bounded labels, take products of their real local approximants at the prescribed translated lattice sites. The product comparison applies to the concatenated list of two clusters. Distinct clusters with integer displacement slopes $i < j$ are eventually disjoint: any coordinate difference has a fixed finite contribution plus $t(j - i)$, which tends to infinity. At the finite period, injectivity after wrapping also holds on the period tail. Thus the same product estimate bounds each of the four combinations of an ordinary cluster and its reflected cluster, on the same extraction order.

Let $f_i^{(k)}$ and $g_i^{(k)}$ be the resulting real product approximants for these two clusters. Set the reflection witness to $f_i^{(k)} - g_i^{(k)}$. Its uniform bound is the sum of their two product bounds. Expand the pair product into its four terms. If their individual comparison costs are $c_k A_{00}, c_k A_{01}, c_k A_{10}, c_k A_{11}$, the triangle inequality gives total error

$$\epsilon_k = c_k(A_{00} + A_{01} + A_{10} + A_{11}) \longrightarrow 0.$$

All four integrals exist by the local approximant estimates and finite lattice moment bounds. This is precisely the common witness family used in the finite-square argument. Fixed approximation index, fixed finite number of copies, large displacement, and the prescribed ordered mesh/period limits retain the ordering stated there.

The source obtains the individual local approximants by first fixing a reporting scale and an admissible finite atlas, then choosing a sufficiently late auxiliary scale. The gain-chart denominator makes the rooted approximation error small: its decaying scale factor dominates the subexponential column and stock factors. It then takes the real part as above and transfers the resulting estimates to the original mesh and period tails, using locality on the common chamber. The following coarse-lens and rooted-approximation argument supplies the gain-chart estimate, with the same local chamber and transported coefficients.

12.48 Coarse lenses, coefficient recovery, and rooted approximation

At a fixed coarse scale k, the lens of a label L is an explicit bounded local function:

$$\mathcal{L}_{k,L}(U) = a_k L_{\mathbf{1}} + 1_{\mathrm{Healthy}(\vartheta_k,\mathcal{F})}\, \mathcal{C}_{k,L}(U). \tag{AP1}$$

Here a_k is the prescribed field scaling factor, $L_{\mathbf{1}}$ its constant coefficient, $\mathcal{F}$ the local stencil footprint, and $\mathcal{C}_{k,L}$ the scaled candidate polynomial of its nonconstant coefficients. The other regular activity

terms of the primary lens are identically zero in this specialization. The candidate is bounded on the healthy compact chart region, uniformly in the finer terminal cutoff and period, and is multiplied by zero outside that region. Thus the bound for (AP1) is $|a_k L_{\mathbf{1}}| + A_{k,L}$. Its read set is contained in the coarse window of radius equal to the stencil radius plus four: each inspected edge is within four sites of a stencil point, so the triangle inequality gives the inclusion. This proves the lens locality and uniform bound used for the real witnesses.

The finite coefficient inversion in the approximation is quantitative. In the coefficient norm $|\cdot|_1$, suppose the recovered coarse coefficient vectors Jb_i satisfy

$$|e_i - Jb_i|_1 \le \tfrac{1}{4}$$

for every coordinate basis vector e_i. Let $Av = \sum_i v_i Jb_i$. Then

$$|v - Av|_1 \le \tfrac{1}{4}|v|_1, \qquad |v|_1 \le 2|Av|_1. \tag{AP2}$$

The first inequality follows by expanding the difference and summing coordinate norms, and the second by the triangle inequality. Hence A is injective and, in finite dimension with equal domain and codomain, surjective. For a terminal label L, solve this linear system and use the injective coefficient map J to recover

$$L = \sum_i v_i \,\mathrm{terminal}(b_i), \qquad |v|_1 \le 2e^{\mathrm{columnLogs}(k)}|L|_1. \tag{AP3}$$

The last factor is the earlier backward coefficient budget. Denote its right side by $\mathrm{stock}_k(L)$.

Each basis lens, lifted back from scale k to the fine lattice, approximates its recovered terminal observable in the root seminorm, with common error at most $\mathrm{stakes}_{J,m}/w$. Their linear combination with (AP3) is a fine-lattice approximant B to O_L, with

$$\|B\|_\infty \le \mathrm{stock}_k(L) A_k, \qquad \|B - O_L\|_{\mathrm{root}} \le \mathrm{stock}_k(L)\frac{\mathrm{stakes}_{J,m}}{w}. \tag{AP4}$$

Back transport preserves the sup bound and sends the coarse window into the prescribed common fine chamber, by the atlas room condition. Subadditivity and homogeneity of the root seminorm prove the second bound after the exact terminal-label decomposition. Thus both locality and the approximation concern the same B.

The basis root bound comes from an actual many-copy insertion. Use the separated atlas copies of the observable difference, with the compulsory mark at each copy, and extract the coefficient containing every copy once. Exact RG transport preserves its Wilson integral. The insertion estimate bounds that coefficient by

$$|M|\, w^q \le e^{qc}, \qquad c = (a_J/a_m)^4 \,\mathrm{treasury}_m + 5,$$

where q is the number of atlas cells and M is the corresponding Wilson moment of the product of copied differences. The term $(a_J/a_m)^4$ follows from the exact cell census, and the remaining five units are the insertion and vault costs in this application. Taking the q-th root gives the root estimate $e^c/w = \mathrm{stakes}_{J,m}/w$. The coefficient identity cancels the same positive Wilson partition function on both sides; it is not an estimate in a different normalization.

The coefficient construction requires

$$w\,|a_m|\frac{4(1+\mathrm{columns}_k)}{\mathrm{gainChart}(k,m)} \le \tfrac{1}{2}\zeta_{d,m}. \tag{AP5}$$

Fix the reporting data J, m, d first. Choose w at the upper value allowed by (AP5). Then the error in (AP4) is a fixed constant times the product of the stock and column factors divided by the gain chart. Those factors have subexponential growth, whereas the inverse gain chart contains the strictly decaying scale factor. Their product tends to zero; hence a sufficiently late k, also satisfying the atlas room condition, makes (AP4) smaller than any prescribed positive tolerance. Fixing that k fixes the sup bound as well, uniformly on all later mesh and period parameters. Taking the real part as above supplies the real bounded approximation needed for the reflection witnesses.

12.49 Transporting the approximation difference with its paid mark

The comparison used in (AP4) retains a single common source history. Starting from the primary coefficient L at scale k, the backward coefficient construction chooses a reference label b and a residual history Z_i, $m \le i \le k$. Its identities at the upper comparison scale are

$$\exp(Z_k) = \exp(\text{primary } L) + \exp(\text{history}(-b)), \qquad \text{forced}(Z_k) = \text{forced}(\text{history}(-b)). \tag{AP6}$$

These are exact identities for the evaluated exponential entry and the forced activity profile. They follow by linearity of the history map and cancellation of the reference label in the primary/reference decomposition. In particular the primary contributes no forced profile here. Its residual coefficient is $L - J_k b$, with norm at most $Q/4$; the backward cone estimate gives this bound, while the terminal/reference coefficient is bounded by $2E$.

The remaining residual cost at scale i obeys

$$\text{cost}(Z_i) + |Z_i^{\text{coeff}}|_1 \le E/\text{gainChart}(k, i).$$

Transport the mark price from w at m by the common multiplicative schedule. The scalar reserve recurrence and (AP5) imply, throughout $m \le i \le k$, that the transported price times the scaled residual cost is at most $\zeta_{d,i}/2$. The same bound pays for the reference-label contribution at k, using its coefficient bound $2E$. These are the two half-reserves in the construction.

Now insert, at each separated atlas copy, the negative terminal reference observable, transport its marked polynomial exactly from N to k, and multiply there by the corresponding coarse lens insertion. Equation (AP6) identifies the resulting first marked slot and forced slot with the residual-history slots. The paid insertion transport theorem then carries those slots from k down to m, along the same migrating copy positions. The separation and occupancy bounds are obtained from the chosen atlas at each intermediate scale. The resulting insertion at m has the served-source bound with support exponent of that scale and the two vault bounds $2, 1$ used in the root extraction. The residual-history estimates pay both bounds; no new history or mark schedule is chosen during this transport.

The coefficient identity relating this insertion to the actual fine-lattice discrepancy is algebraic. For a finite census, the new coarse insertions multiply as $\prod_i (1 + \mathbf{z}_i Q_i)$. Extracting a prescribed multi-index coefficient sums over the subset of slots taken from this product; every other slot is the coefficient transported from the negative fine observable. Coefficient extraction commutes with each finite RG push by linearity and integrability. The resulting subset sum is exactly the joint fine/coarse difference polynomial. Integrating the push does not change its integral, and restoring the fine/coarse coupling gives the positive Wilson partition function times the discrepancy moment. This is the identity used before taking the all-copy coefficient in (AP4). Thus the comparison bound, its normalization, and its locality all refer to the same transported difference.

The backward-cone estimate used here follows from the one-step coefficient and remainder bounds, as follows. Use the coefficient norm $|u|_1$, including its unit coordinate. At a step of block

size L, put $q = L^{1/8}$ and $e = q^{-3}$. If v is the finer coefficient, let

$$y_\alpha = L^{d-\deg(\alpha)} v_\alpha, \qquad Y = |y|_1.$$

All exponents are nonnegative, so $|v|_1 \le Y$. The step supplies an invertible linear map B, commuting with deletion of the unit coordinate, with $\|B\|_1, \|B^{-1}\|_1 \le K$, $K \ge 1$. Its coefficient error and remainder costs $r, t \ge 0$ satisfy

$$|u - By|_1 \le e(|\operatorname{core} v|_1 + t), \qquad r \le K|\operatorname{core} y|_1 + et.$$

Applying the inverse bound, the triangle inequality, and $|\operatorname{core} v|_1 \le |v|_1 \le Y$, gives

$$Y \le K\big(|u|_1 + e(Y + t)\big), \qquad r \le KY + et. \tag{AP7}$$

Let $C_i \ge 1$ be the column allowance at scale i, and put $a_i = 32(1 + C_i)^2$. The room and bank choices give

$$a_i \ge 4K^2, \quad 4Ke \le 1, \quad 4Ka_{i+1}e \le 1, \quad (4K^2 + 1)qe \le 1. \tag{AP8}$$

For precision, the bank condition is

$$q \ge 100000\, M^6 (1 + C_{i+1})^2,$$

where $M \ge K$ is the neighboring price bound. It implies $q \ge 128K^2(1 + C_{i+1})^2 \ge 1$. Use $e = q^{-3} \le q^{-1}$ and $qe = q^{-2} \le q^{-1}$ to get the last three inequalities in (AP8). The room condition $C_i \ge 64K^2$, with the stipulated monotonicity of stage prices, gives the first.

Suppose $a_i|u|_1 \le r$. By (AP7) and $K \ge 1$,

$$4K|u|_1 \le Y + et.$$

Substitute this bound on $|u|_1$ back into (AP7). The first smallness inequality in (AP8) absorbs the resulting Y-term and gives $Y \le 4Ket$. Consequently

$$a_{i+1}|v|_1 \le a_{i+1}Y \le t, \qquad qr \le q(KY + et) \le t. \tag{AP9}$$

The same argument with $e = 0$ also holds directly. Thus the cone $a_i|u|_1 \le r$ propagates from the coarser endpoint to the finer endpoint, while the remainder cost gains the factor q in the opposite direction.

To obtain the history to which this estimate applies, start with the given terminal formatted insertion and its remainder bound. Descend the finite interval one scale at a time, using the same one-step marked transport and its linear coefficient decomposition. At each step update only the newly reached component of the history. Previously obtained components and their step identities are retained. This finite backward induction constructs the entire matched history; its terminal density is exactly the prescribed insertion plus the reference history. Subtract the reference coefficient determined at the coarsest endpoint. Linearity makes the residual coefficient at that endpoint zero, while preserving every density and step identity. Its nonnegative remainder cost therefore satisfies the cone there. This explains why the cone is available for the residual history, rather than being an assumption on an arbitrary primary history.

Let c_i be its remainder cost on the interval $[l, k]$, and v_i its residual coefficient. Iterating (AP9) gives

$$a_i|v_i|_1 \le c_i, \qquad \operatorname{gainChart}(k, i)c_i \le c_k, \qquad \operatorname{gainChart}(k, i) = \exp\left[\frac{\operatorname{scale}(k) - \operatorname{scale}(i)}{8}\right]. \tag{AP10}$$

The identity used in this iteration is $\mathrm{gainChart}(k,i) = \mathrm{gainChart}(k,i+1)L_i^{1/8}$. The terminal pairing bound is

$$c_k \le (1 + C_k)w + C_k|v_k|_1.$$

Since the cone implies $2C_k|v_k|_1 \le c_k$, it follows that $c_k \le 2(1 + C_k)w$. Using $|v_i|_1 \le c_i$ in (AP10) yields the claimed residual estimate in its explicit form:

$$\mathrm{gainChart}(k,i)(c_i + |v_i|_1) \le 4(1 + C_k)w. \tag{AP11}$$

The constant E used in the approximation argument is this terminal bound with its fixed coefficient and stock factors.

The numerical choices required for (AP8) can be made before selecting a history or a volume. At each stage choose its column allowance above the room threshold, and choose its block lower bound so that $L^{1/8}$ exceeds the displayed bank threshold. This is possible since $L^{1/8} \to \infty$; enlarging either the prescribed block lower bounds or the later delays preserves the choice. Equations (AP6)–(AP11), followed by the paid marker transport and root extraction, therefore establish the bounded local approximants on the original common family.

13 Separated and nested operator products

An operator-product expansion is an approximation at a fixed finite degree horizon. One homogeneous coefficient reader is chosen on the physical filtration and used for every phase and truncation. Lower-degree terms are subtracted with that same reader, which is crucial when an expansion is inserted into a second expansion. The estimates keep the cluster separation and spectator separation parameters explicit.

13.1 The separated expansion and its fixed homogeneous channels

Let E_d be the finite-dimensional field space of degree at most d, with the coherent inclusions already constructed. A reader is chosen once for this filtration, its real structure, and its Euclidean rotation action. For finite input degrees d_i, write $\Delta = \sum_i d_i$. The expansion proved below has the following precise form. If

$$\|x_i - o\| \le Mr, \qquad \|x_i - x_{i'}\| \ge \eta r \quad (i \ne i'),$$

and the spectator points satisfy

$$\|o - y_j\| \ge c, \qquad \|y_j - y_{j'}\| \ge c \quad (j \ne j'),$$

where $M \ge 0$ and $\eta, c > 0$, then for each fixed truncation D and slack $\epsilon > 0$, all sufficiently small $r > 0$ satisfy

$$\left| S_s(a(x)b(y)) - \sum_{j=0}^{D} \sum_{\alpha \in \mathcal{C}_j} C_{j,\alpha}(x,a;o)\, S_s(e_{j,\alpha}(o)b(y)) \right| \le r^{D+1-\Delta-\epsilon}. \tag{OE1}$$

Here s may be the unselected hierarchy or any character component. The threshold precedes every allowed x, o, y, s. The same coefficients serve every D, spectator and viewer. For each fixed channel,

$$|C_{j,\alpha}(x,a;o)| \le r^{j-\Delta-\epsilon} \tag{OE2}$$

on its corresponding small-radius tail. These are fixed finite-order asymptotic statements; no convergence as $D \to \infty$ is asserted.

To define the homogeneous pieces, begin with the real, rotation-averaged finite reader at level m. Apply it to one field at the origin to obtain linear maps $L_{m,q} : E_q \to E_m$. They respect inclusions in the second index, and for $q \le m$ reproduce that included field. Define recursively

$$K_m = 1 - \sum_{i<m} \iota_{i,m} K_i L_{i,m}, \qquad \pi_{m,q} = K_m L_{m,q}.$$

Thus, whenever $q \le m$,

$$\pi_{m,q}\upsilon = \iota_{q,m}\upsilon - \sum_{i<m} \iota_{i,m}\pi_{i,q}\upsilon, \qquad \sum_{i=0}^{m} \iota_{i,m}\pi_{i,q}\upsilon = \iota_{q,m}\upsilon. \tag{OE3}$$

For $q < m$, the second identity at level $m-1$, included into E_m, makes the first right side zero. Hence $\pi_{m,q}$ annihilates all lower-degree fields. Inclusion compatibility follows directly from that of $L_{m,q}$.

Set

$$H_m = \bigcap_{i<m} \ker \pi_{i,m}.$$

Induction in the recursive formula proves $K_m(E_m) \subset H_m$: apply each lower projection, use its inclusion compatibility, the induction hypothesis on the earlier homogeneous pieces, and cancellation of its one matching term. The same identities show that the projections fix their own homogeneous space and annihilate the others after inclusion. Choose the prescribed real compatible basis $e_{m,\alpha}$ of H_m, $\alpha \in \mathcal{C}_m$.

For an input cluster let $A_m(x, a; o) \in E_m$ be its finite averaged reader value, and define its new homogeneous output

$$h_m(x, a; o) = K_m A_m(x, a; o) \in H_m.$$

The coefficients $C_{m,\alpha}$ are exactly its coordinates in that fixed basis. Thus

$$h_m = \sum_{\alpha} C_{m,\alpha} e_{m,\alpha}.$$

The finite packet at horizon D is $\sum_{m \le D} \iota_{m,D} h_m$. Applying $\pi_{j,D}$ gives exactly h_j for $j \le D$, independent of D. For a single field at its center the packet is precisely its inclusion, by (OE3). Pairing the packet with a spectator correlation gives the finite sum in (OE1), by linearity and the component inclusion identities. These are exact compatibility identities, not asymptotic approximations.

The ground-algebra error is defined before selecting a viewer: it is the compressed whole cluster-and-spectator word minus the finite sum of compressed centered homogeneous outputs with those spectators. Evaluating it in the vacuum or any character gives exactly the scalar remainder in (OE1). Contractivity of all these evaluations and the Gelfand norm give equivalence between a common algebra norm bound and the uniform bound over character viewers. This is how the same error estimate will serve all phases.

The averaged reader used here has an explicit meaning. Write $P_{\le m}(x, a; o)$ for the old finite physical packet. For normalized Haar measure on the compact orthogonal group, put

$$A^0_m(x, a; o) = \int_{O(4)} X_m(k) P_{\le m}(k^{-1}x, X(k^{-1})a; k^{-1}o)\, dk,$$
$$A_m(x, a; o) = \tfrac{1}{2}\{A^0_m(x, a; o) + \overline{A^0_m(x, \bar{a}; o)}\}.$$

The first formula is used on the open set where every rotated readout configuration is admissible; outside that set the source assigns zero. All subsequent asymptotic uses eventually lie in the open

admissible set. Each packet is a finite sum, and its integrands and all positional derivatives are continuous uniformly on the compact rotation group on a sufficiently small neighborhood. Thus integration preserves smoothness and multilinearity there. Translation and permutation invariance follow term by term. Haar invariance, with k replaced by gk, gives orthogonal covariance. The second formula commutes with conjugation. For one field of degree $q \le m$ placed at the center, coherence of the old readouts gives $P_{\le m} = \iota_{q,m} a$. Equivariance makes the integrand constant after the exterior rotation; reality of the inclusions makes the second averaging leave this value unchanged. This proves the exact single-field property used in (OE3).

Induction in the triangular formula for K_m now proves that K_m commutes with rotations and conjugation: each earlier projection does so, and the inclusions intertwine both operations. Consequently H_m is invariant under them and h_m has the stated translation, permutation, reality and rotation laws. The real basis invoked above exists for a concrete elementary reason. If C is the conjugate-linear involution on H_m, then

$$v = \frac{v + Cv}{2} + i\,\frac{v - Cv}{2i}$$

expresses every vector as a complex combination of fixed vectors. A real basis of the finite-dimensional fixed space is complex independent: apply C to a complex linear relation and add and subtract to obtain separate real relations for its real and imaginary coefficients. It therefore is a complex basis of H_m consisting of real fields. This choice is made once at each m, before any truncation or exterior word is specified.

Here is the precise transfer of remainder rates through these choices. Use the source's logarithmic parameter $T \to +\infty$; a rate at most λ means that for every $\epsilon > 0$ the norm is eventually at most $e^{(\lambda+\epsilon)T}$. Suppose the cluster is of size e^{-aT}, is internally separated at that rate, and $a > 0$. If $Z_H = P_{\le H}(x, a; o)$, the reader comparison gives

$$A_n(x, a; o) - A_n(o, Z_H; o) \quad \text{rate at most} \quad (|d| - H - 1)a. \tag{OE4}$$

This statement holds with every varying rotation, not just each fixed rotation. Integration therefore preserves the rate uniformly: otherwise choose a violating rotation at each violating parameter. The conjugate comparison supplies the second half of the real average. Applying the fixed linear map K_n yields

$$h_n - \pi_{n,H} Z_H \quad \text{rate at most} \quad (|d| - H - 1)a. \tag{OE5}$$

In particular, at $H = n$, all old grades below n are annihilated by $\pi_{n,n}$. The remaining old grade n has rate $(|d| - n)a$; the error (OE5) has the better rate $(|d| - n - 1)a$. Finite summation proves the same rate $(|d| - n)a$ for h_n, and hence for every coordinate $C_{n,\alpha}$.

For completeness, the scalar remainder bookkeeping does not lose a grade. Suppose cluster-to-exterior separation has rate b, exterior internal separation rate c, with $0 \le b \le c$ and $b < a$. A centered field of degree j, paired with the exterior labels of total weight $|e|$, costs $jb + |e|c$. For a varying field of rate t, its cost is $t + jb + |e|c$: expand in a fixed finite basis and use the fixed-field bounds on each coordinate. Reconstruct Z_D by (OE3). The exact identity for the new remainder is

$$W(x, a; y, b) - \sum_{j=0}^{D} W(o, h_j; y, b) = [W(x, a; y, b) - W(o, Z_D; y, b)] + \sum_{j=0}^{D} W(o, \pi_{j,D} Z_D - h_j; y, b).$$

The old remainder has rate

$$\xi = |d|a - (D+1)(a-b) + |e|c.$$

By (OE5), the j-th correction has rate $(|d| - D - 1)a + jb + |e|c \le \xi$. The finite sum thus has rate at most ξ, uniformly for every varying viewer. This last quantifier is what permits passage to the ground-algebra norm, rather than merely pointwise character bounds. In the separated regime $T = -\log r$, $a = 1$, $b = c = 0$, these rates are precisely (OE1) and (OE2), after spending the arbitrary positive slack ϵ. To make the threshold uniform over all admissible geometries, a failure would select a bad geometry at arbitrarily small radii; restricting to those radii contradicts the just-proved rate along that selected path. This also explains why the coefficient family is common to every truncation and phase.

The finite readout comparison in (OE4) follows directly from testing the old scalar remainder. Each readout is a finite sum of fixed exterior configurations with fixed scalar weights. After a varying rotation, these exterior configurations remain a fixed positive distance from the shrinking center and from one another, and their rotated labels and loaded output vectors have rate zero, by compactness of the rotation group. The old scalar remainder at these moving tests has rate $(|d| - H - 1)a$. To allow their moving labels, expand every exterior label in a fixed finite basis; the coordinates have rate zero and multilinearity preserves the rate in a finite sum. Summing the fixed detector weights and output vectors proves the rotated readout comparison. Conjugating the whole correlation and using the hierarchy's reality identity proves its conjugate version with the same norm. This proves (OE4), including the necessary arbitrary varying rotations.

We detail the grading argument supplying the old estimates. At logarithmic scale T, choose lattice levels with mesh exponents

$$0 < b < a < x < d,$$

where the actual cluster has exponent x. Put

$$\tau = |g|d + 4\,|I|(d-a).$$

The physical approximation supplies, for an auxiliary horizon q, a centered label Q of degree q, an approximation amplitude, and a survival statement for their actual continuum limits. Its label budget has exponent $\tau - qa$, and its mismatch against all allowed exterior tests has exponent at most

$$\tau - (q + \tfrac{1}{16})(a-b). \tag{OE6}$$

The same mismatch is used at every detector; no refitting of the centered field to a detector or phase is allowed.

For a mode of degree $h \le q$, applying its physical readout to Q costs only $\tau - ha$. For $h = 0$ this follows from the whole-label budget. For $h > 0$, the approximation-to-lower-degree estimate gives a label s belonging to filtration degree $h-1$, with distance from Q bounded by its degree-$(h-1)$ fare times the original budget. The degree-h readout vanishes on s, by coherence and inclusion compatibility. Its value on Q therefore equals its value on $Q - s$; the bounded linear readout and the fare's exponent $(q-h)a$ give the asserted exponent. This subtraction is essential: using just the norm of Q would lose all of this graded improvement.

Choose $q \ge h$ large enough that

$$(h+1)a \le (q + \tfrac{1}{16})(a-b).$$

Testing (OE6) at the finite detector list and adding the preceding readout bound proves that the actual degree-h coefficient has exponent at most $\tau - ha$ on the lattice screen. For a truncation at

degree N, choose instead $q \geq N$ with

$$(N+1)a \leq (q+\tfrac{1}{16})(a-b).$$

Let the exterior clearance exponent be $c < b$, its internal exponent be $u \geq c \geq 0$, and its total weight be $|r|$. The centered degree-h correlation costs $hc + |r|u$. Reconstructing the centered Q through all its readouts, split its finite expansion into degrees $h \leq N$ and $N < h \leq q$. For the first group subtract the true coefficients; (OE6) pays each difference and the centered exterior correlation. For the second group use the graded budget

$$\tau - ha + hc + |r|u \leq \tau - (N+1)(a-c) + |r|u.$$

The original product-to-Q mismatch has this same upper exponent by the choice of q. The exact sum of these three differences is the desired product minus its degree-N truncation. Thus the lattice remainder has rate at most

$$\tau - (N+1)(a-c) + |r|u. \tag{OE7}$$

All statements retain arbitrary varying viewers throughout.

The passage from these lattice screens to arbitrary points uses the smooth interpolation estimate, applied to the coefficient, or to the product minus its fixed truncation. Its hypotheses consist of the rough correlation bound, the screen bound just proved, and the positional derivative bounds for these same functions. For coefficients the rough rate is $x|g|$, since the detector exteriors stay at positive distance. For remainders it is at most $x|g| + |r|u + Nc$, obtained by separately bounding the original product and the finitely many truncated terms. The derivative bounds are those of the same hierarchy and finite detector sums. Interpolation gives (OE7) with the displayed τ at arbitrary approaching configurations. Finally choose $a = x - \delta$, $d = x + \delta$, and an intermediate b strictly between the exterior exponent and a. Mesh levels with each of these positive exponents exist on the supplied clock. As $\delta \downarrow 0$, the coefficient exponent tends to $(|g| - h)x$, and (OE7) tends to

$$|g|x - (N+1)(x-c) + |r|u.$$

For any strict requested upper exponent choose $\delta > 0$ small enough first, then the mesh levels and auxiliary horizon. This proves the old rate estimates with their exact quantifier order.

Translation to a moving center changes no distances or readouts. The empty cluster requires no approximation: its coefficient is the readout of the degree-zero unit, its positive-degree coefficients vanish, and its truncation is exactly the identity insertion in every exterior word. This covers the cases excluded when an anchor in the cluster was chosen above.

Here are the details of the smooth interpolation step. Consider a family $f_T(p, z)$ of scalar functions, whose label list p has total degree $|p|$, and suppose every positional coordinate derivative replaces one label by its prescribed derivative, raising total degree by one. Assume that in a corridor of configurations of size and separation exponent a,

$$f_T(p, z) \text{ has rate } \leq a|p| + c,$$

while on screened points of mesh exponent $b > a$ it has rate $\leq b|p| + d$. Choose an integer n such that $(b-a)n \geq c - d$. We prove the better rate for every degree by descending induction with the statement that it holds whenever $|p| + k \geq n$. At $k = 0$ the rough bound suffices. For the induction step move z to a screened mesh configuration z' at distance of rate $-b$. Every point in the ball of that radius remains in the same corridor because $b > a$. Each derivative there is a function with

label degree $|p|+1$, so the preceding induction statement bounds its rate by $b(|p|+1)+d$. The finite coordinate basis bounds the operator norm of the derivative by the sum of its coordinate norms. The mean-value inequality on the ball gives

$$f_T(p,z') - f_T(p,z) \text{ rate } \le b(|p|+1)+d-b = b|p|+d.$$

Adding the screened value proves the induction step. Taking $k=n$ includes every original degree.

The required uniform derivative bound on this moving ball follows from the pathwise hypotheses without assuming uniformity: if it failed arbitrarily late, select a violating point in each such ball, and use that selected path in the derivative hypothesis. The same selection argument supplies eventual differentiability at every point of the ball. Screened approximation itself is elementary finite-dimensional rounding with finitely many forbidden coincidences: a fixed number of mesh shifts suffices to avoid them, with distance bounded by $2K$ times the mesh size, where K depends only on the number of cluster positions. Thus the displacement rate is indeed $-b$. Finally, the derivative replacement identity holds for the actual functions in question. It holds for full correlations by the hierarchy's derivative law, for a readout by its finite sum of full correlations, and for the remainder by subtraction and the finite product rule (the centered exterior factors do not vary with the cluster positions). No differentiability of a chosen approximating label is used.

We also specify the compactness and exterior-price passage in the physical matching step. A finite plot at levels $k \le m < m+1 \le j$ satisfies

$$8\|x_i\| < a_{m+1}, \qquad \|x_i - x_{i'}\| > 2a_j B_j^2 \quad (i \ne i').$$

At sufficiently large periods its wrapped positions retain their actual locations and are separated by B_j^2 lattice units, while every position belongs to the same central block through level $m+1$. The matching estimate therefore produces, along the original extraction order, degree-q labels whose products differ from the original cluster by a paid local insertion of amplitude A, and whose quota at level m is at most p. The quota comparison bounds their norm in the fixed finite-dimensional reference space by Kp. Finite-dimensional bounded compactness gives a nontrivial refining filter and a limit label Q. Both the paid mismatch and the quota bound persist on that filter. This is the precise survival statement used above: it retains the actual refinement filter, geometry, paid insertion, reference horizon and quota, in addition to convergence of the labels.

To test that same mismatch against an exterior word, first fix arbitrarily small positive slack ρ. The bounded-copy construction for an exterior field of degree e_j gives a per-field envelope

$$\mu_j = (1+Y_j)a_\ell^{-e_j-\rho} + 1.$$

At a mesh exponent ξ, its product has rate $(|e|+\rho|J|)\xi$. Choose auxiliary mesh exponents strictly between the exterior clearance exponent and the matching support exponent, then a further pair strictly above the exterior internal exponent. Their strict gaps ensure all finite support widths and separation requirements for the paid-insertion estimate; any fixed number of fields costs no exponent. The common survival filter makes the paid insertion estimate apply to the very same Q. After the continuum limit the compressed exterior mismatch has rate at most

$$\delta + (|e|+\rho|J|)\xi,$$

where δ is the amplitude's rate. Choose these auxiliary exponents with $\xi \downarrow u$, then $\rho \downarrow 0$, before requiring any eventual bound. This gives exactly $\delta + |e|u$. Contractive evaluation gives the identical bound for arbitrary varying phase viewers. The two-group full-correlation bound follows by the

same argument with the product of the cluster and exterior envelopes, giving engineering rate $|g|a + |e|u$. These arguments justify all exterior costs used in the grading and interpolation proof.

For the approaching plots, the strict exponent inequalities ensure their two geometric conditions eventually: the cluster size is smaller than $a_{m+1}/16$, its internal spacing exceeds $2a_j B_j^2$, and the selected level indices are ordered. Here B_j is subpower relative to the mesh, so it consumes none of the strict gaps. Consequently the compact matching limit exists at each approaching parameter. Selecting one such limit label pointwise gives the family Q used in (OE6), with the same amplitude and quota rates.

The matching accuracy is funded as follows. Let the chosen block level have exponent a, the input mesh exponent $d > a$, and the exterior support level exponent $b < a$. The individual input budgets have product rate $-|g|d$. The many-copy initial amplitude includes the volume rescaling factor of rate $-4|I|(d-a)$, hence is a positive quantity x_T of rate $-\tau$, with the same $\tau = |g|d + 4|I|(d-a)$ as above. Its degree-q terminal cost is explicitly

$$e_T = 4(1 + \mathrm{columns}_m)\, 2a_{m+1}^q \ \mathrm{bling}_m \ \frac{4}{x_T}.$$

The column count, the displayed chart factor, and the copy controls have rate zero; the adjacent level $m+1$ has the same mesh exponent as m. Thus e_T has rate $\tau - qa$. The output quota is $2e_T$.

Choose

$$z_T = \exp\Big(\big[(q + \tfrac{1}{16})(a-b) - \tau \big] T \Big).$$

The stock-opening level has the same exponent b as its starting level, and its stakes have upper rate zero. The error amplitude is

$$A_T = 2\, \frac{\mathrm{stakes}}{z_T},$$

which has the upper rate in (OE6). The two affordability conditions in the matching construction hold eventually. For the march condition, the available exponent is the full gain $(q+1)(a-b)$; for the chart condition it is $(q + \frac{1}{8})(a-b)$. The chosen expenditure uses only $(q + \frac{1}{16})(a-b)$. Both leave strict positive margins. All numerical constants, stocks, columns and chart prices have rate zero, so they are absorbed by those margins. The remaining room condition also follows from $b < a$: the clock satisfies $J_{m+1} \le J_m + 1$, while each block scale increment is at least $\log 3$. Hence $J_m + 5 \le J_0 + 5 + \mathrm{scale}_m / \log 3$ and $\log(J_m + 5)/\,\mathrm{scale}_m \to 0$. A fixed readout radius and this subpower extra width are dominated by the strict mesh gap.

These choices are inserted in the finite matching construction before the continuum limit. It yields a centered degree-q label with quota $2e_T$, a mismatch supported in the prescribed central chamber, and root-test bound A_T. The input positions are required to remain in that chamber through level $m+1$, exactly the plot conditions already checked. The finite matching construction itself is the sweep-and-replica construction; its remaining formulas are addressed below.

There is also no hidden independence assumption in the exterior test passage. On a common finite lattice the original cluster register and its centered replacement are multiplied by precisely the same exterior register. The algebraic identity is

$$(\mathrm{register}_x - O(Q,0))\, \mathrm{register}_y = \mathrm{register}_{x \sqcup y} - O(Q,0)\, \mathrm{register}_y\,.$$

The paid insertion bound applies to its left side as one main insertion and the given number of exterior insertions. Taking the original extraction limit along the survival filter gives the two terms on the right with the actual limiting label Q. On screened configurations this proves the norm bound for their compressed difference. For arbitrary exterior positions, approximate by screened lattice

configurations avoiding the finite set consisting of the center and all cluster positions, coordinate by coordinate. Their separation inequalities persist eventually, and the hierarchy is smooth at both noncoincident limiting configurations. For each character the scalar inequalities therefore pass to the limit. Taking the supremum over characters gives the compressed algebra norm bound. The same approximation, without a replacement, proves the corresponding bound for full products. Thus the extension to arbitrary exterior points uses actual continuity of the already constructed hierarchy and one fixed replacement, not a fresh replacement chosen after the test.

We describe the geometry and algebra of that finite matching construction more precisely. Make one copy of the input list in every cell of a fixed coarse tiling, with the prescribed reflection and translation of that cell. Write $F_{v,i}$ for these copied insertions and $q_v = \prod_i F_{v,i}$. The anchor copy is unchanged. At level $m+1$ all members of the original list have parent zero, so the parents of the copied list are exactly the tiling centers. At every finer level the copied position has the same coarser parent as its center and lies within half the intervening block length of the embedded parent. Two distinct centers are separated at the parent level by at least B_{m+1}^2+1. If the intervening enlargement factor is L, their embedded separation is at least $L(B_{m+1}^2+1)$; the two deviations cost strictly less than L. The propagated budget satisfies $B_i^2 \le LB_{m+1}^2$. The triangle inequality therefore gives the required cross-copy separation at every intermediate level. Within a copy the given finest-level spacing supplies the remaining separation condition.

Represent the copied insertions by a polynomial with one formal marker per occurrence (v,i). Descending integration preserves this polynomial and its integrability. Grouping markers within a copy into one marker replaces $F_{v,i}$ by q_v. Grouping in stages or all at once has identical coefficients: the coefficient of a specified marker set is the product over its full inverse image. Since grouping is a fixed linear map on a finite-dimensional polynomial space, it commutes with each conditional integral. This proves the conservation identity between the descended polynomial of the products q_v and the grouped descended polynomial of the individual $F_{v,i}$.

The descending sweep starts with positive single-field allowances

$$w_i = \delta_j \phi_j a_i/8.$$

At each level it joins precisely the fibers that have acquired a common parent, retaining one anchor in each fiber; thus the final number of fibers is at most the original arity. Its available normalization is bounded below by the compensation of $\prod_i w_i$, including the four-dimensional volume ratio between levels $m+1$ and j, and the licenses used in these joins. The initial allowance is

$$\alpha^{|I|}\,\mathrm{compensation}\left(|I|, m+1, j, 0, \alpha, (\delta_j\phi_j/8)^{|I|}\prod_i a_i\right).$$

The chosen $\alpha \le 1$ is no larger than any intermediate license, so fewer surviving fibers only increase this lower bound. This is the finite quantity whose rate was denoted x_T above.

The licenses can be chosen uniformly down to every preceding level while preserving rate zero. If $u_i > 0$ is the license and $\log u_i/\operatorname{scale}_i \to 0$, set $\alpha_i = \min(1, u_0, \dots, u_i)$. Then $0 < \alpha_i \le 1$ and $\alpha_i \le u_j$ for all $j \le i$. For any $t < 0$, the rate-zero hypothesis gives $u_j \ge e^{t\,\mathrm{scale}_j}$ for large j. Monotonicity of the scale implies $\alpha_i \ge \min(\alpha_{i_0}, e^{t\,\mathrm{scale}_i})$. Dividing logarithms by the scale and letting $t \uparrow 0$ proves that α_i still has rate zero. The halo purchase bounds the logarithm of the geometric part of each license by $16\sqrt{\widehat{\beta}_i}$, whose ratio to the scale tends to zero. The other license factors have the same rate-zero property. Thus this minimum introduces no power loss.

At the last join, let $z > 0$ be at most the surviving normalization and let $x \ge 1$ be the purchased local coefficient constant. The terminal normalization estimate gives a local analytic coefficient with

remainder bound $4/z$. Prune its insertion allowance to

$$I'_v = z/(16x).$$

This is no larger than the available normalization. It satisfies the two elementary fusion inequalities

$$2I'_v\, 2(4/z) \le 1, \qquad I'_v\, |\lambda^{-1}|\, x(4/z) \le 1,$$

where $\lambda = (\ell_m)^q \in \mathbb{C}$, and $\ell_m \ge 2$ is the integer block size at this step. This λ is a scalar multiplying every coefficient in this degree-q fusion; it is not an arbitrary linear change of frame. Since $q \ge 0$,

$$|\lambda| = \ell_m^q \ge 1, \qquad |\lambda^{-1}| = |\lambda|^{-1} = \ell_m^{-q} \le 1.$$

For $q = 0$ both moduli equal one. Equivalently, scalar multiplication by λ^{-1} has operator norm at most one on the coefficient space. Thus the first smallness expression is $1/x \le 1$, and the second is $|\lambda^{-1}|/4 \le 1/4$. A lower bound on the norm of a general operator would not imply this inverse bound; the scalar identity is what proves it here. The target march allowances satisfy $W_v \le \delta_m I'_v$, so their products meet the insertion substitution bound. Applying the local fusion estimate yields a coefficient Q_0 with

$$\|Q_0\|_{\mathrm{coeff}} \le p' = a_{m+1}^q x\, 4/z$$

and the same remainder budget, representing the descended products simultaneously in every copy. All transformations of one copy into another are inherited from those of the original $F_{v,i}$; the anchor is at zero. Restoring the polynomial coefficient, then applying the bounded-copy comparison to its actual local section, gives a physical label with quota at most $2e$ and root-test error 2 stakes $/b_0$, where b_0 is the chosen final march amplitude. These are exactly the two outputs used in funding, with $b_0 = z_T$.

The common cluster purchases used in this construction are made for every finite arity on the original family. They comprise the same positive activities, smallness inequalities

$$(2B_i + 1)^4\ \mathrm{scurry}_i \le \tfrac{1}{2}, \qquad e^{\mathrm{scurry}_i} - 1 \le 1,$$

the normalization lower bound, and the local fusion and analytic remainder estimates at every sufficiently late degree reference. The degree reference ensures the clock has reached at least degree plus three. These clauses are supplied by the already chosen scheduling and local analytic estimates; they do not entail replacing the Wilson family when the arity or output degree changes.

For the paid-product limit, a finite linear combination of screened words is encoded with its actual moving lattice labels. Its continuum compression and its lattice expression use the same scalar coefficients. Products of registers agree identically with the register of the concatenated list. Differences use the coefficients $1, -1$. Applying the paid-product inequality to this exact expression and passing to the selected continuum limit gives the compressed norm bound by the product of its insertion allowances. In particular it applies to the single mismatch times all exterior factors. This explains why the subtraction stays intact through the limit; estimating its two terms separately would discard the approximation accuracy.

The compensation inequality underlying the sweep can be written without any scale shorthand. If H is the original number of fields, s is the number of remaining groups, $0 < \alpha \le 1$, and $w > 0$, define

$$\mathcal{K}(i, j, s) = \left(\frac{a_j}{a_i}\right)^{4H} \left(\frac{\alpha}{8}\right)^{H(H-s)} w. \tag{OE8}$$

The sweep maintains $\mathcal{K}(i,j,s) \le \prod_{\gamma=1}^{s} W_\gamma$. At the finest level, $s = H$, so it is equality with $w = \prod W_\gamma$. When the groups stay separated at the next coarser level, their allowances become W_γ/L_i^4, where $a_i = L_i a_{i+1}$ and $L_i \ge 2$. Since $s \le H$,

$$\mathcal{K}(i,j,s) = L_i^{-4H}\mathcal{K}(i+1,j,s) \le L_i^{-4s}\prod_{\gamma} W_\gamma.$$

This pays the entire one-step transport without a loss that accumulates separately in each tiling copy.

If separation fails, select a nontrivial grouping, preserving the distinguished anchor, whose representatives are separated at the coarser level. Let its size be $s' < s$, and let ℓ_i be the purchased cover bound for its groups. The local grouping estimate gives the exact product of new group allowances as the old product times $\mathrm{license}(H,\ell_i,i+1)^s$; before transport each group also spends a factor 8. Since α is no larger than this license and

$$H(H-s') \ge H(H-s) + H, \qquad \alpha^H \le \alpha^s, \qquad 8^{s'} \le 8^H,$$

we obtain

$$\mathcal{K}(i+1,j,s') \le \mathcal{K}(i+1,j,s)(\alpha/8)^H \le 8^{-s'}\mathrm{license}(H,\ell_i,i+1)^s \prod_{\gamma} W_\gamma.$$

The preceding transport inequality then pays the coarser step. This proves (OE8) by finite descending induction, in both cases. A regrouping strictly decreases the number of groups, so the exponent $H(H-s)$ pays at most the fixed arity's worth of grouping events, independently of the number of scale steps and tiling cells. Grouping the marker polynomial and then pushing it is identical to pushing first and grouping afterward, by the finite-dimensional integral identity already proved. Integrability is likewise retained under each such linear grouping.

Every step also retains the actual representatives, not just their count: their parent relation is inherited from the original positions, their copy transformations are inherited from the original collars, and their anchor representative is the original anchor. If a group is selected at a coarser level, its cover is bounded by the purchased ℓ_i. These are exactly the data needed for the local grouping estimate and the final normalization. Thus the induction does not require consistency choices beyond the existing finite marker grouping.

For clarity about the norm used throughout, a compressed word at noncoincident positions is its ground-space multiplication operator M; it is assigned zero elsewhere only to make the notation total. For a character χ, $\chi(M)$ is its phase correlation. For the unselected hierarchy its evaluation is $\langle \Omega, M\Omega\rangle$, because translating the entire finite word into a positive-time chamber leaves its correlation unchanged and the compression formula realizes that translated word. Both evaluations are contractive: the latter by $\|\Omega\| = 1$, the former by the character norm. Conversely, the Gelfand norm is the supremum of the character values, so a common nonnegative scalar bound on all viewers is equivalent to the compressed operator norm bound. This identifies the all-phase norm in the OPE with the ground algebra already constructed.

The representative selection in the grouping step is finite and explicit. For the positions in the anchor copy choose a subset of maximum cardinality containing the distinguished anchor and having pairwise lattice separation at least B^2. Every omitted position is within separation $< B^2$ of one of the chosen positions, since otherwise it could be adjoined. Assign each position to such a representative, assigning representatives to themselves. This gives a surjection with a section fixing the anchor. When separation failed, the subset is proper, so its cardinality strictly decreases. Repeat this same relative selection in every copy. Reflections and translations preserve lattice

separation, and the cross-copy separation was already proved; hence the entire selected configuration is separated simultaneously.

The necessary finer-level two-point cover bound follows from the parent relation. If x, y have parents at separation $< B_m^2$, the distance from either point to its embedded parent is less than half the block length L_m. Thus

$$\operatorname{sep}(x, y) \le L_m(B_m + B_{m+1} + 1)^2 =: c_m.$$

A pair at distance at most c_m is contained in a halo of its first point. Covering this halo by B_{m+1}-blocks costs at most

$$(2c_m + 1)^4 \le 16(1 + c_m)^4.$$

This is the purchased carpet bound. It applies separately to each point and its representative; the construction does not assume that the entire group lies in a single small halo.

We give the remaining moment argument for the compressed norm passage. Let E be a fixed finite linear combination of screened words and let its labels vary along a refining regulator filter, converging in their fixed finite-dimensional label spaces. For each regulator pair, freeze the labels at that pair and use this ordinary static word in the finite even-moment expansion. For an even order $2h$, the expansion is a finite sum over choices of one word from E at each of the $2h$ occurrences. Its coefficients are the products of the scalar coefficients and their conjugates, and its terms are the corresponding translated and reflected full lattice correlations. This is an exact expansion of the even-moment expression, including every subtraction in E.

Each such lattice correlation converges along the refining filter with the moving labels. Indeed expand each label in a fixed finite basis, expand the multilinear correlation as the finite sum over all basis choices, and use convergence of the finitely many coordinate products and the fixed-label point limits. For every sufficiently large replica displacement the limiting positions are generic, so all these fixed-label limits apply. Therefore the finite even-moment expression converges first along the regulator filter to its continuum replica expression. Sending the replica displacement to infinity gives the vacuum expectation of the corresponding even power of the ground compression, by the long-cluster moment formula already proved.

If the finite paid-product estimate bounds the whole even-moment expression in norm by c^{2h}, its real part is at most c^{2h}. Both limits preserve that inequality, for every h. The ground-algebra even-moment norm argument then gives $\|M_E\| \le c$. Equivalently, one may apply the same moment inequality to each character and use the Gelfand norm. Here c is the product of the main and exterior insertion allowances; it is nonnegative because each paid allowance is nonnegative. Applying this construction to the exact register identity for the mismatch proves its norm bound. This proof uses finite-dimensional convergence of the moving label along the actual refinement filter; it requires no separate inner regulator limit for that label.

The local grouping allowances just used are

$$\lambda_i(H, \ell) = \frac{(\delta_i/H)\delta_i}{\exp((e_i + \frac{1}{2})\ell)}, \qquad W'_\gamma = \prod_{j \in \gamma} \lambda_i(H, \ell) W_j, \quad e_i = \text{supportExponent}_i\,. \tag{OE9}$$

Indeed each occurrence first receives $\delta_i W_j$, the allocation among at most H occurrences gives the factor δ_i/H, and moving its support to the chosen representative costs $\exp((e_i + \frac{1}{2})\ell)$. The occurrence sets of distinct groups partition the old occurrences, so $\prod_\gamma W'_\gamma = \lambda_i(H, \ell)^s \prod_j W_j$ exactly. This proves the group-product formula used in (OE8).

To obtain the grouped local functions, apply the finite marker fusion estimate to those occurrence sets, using the cover bound ℓ, the occupancy bound, and the two smallness inequalities

displayed above. The original insertion allowance W_j controls its actual local analytic norm; multiplication by $\delta_i \le 1$ preserves that bound. The positive background density permits normalization, and the representation of each lattice stage identifies this normalized density with the stage's actual background. Fusion produces one local function per group, together with dominated activity profiles, representing the exactly grouped marker polynomial. Their normalized remainder bound is $4/W'_\gamma$. Pruning each group allowance to $W'_\gamma/8$ rewrites this as

$$4/W'_\gamma = \frac{1}{2(W'_\gamma/8)},$$

the bound required to transport the group to the next scale. The leading counterterm of a newly grouped function is zero. The copied functions and profiles retain their common translations and reflections, since the same groups and representatives are used in each copy. For the terminal grouping there is just one group in each copy; no factor $1/8$ is needed before normalization, and the common allowance is $\prod_j \lambda_i(H,\ell)W_j$. Its remainder bound is 4 divided by this allowance, exactly the input to the terminal step above.

Finally, the paid even-moment bound can be seen directly from its lattice layout. The two groups of fields have separation scales a_b and a_ℓ. Choose buffer and insertion levels satisfying

$$b \le f \le \ell \le n, \qquad f \le Q,$$

and the finite packing inequalities

$$|I|\, 10a_f < a_b, \qquad |J|\, 6a_n < a_\ell.$$

The seating lemma places every original insertion with the required interior margins in the appropriate translated reflection chamber. Translation copies are displaced by the prescribed replica itinerary. Within one replica the original spacing is unchanged; between different replicas all distances tend to infinity with the displacement. At sufficiently large displacement, and then sufficiently fine regulator, every required lattice separation exceeds its chamber scale. The same interior margins hold in all replicas because the itinerary preserves the chamber margins.

At the $2h$ occurrences replace each field by its adjoint on the alternating even positions and keep it on the odd positions. Adjoining preserves its paid allowance. The mixed chessboard inequality on the two nested chambers now bounds the integral of the product of all these translated fields by the product of all their allowances, namely

$$\left[\left(\prod_{i\in I} B_i\right)\left(\prod_{j\in J} D_j\right)\right]^{2h}.$$

The register translation identity and conjugation of finite products identify this integrand with the even-moment expression of the entire registered insertion Z. For a mismatch, Z is the mismatch times the exterior register, so the main group has exactly one occurrence per replica with allowance A. Nothing expands that occurrence into two independently estimated fields. This proves the finite bound required by the compressed-norm limit.

We justify the nested-chamber inequality used in the last paragraph. A paid field has an interior support in its original central chamber and a bound A for that chamber's root test. Any permitted diagonal reflection and conjugation preserves both properties and its price. If its placement has a margin of three original chamber scales inside a larger chamber, the translated field is a copy, in the appropriate cell of the larger tiling, of another interior function with the same bound A. This is the offset-and-reflection identity for the tiling copy map.

First collect all finer fields assigned to one larger cell. Copy this collected product to every cell of the larger tiling. Within the finer tiling all individual copied fields occupy distinct cells: distinct larger cells cannot meet in a finer cell, and two original finer placements in one larger cell are farther apart than the finer cell's extent. Reflection preserves that separation. The finer chessboard inequality bounds the full copied product by the product of the finer root tests. Taking the root of order equal to the number of larger cells gives the collected product a larger-chamber root test at most the product of its individual prices. This root extraction is legitimate because the prices and root tests are nonnegative and the number of cells is positive.

Now the coarser main placements occupy distinct larger cells, and no one of them occupies a larger cell containing a finer exterior placement: otherwise the cell's extent would contradict the assumed main-main or main-exterior separation. Put each main field in its own cell, and each collected exterior product in its occupied cell. One application of the larger chessboard inequality bounds the full mixed product by the product of their root tests. Multiplying the collected bounds gives exactly $\prod B_i \prod D_j$. Thus the different insertion scales are retained; replacing them by a single common scale is unnecessary.

The common seating translation exists by two finite avoidance steps. First choose a translation placing the main positions at distance five buffer scales from the coarse cell boundaries; the total excluded intervals fit because $|I|10a_f < a_b$. Then shift by at most one intermediate cell scale to place the exterior positions at distance three fine insertion scales from their cell boundaries, using $|J|6a_n < a_\ell$. The second shift consumes at most two of the five main buffer scales, leaving the required three. The nested cell periods divide the larger periods. Replica displacements are integer multiples of these periods, so they leave every boundary residue and hence all margins unchanged. Actual torus separations follow from the Euclidean ones for sufficiently large periods: the physical distance is at most twice the mesh times the wrapped lattice separation. Subsequent embedding multiplies lattice separations by the mesh ratio, proving the required inequalities at every finer cutoff.

The one-step sweep transport uses allowances $W'_\gamma = W_\gamma / L_i^4$. Its local driving estimate multiplies the analytic remainder bound by $6L_i$; this fits since $6L_i \le L_i^4$ for $L_i \ge 2$. Thus

$$(6L_i)/(2W_\gamma) \le 1/(2W'_\gamma).$$

All marker tolls are finite sums of nonnegative products of allowances. Decreasing the allowances therefore preserves every toll estimate. The separated representatives give the support substitution bound, the occupancy hypotheses hold before and after taking parents, and the prescribed transition density gives the exact descended polynomial. The leading counterterms retain their copied scalar transformations, and $W'_\gamma |c_\gamma| \le 1$ follows from $W'_\gamma \le W_\gamma$. At the initial fine scale, the solved activity representation gives remainder bound at most $4/W_\gamma$; taking $8W'_\gamma \le W_\gamma$ converts this to $1/(2W'_\gamma)$. These are precisely the numerical transport and initialization bounds used in the descending compensation induction.

The marker fusion has the following quantitative construction. For a surjection of old occurrence labels onto new group labels, with a chosen representative of each group, extend every old support by the fixed connecting cover from its occurrence to its representative. If this cover costs at most ℓ, the new allowance per old occurrence is

$$\widetilde{W}_i = (\eta/H) W_i e^{-\alpha \ell}, \qquad \alpha = e + c.$$

An old coefficient involving a set A of occurrences therefore gains $(\eta/H)^{|A|}$ in its weighted envelope, after paying all connecting covers. The group allowance is the product of these $\widetilde{W}_i$ over its full fiber. Distinct original occurrences retain distinct markers, so whenever a product is eligible its

marker sets are disjoint; the product of their tolls is the toll of their union. The latter is exactly the product of the full group allowances for the selected new markers. This is the algebraic reason fiber sizes can vary without changing the coefficient identity.

The local coefficient inventory consists of two scalar-entry contributions and the activity-entry contributions. If the occurrence occupancy is at most s, the scalar coefficients have weighted bound a, the activity coefficients bound b, and $W_i|c_i| \le T$, extending their supports and charging representatives gives total sitewise envelope

$$\gamma = 2(1+s)\eta(sTe^{\alpha} + a) + (1+s)b. \tag{OE10}$$

For nonempty marker sets, $(\eta/H)^{|A|} \le \eta/H$. When a chosen representative is charged, there are at most H old siblings in its fiber; the factor H^{-1} cancels precisely that multiplicity. The representatives themselves satisfy the same occupancy bound, because the section is injective and their locations are old occurrence locations. These observations give the $1 + s$ incidence factor in each contribution. The scalar-entry inventory has two copies, producing the first factor 2; the activity inventory is charged once. Summing their nonnegative envelopes proves (OE10).

Expand the grouped coefficients as finite eligible products of these inventory entries. Give a product the product of its entry envelopes. The disjoint-marker identity in the preceding paragraph proves that this is a valid weighted coefficient bound. The connected-set forest estimate, at excess support exponent c, applies provided

$$(2B+1)^4\gamma \le c.$$

It bounds both the grouped ordinary coefficients and grouped activity coefficients by $e^{\gamma} - 1$ at support exponent e. This is the same connected-set estimate already proved for the general insertion calculus, now with the extended supports and group tolls. Selecting $e^{\gamma} - 1 \le Z$ gives the required two coefficient budgets Z.

Read the singleton coefficient of the grouped ordinary polynomial as the new local function and the singleton coefficient of the grouped activity polynomial as its activity profile. The singleton restoration identity and the normalized nonvanishing background give their exact restored polynomial and a local remainder bound

$$2Z/W'_{\gamma} + 2Z/W'_{\gamma} = 4Z/W'_{\gamma}.$$

The domination construction for the activity singleton retains the actual background activity coefficient at the empty marker. Hence it yields an admissible dominated activity profile, not only a numerical estimate on one coefficient. The higher coefficients give the simultaneous served-insertion representation. In the sweep specialization $c = 1/2$, $Z = 1$, these are exactly the bounds and hypotheses used in (OE9).

Copy covariance of these new singleton functions follows at the level of every grouped marker coefficient. A marker set focused in a given copy involves only inventory entries in that copy; its transformed set is focused in the transformed copy. The copy permutation is a bijection on these inventories, takes extended supports to their transformed supports, and conjugates each coefficient by the prescribed scalar transformation. Reindexing the finite sums and products proves covariance for the grouped ordinary and activity polynomials. Reading their singleton coefficients then proves covariance of the local functions and profiles. The same argument includes the empty marker and all interactions internal to a copy, so repeated regrouping preserves the hypotheses required at the next scale.

We spell out the inventory underlying this restoration. Write c_i for the scalar marker entry, A_S for the ordinary polynomial on a support S, and $K_{S,F}$ for the activity polynomial on a support and

its flagged subset. For every scalar c_i, take a singleton-marker entry on its occurrence location; for every nonempty marker coefficient of A_S, take that coefficient on S. Extend each of these supports by the representative locations of the groups mentioned by its markers. Split its scalar value into two entries: one equal to the value when this extended support is healthy and zero otherwise, and the other equal to the value when it is unhealthy and zero otherwise. Their sum is exactly the original value, at every configuration. Keep every allowed coefficient of $K_{S,F}$, with the same extended support and its original flagged data. These are the two scalar-entry inventories and the single activity inventory in (OE10). They are finite on each lattice.

Summing the two scalar entries recovers $\sum_i c_i z_i + \sum_S A_S$ coefficient by coefficient. Collecting activity entries by their original support-and-flag label recovers $K_{S,F}$, also coefficient by coefficient. Each nonzero coefficient has its marker locations inside its support by locality, so every term required for either recovery is present. The grouped construction forms eligible connected collections of these entries and assigns their union support, union markers, and flag intersection to the resulting coefficient. Expanding the marker exponential and activity product, then partitioning a finite collection into its connected components, gives

$$\mathcal{R}_\Delta\left[\exp_z\left(\sum_i c_i z_i + \sum_S A_S\right)\mathrm{Gas}_z(K)\right] = \exp_\zeta\left(\sum_S \widetilde{A}_S\right)\mathrm{Gas}_\zeta(\widetilde{K}). \tag{OE11}$$

Here $\mathcal{R}_\Delta$ reads the old coefficient on the full inverse image of each new marker set, and the exponentials are the finite exponentials in the nilpotent marker algebras. The eligibility condition requires disjoint old markers and complete group fibers. It is exactly what identifies the left coefficient with the right one; the connected-component partition is unique. The healthy/unhealthy split routes a component to the ordinary polynomial or the flagged activity polynomial without changing its scalar value. This is a finite polynomial identity, so it holds regardless of zeros of any individual entry.

The empty activity coefficient is preserved in (OE11). A connected eligible collection with no new marker can contain no scalar entry, since each scalar entry has a nonempty old marker. Starting from any one of its activity entries, a connection to a second entry would have to use a group marker or intersect a flag support. There is no group marker, and distinct old empty-marker activity entries have disjoint prescribed flagged supports. Thus the collection consists of a single old empty-marker activity entry. Conversely each such entry is eligible and contributes its original value and its original support-and-flag label. This bijection proves preservation of the empty coefficient before any envelope estimate. After restoring the nonzero background, (OE11) is the served-polynomial identity used to construct the new local functions.

The sitewise scalar-entry estimate in (OE10) can also be read directly from this inventory. At a site there are at most s singleton scalar entries, each of weighted size at most T, and a singleton support has cover number at most one. Their contribution is at most sTe^{α}. The nonempty coefficients of the ordinary polynomials contribute at most a. The two healthy/unhealthy copies are each bounded by this same envelope. Activity coefficients contribute at most b. Extending the supports and reallocating to representatives gives the $1 + s$ factors already described and the occurrence factor η, proving (OE10) with no omitted unit term.

For precision, the eligibility and connectedness in (OE11) are as follows. An old marker set A is retained by $\mathcal{R}_\Delta$ exactly when it is a union of complete fibers. Its new marker set is $\Delta(A)$. For any old monomial,

$$\mathcal{R}_\Delta(z_A a) = \begin{cases} \zeta_{\Delta(A)} a, & A = \Delta^{-1}(\Delta(A)), \\ 0, & \text{otherwise.} \end{cases}$$

For a product of retained monomials the new marker sets are disjoint exactly when the old ones are disjoint: an overlap in a new group would force its entire nonempty fiber into both old sets. Thus reading the full-fiber coefficients commutes with such eligible products.

Connect two inventory entries when their group rosters meet, or when their extended supports meet and both entries are of the type that must be collected into an activity. The latter types are the unhealthy scalar entries and all old activity entries. A grouped ordinary entry is a nonempty connected eligible collection containing no such type. A grouped activity entry is a connected eligible collection containing at least one old activity entry, with its total flagged intersection prescribed to be the union of the old flagged subsets. Its coefficient is the product of the inventory coefficients when this flag condition holds, and zero otherwise. The old activity supports inside a collection are required to be pairwise disjoint, as they were in the original gas expansion.

Partition an eligible old collection into components for this relation. A component containing an unhealthy scalar entry but no activity entry has zero contribution whenever the old flag pattern is selected: the unhealthy entry requires a flag in its extended support, and that flag must belong to an old activity support, which would connect that activity to this same component. Thus every contributing component is either an ordinary entry or a grouped activity entry as defined above. Conversely, a selected collection of grouped entries has disjoint group rosters, and its activity supports obey the selected disjoint-support pattern. Its expanded components therefore have no additional connections to each other. Taking their union and then taking components are inverse operations. The union of the old flagged subsets equals the union of the grouped flagged subsets. Finally, under the selected pattern, a grouped activity's own flag condition holds separately: a flag in its support assigned to another grouped activity would contradict disjointness of the grouped activity supports. The coefficient product is unchanged by this bijection. Summing proves (OE11).

Locality and measurability also survive this construction. Every grouped support contains its mentioned representatives. Each entry depends only on the original variables and flag pattern in its extended support; the product depends only on their union. Its flag indicator is measurable and gauge invariant, and every underlying coefficient is measurable and gauge invariant. A nonzero ordinary grouped entry is healthy on its entire support, since all its scalar entries are healthy there. A nonzero activity entry has exactly its recorded flag intersection. Thus the new ordinary and activity polynomials satisfy the same support and flag conditions as the old ones, enabling the next grouping step.

For the final identification of one-step transport, extract each singleton coefficient of the fine served polynomial. The singleton restoration law identifies it with the normalized local insertion, including its dominated activity correction. The local analytic transport estimate gives at the anchor one coarse local function and a dominated coarse activity profile, with the same leading counterterm and remainder multiplied by the local driving constant. Define the functions in every other copy by the prescribed reflection and translation of this anchor function; reindex the activity labels by the actual activity-motion bijection. The conditional integration intertwines these transformations, so their normalized insertions agree with the corresponding singleton coefficients of the pushed polynomial. Coefficient extraction commutes with this integration, since the fine polynomial is integrable. The marked transport theorem then gives the full coarse activity representation at weights δI, not just singleton identities. Finally the support substitution estimate with target weights J, the occupancy bound, and the one-step duty inequalities restores the served representation at J. These are exactly the analytic transport, singleton cancellation and served restoration estimates developed above, with occurrence set $V \times I$. No two-element property of the earlier occurrence set is used: all coefficient and support estimates are stated for an arbitrary finite marker set. This completes the transport identification required in the sweep.

13.2 Nested expansions

Nested expansions require an order of scales. The inner cluster is replaced first with the common coefficient reader, and its remainder is estimated before the outer replacement is summed. The scale-ratio gain applies on the declared nested geometric domain. Terms with zero exponent are retained with their stated logarithmic or subpower qualification, rather than being assigned a nonexistent positive power gain.

Nesting keeps one reader At a nested node, the output of the inner expansion is an input to the outer one. Both therefore use the same homogeneous channel labels and finite projections on the physical filtration. The lower-degree terms are retained before the next substitution; otherwise the outer reader would see a different coefficient. The estimates distinguish a small ratio of nested radii from a fixed geometric separation, and keep the zero-exponent cases in their stated scope. Each nested tree and degree horizon is finite.

Let t be a finite rooted tree, with leaves carrying input fields and with every internal vertex carrying a center z_v and radius $s_v > 0$. A leaf also has its insertion position. Vertices may have zero or one child; all finite arities are allowed. At every internal vertex all positions in its subtree lie within Ms_v of its center, and positions in distinct immediate child subtrees are at distance at least ηs_v. Radii of internal children are at most their parent's radius. The root radius is $r \downarrow 0$. There is no lower bound on any child-to-parent ratio other than positivity.

These conditions may equivalently be obtained from bounds on immediate child centers after requiring a sufficiently small uniform upper bound on child-to-parent radii. To see this, induct over the finite tree. Let C_i be the already obtained displacement constants of the child subtrees and $C' = \sum_i C_i$. Choose

$$0 < \beta \le 1, \qquad \beta \le \eta/[4(C'+1)]$$

and below every child's previously chosen ratio bound. The displacement of any descendant from its immediate child center is at most $C'\beta s_v$. Two such deviations reduce the immediate sibling separation by at most $2C'\beta s_v \le \eta s_v/2$. The displacement from the parent center is at most $(M + C')s_v$. This proves the stated full-subtree geometry recursively, with separation $\eta/2$.

Write w_ℓ for each input degree, $W_v = \sum_{\ell \text{ below } v} w_\ell$, and

$$q_v = s_v/s_{\mathrm{parent}(v)}$$

for a nonroot internal vertex. The virtual parent radius of the root is 1, so $q_{\mathrm{root}} = r$. For fixed output degree m, define the common price

$$\mathcal{T}_{m,\epsilon}(s) = r^{\,m-\sum_\ell w_\ell-\epsilon} \prod_{\substack{v \text{ internal} \\ v\neq\mathrm{root}}} q_v^{-W_v-\epsilon}. \tag{NE1}$$

Choose fixed truncation degrees D_v at the internal vertices, before taking $r \to 0$. The direct coefficient is the coefficient $C_{m,\alpha}$ of the entire leaf cluster at the root center. The iterated coefficient is the finite sum obtained by expanding each nonroot child cluster through its assigned degree, substituting its homogeneous basis fields at that child's center, and repeating upward to the fixed root output m, α. Multiplication of the successive coefficients and summation over all intermediate degrees and channels is meant literally. These are the same C's and homogeneous bases used in (OE1), with no new reader choice.

Uniformly in the admissible positions and positive radii, for sufficiently small r,

$$|C_{\mathrm{direct}}| \le \mathcal{T}_{m,\epsilon}(s), \qquad |C_{\mathrm{direct}} - C_{\mathrm{iterated}}| \le \mathcal{T}_{m,\epsilon}(s) \sum_{\substack{v \text{ internal} \\ v\neq\mathrm{root}}} q_v^{D_v+1}. \tag{NE2}$$

If there are no nonroot internal vertices, the iterated and direct coefficients agree exactly; the sum on the right is zero. The degree assigned to the root does not enter that sum, since the root output has already been fixed.

For exterior fields whose positions stay at least $\delta > 0$ from all tree positions and from one another, the full leaf product minus its direct root expansion through degree m has ground-algebra norm at most

$$\mathcal{T}_{m+1,\epsilon}(s). \tag{NE3}$$

The same estimate holds for the unselected law and every character, uniformly in their choice. The coefficients, the finite iterated sums and the tested expressions are smooth on the open set where all their finitely many reader configurations are admissible. For every fixed collection of requested degrees the geometry above eventually belongs to this open set. This follows recursively from admissibility of each child-center list and the direct leaf list; their diameters tend to zero and their required distinctness follows from the sibling separation.

We explain the scale accounting used to prove these assertions. If s_v has logarithmic rate $-a_v$, with nondecreasing a_v down the tree and virtual parent exponent zero, set $\Delta_v = a_v - a_{\mathrm{parent}(v)} \geq 0$. Then the engineering cost of the original leaves is

$$\mathcal{E} = \sum_{v \text{ internal}} W_v \Delta_v. \tag{NE4}$$

This is also the sum over leaves of their degree times their nearest internal parent's exponent: each leaf's path telescopes. Thus a direct output of degree m has rate $\mathcal{E} - m a_{\mathrm{root}}$, and a tested remainder through degree m has rate $\mathcal{E} - (m+1)a_{\mathrm{root}}$. The price (NE1) has exactly this engineering rate plus $\epsilon \sum_v \Delta_v$. The latter is strictly positive after normalizing $\sum_v a_v = 1$: at least one $a_v > 0$, and each a_v is at most the sum of the nonnegative advances along all vertices. Therefore it absorbs every fixed multiplicative constant, even when individual advances are zero.

We prove the recursive rate assertions, keeping all truncations finite. A leaf has just one alternative: its original field at its position, with coefficient 1. An internal child's alternatives through horizon D are the homogeneous fields of degrees $0, \ldots, D$, at its center, weighted by their direct coefficients. For any specified alternatives of the children, the parent coefficient is the product of their coefficients times the parent OPE coefficient for their resulting fields. A bud is never replaced by a degree-zero projection: its stage set is a singleton carrying its actual input degree. This convention includes arbitrary input labels correctly.

The main inductive assertion has two parts. Every raw subtree can be represented, up to any prescribed tested error rate, by a finite family whose positions all have separation at its parent's exponent and whose coefficient amplitudes pay its total engineering cost minus the total resulting field degree times that parent's exponent. The same raw subtree has direct coefficient and full tested remainder bounds at its own exponent, with its engineering cost. For a leaf both statements are immediate. At an internal vertex, combine the child families by finite products. Their costs add exactly to (NE4). Sibling separation makes the combined positions admissible.

If a child's exponent equals its parent's, retain the child's comparable finite family obtained from its first inductive assertion. If it is strictly larger, use its second assertion to replace it through a sufficiently high auxiliary finite horizon H. Specifically, if its engineering cost is Ω, the parent exponent is k, its own exponent is $a > k$, and later exterior costs are bounded by Wb, the tested error rate is at most

$$\Omega - (H+1)(a-k) + Wb.$$

Given any requested rate ρ, choose

$$H > (\Omega + Wb - \rho)/(a-k),$$

and above any previously requested horizon. This makes the error smaller than requested. Perform the finitely many child replacements in order. At each step the preceding finite replacements and their field degrees have already been fixed, so choose the next auxiliary horizon to pay their actual exterior cost. The exact telescoping sum contains one replacement error at a time. This triangular choice avoids circular dependence among the horizons. Finite readouts of the tested errors then give the same comparison for coefficients. Applying the single-center estimate to the resulting comparable finite family establishes both assertions for the parent.

To compare with prescribed, rather than auxiliary, truncations, fix a gain $\gamma \geq 0$ such that

$$\gamma \leq (D_v + 1)\Delta_v \quad \text{for every nonroot internal vertex } v. \tag{NE5}$$

Inductively, each recursively composed child coefficient has its direct coefficient's engineering rate, and their difference improves that rate by γ. The product identity

$$\prod_j x_j - \prod_j y_j = \sum_j (x_j - y_j) \prod_{i<j} y_i \prod_{i>j} x_i$$

therefore improves the product rate by γ, not by a multiple of γ. Multiplication by the parent coefficient preserves this improvement.

For $\gamma > 0$, every nontrivial immediate child has a strictly positive exponent gap, by (NE5). Replace all such children through auxiliary horizons $H_j \geq D_j$ with error of any desired rate, using the preceding triangular argument. Subtract the finite parent sum through the desired horizons. Its difference from the auxiliary sum is exactly the sum over omitted channel choices. In every omitted choice at least one child degree is at least $D_j + 1$. That child's coefficient costs its degree at the finer exponent and the parent coefficient returns it at the parent exponent; hence the choice gains at least $(D_j + 1)\Delta_j \geq \gamma$. All other degree-gap contributions are nonnegative. Thus the omitted finite sum gains γ. Replacing the remaining direct child coefficients by their recursive coefficients gains the same γ by the product identity. Adding these three differences proves that the direct-to-iterated coefficient error has rate at most

$$\mathcal{E} - ma_{\mathrm{root}} - \gamma. \tag{NE6}$$

If $\gamma = 0$, direct and iterated coefficients already have the common raw rate, so their difference has that rate by the triangle inequality; no high-accuracy replacement is needed at a zero exponent gap. If there are no nonroot internal vertices, there is one choice at every child, its coefficient is 1, and permutation invariance identifies the parent sum with the direct coefficient exactly.

Take γ to be the minimum of the finitely many $(D_v + 1)\Delta_v$. Choose a vertex attaining this minimum. Its ratio power $q_v^{D_v+1}$ has rate $-\gamma$, while the positive slack in (NE1) absorbs the fixed constants in (NE6). Its single positive ratio term is bounded by the sum in (NE2), proving the claimed sum of gains. The direct coefficient and exterior remainder rates similarly give (NE2)'s first inequality and (NE3). Contractivity and the character supremum give the latter in the common ground-algebra norm.

The finite-family calculus in that induction has explicit quantifiers. An alternative is a finite word at selected subtree positions, multiplied by a scalar amplitude of some finite upper rate. A family is eligible at exponent b when each alternative's positions are separated at b. Two families agree to rate ρ, up to exterior weight W, when their complete correlations differ with rate at most ρ against every fixed exterior label list of total degree at most W, with arbitrary admissible position paths and arbitrary viewer paths. A supply at parent exponent k and cost M gives such an eligible family for every sufficiently large exterior exponent b; for every W, ρ it gives a comparable replacement agreeing in this sense. These choices concern exact continuum correlations throughout.

To change one branch while retaining all other branches, fix an alternative of the other branches. Let its field degree sum be d and its amplitude have rate ζ. Include its fields in the exterior test of the changed branch. Request agreement through weight $W + d$ and rate $\rho - \zeta$. Multiplying by that alternative's amplitude returns rate ρ. The finite sum over all other alternatives has the same rate. Since there are finitely many alternatives, one may take the maximum of their $W + d$ and ζ first. Sibling separation and the old eligibility show that all these augmented exterior tests have the required separation. This proves the single-branch replacement formula. Induction over the finite branch set proves the triangular replacement assertion: at a given replacement only the previous choices and unchanged later branches enter its finite maxima. Transitivity follows from adding the complete correlation differences.

A tested agreement transfers to coefficient agreement by the actual finite reader. For a requested horizon m, take W at least the maximum degree sum of every exterior word in all detector supports through that horizon. These supports form a finite union. The rotating detector words retain their degree sums and have bounded label coordinates; their positions stay a fixed positive distance from the center and from one another. A shrinking eligible family is therefore admissible against every such moving detector eventually. Write each rotated reader as its finite sum of these tested correlations times its fixed weights and rotated loaded vectors. Apply the agreement estimate and then average over the compact rotation group. The arbitrary-path quantifier makes the bound uniform in the rotating group element, as proved for (OE4). Apply the real symmetrization and K_m, then a basis coordinate; these fixed finite-dimensional linear maps preserve its rate. This proves the coefficient transfer with the same ρ.

There is a further debit when transferring a whole remainder, which is why agreement of products alone would not suffice. Let the desired remainder rate be

$$\rho_0 = M - (m + 1)(k - u),$$

before exterior cost, where the center-to-exterior exponent is $0 \le u \le k$. Choose a comparable replacement agreeing at rate $\rho_0 - mu$, with enough exterior degree to cover both the actual exterior fields and the finite detectors through m. Its full-product difference has at most this rate, hence at most ρ_0. The coefficient differences through degree $j \le m$ have rate $\rho_0 - mu$; their centered exterior correlations cost $ju + |e|b$. Each resulting term therefore has rate at most $\rho_0 + |e|b$. Add the comparable replacement's own complete remainder at that rate and subtract these coefficient differences. The exact resulting expression is the original family's complete remainder. This proves the second inductive supply assertion with all coefficient subtractions retained.

The comparable-family estimate used above includes the boundary exponent cases. An alternative of field degree d has amplitude of rate $M - dk$, and its positions have size and separation exponent k. For $k > 0$, its degree-m coefficient costs $(d - m)k$ by the homogeneous reader bound. Multiplication gives $M - mk$. If $k = 0$, test its amplitude times the word against the finite detector exteriors. The word has rate zero, the amplitude rate M; the readout comparison with the empty family gives rate M for its coefficient as well. Actual shrinking of its positions ensures eventual reader admissibility even though its exponent is zero. For the full remainder, when the exterior exponent $u < k$, multiply the single-center remainder rate by the amplitude to obtain

$$M - (m + 1)(k - u) + |e|b.$$

If $u = k$, estimate the whole correlation and each finite centered term separately; the engineering degrees cancel and give exactly $M + |e|b$, the same expression at zero gain. Summation over the finite comparable family proves the estimate in every case.

We now remove the auxiliary assumption that radii already possess logarithmic rates. Along any approaching admissible family all internal radii tend to zero, since they are positive and bounded above by the root radius. Put

$$T = \sum_{v \text{ internal}} \log(1/s_v).$$

Eventually every summand is nonnegative, $T > 0$, and $T \to \infty$. The vector with coordinates $\log(1/s_v)/T$ belongs to the finite compact cube $[0,1]^{\mathrm{Inner}(t)}$. On a nontrivial refining filter it converges to a; summing the coordinates proves $\sum_v a_v = 1$. Hence each s_v has rate $-a_v$ on that filter. Monotonicity $s_v \le s_{\mathrm{parent}(v)}$ implies $a_v \ge a_{\mathrm{parent}(v)}$, since the opposite strict inequality would eventually force the smaller radius to exceed the larger one by the exponential-rate comparison. Take the virtual parent exponent to be zero.

The geometric size estimate Ms_v gives the required envelope rate at vertex v, and also actual convergence of every descendant displacement to zero. The sibling lower bound ηs_v gives separation at every exponent strictly larger than a_v. Induction over the tree thus gives all the regularity hypotheses of the preceding rate proof, including vertices with $a_v = 0$. Since $\sum a_v = 1$, the slack $\epsilon \sum \Delta_v$ is strictly positive, as required in (NE4). Applying the rate estimates gives (NE2) and (NE3) eventually on this refining filter.

If a single common small-radius threshold for all admissible scales and positions failed, choose a violating scale and position array at arbitrarily small radii and restrict to those radii. The compact layer extraction just proved gives a nontrivial refining filter of this very array where the desired inequality holds eventually, contradicting its selected violations. The same argument handles the exterior geometry and all viewers; for viewers one first selects a violating viewer at each parameter. Thus the final thresholds depend only on the fixed tree, field labels, horizons, M, η, ϵ, and exterior clearance δ, with no lower bounds on the ratios. This proves the uniform quantifiers in (NE2)–(NE3).

The bounds admit arbitrary joint weighted integration. Let (K, μ) be any measure space, $u \in L^1(\mu; \mathbb{C})$, and $\int |u|\, d\mu \le B$. Let $z : K \to \mathrm{Plot}(t)$ be almost everywhere measurable, with admissible geometry wherever $u \neq 0$, at one fixed array of radii s. The direct coefficient and the finite iterated coefficient are smooth on their common open reader domain. Extending them by zero outside that domain gives measurable functions. If $|g(z(i))| \le C$ on the support of u, then

$$|u(i)g(z(i))| \le C|u(i)|$$

almost everywhere, so this product is integrable and

$$\left| \int u(i)g(z(i))\, d\mu(i) \right| \le BC. \tag{NE7}$$

Apply this first to the direct coefficient with $C = \mathcal{T}_{m,\epsilon}$, then to its difference from the iterated coefficient with $C = \mathcal{T}_{m,\epsilon} \sum_{v \neq \mathrm{root}} q_v^{D_v+1}$. The iterated coefficient is integrable because it is the difference of those two integrable quantities. Subtracting their integrals proves the weighted version of (NE2). No product structure or absolute continuity assumption on μ is used.

For jointly varying exterior positions $y(i)$ and phase characters $\chi(i)$, the same conclusion applies to the complete tested remainder with $C = \mathcal{T}_{m+1,\epsilon}$. Its measurability follows on the common open guard domain from smoothness in positions and continuity of the character evaluation of each fixed ground-algebra element. The character space is compact metrizable, and these separate properties give joint measurability. Almost everywhere measurable z, y, χ may therefore be composed with it. The full raw correlation is integrable too: its bound is $\mathcal{T}_{0,\epsilon}$, obtained by adding the finite degree-zero terms to the degree-zero remainder and applying the engineering rates. The truncated expression

is then the difference of two integrable expressions. Consequently the finite sum may be integrated term by term and the norm of the full integrated discrepancy is at most $B\mathcal{T}_{m+1,\epsilon}$. A probability mixture of characters is the special case $u = 1$, $B = 1$; its expansion uses the identical coefficients outside the phase integral.

Finally, covariance concerns the full homogeneous vectors, rather than individual scalar basis channels. At a leaf the emitted vector is its actual label. At an internal vertex first form the emitted vectors of the children, include each in its fixed horizon space, and apply the multilinear h_j to their centers for each j through the parent's horizon. The finite sum of these included outputs is the emitted vector of that subtree. Expanding multilinearity shows that its coordinates are exactly the finite iterated channel formula. Translation invariance, conjugation compatibility and orthogonal covariance now follow by induction from the corresponding laws of h_j and the inclusions. Applying just h_m at the root proves the same laws for the degree-m iterated output. The basis can transform nontrivially within H_m; no assertion that its vectors are individually rotation invariant is needed. The full tested discrepancy inherits these laws from the hierarchy and these homogeneous outputs.

The separated and nested expansions, their compatible coefficients, all-viewer estimates, joint weighted versions and covariance have now been translated.

14 Extension across collision sets

Extending a separated correlation across collisions requires a choice of subtraction device. Taylor jets normal to a collision set are subtracted on dyadic shells; the cancellation-bearing shell expressions then have an absolutely summable L^1 bound. Forests organize overlapping collisions with finitely many branches at a fixed degree and arity. The construction proves compatibility for the chosen devices and finite cutoffs; independence of arbitrary subtraction prescriptions is not an additional conclusion.

The additional contact conclusion concerns Euclidean distributions, constructed with a specified subtraction system using the same reader prescription. Its scope is the following. Restrict the k-point configuration space by a real linear isometry $P : E \to (\mathbb{R}^4)^k$, where E is finite-dimensional and carries its Haar volume. A collision group specifies a subset of the k labels and its collapsed configurations. The permitted forest consists of compatible collision groups together with local continuation charts, support windows, finite Taylor orders and projection choices. On every admissible such forest, the construction supplies an integrable Taylor-subtracted pairing for a compactly supported smooth test function in the chart. Its expansion into the prescribed annular pieces is absolutely summable in integral norm and sums to that pairing. Iterated expansions over any finite ordered list of centered collision groups have the same property. On annuli away from total collapse the result agrees with the punctured correlation.

The public endpoint also bounds the resulting distribution on every permitted isometric collision plane containing the collision directions, at the stated finite total test order. The subtraction devices carry the actual rotation action of the reconstructed hierarchy, and their finite projection basis is exactly the reader prescription used for the OPE. A real system basis and compatible positive threshold choices exist simultaneously. These contact assertions must therefore be proved with these specified devices and admissible forest conditions; they are not a claim of arbitrary subtraction-independent contact values.

For a nonempty label set g, its normal displacement space in the full configuration space is

$$N_g^{\mathrm{full}} = \{x : x_j = 0 \ (j \notin g), \ \sum_{j \in g} x_j = 0\}.$$

Its orthogonal complement contains the collapsed configurations of the group, including its centroid. In E use $N_g = P^{-1}N_g^{\mathrm{full}}$. The admissibility assumption is $N_g^{\mathrm{full}} \subset \operatorname{ran} P$; hence P identifies these normal spaces isometrically. Omitting any one member of g identifies N_g^{full} linearly with $(\mathbb{R}^4)^{|g|-1}$: the omitted coordinate is minus the sum of the others. In particular

$$\dim N_g = 4(|g| - 1) =: \nu_g. \tag{CE1}$$

For disjoint groups the normal spaces are orthogonal. Centroids of two disjoint nonempty proper groups are different almost everywhere in these normal variables: their equality is the kernel of a nonzero linear map. It is nonzero by placing a vector at one member of each group with opposite signs. Its proper linear kernel has Haar measure zero. This permits the subsequent almost everywhere shell identities without assigning a correlation directly on a coincidence plane.

Write π_g for orthogonal projection onto N_g. The degree-j normal Taylor term of a smooth test $f : E \to \mathbb{C}$ is

$$T_{g,j} f(x) = \frac{1}{j!} D^j f(x - \pi_g x)[\pi_g x, \dots, \pi_g x].$$

For a fixed positive subtraction length s, choose the prescribed smooth radial cutoff $\chi_s(\pi_g x)$, equal to one near zero and zero for $\|\pi_g x\| \geq 2s$. The subtraction operator through order m is

$$\mathsf{T}_{g,s,m} f = \chi_s(\pi_g x) \sum_{j=0}^{m} T_{g,j} f, \qquad \mathsf{R}_{g,s,m} = 1 - \mathsf{T}_{g,s,m}. \tag{CE2}$$

Tangential coordinates remain parameters in this Taylor expansion. The subtraction depends only on the normal jets through m along $N_g^{\perp}$, and vanishes on a test supported away from that subspace. Its support is contained in

$$\{x : x - \pi_g x \in \operatorname{supp} f,\ \|\pi_g x\| \leq 2s\}.$$

These facts follow directly from the displayed derivative formula and the cutoff. Orthogonal groups have commuting Taylor operators: their normal projections annihilate one another, their normal derivatives commute, and their cutoff factors depend on different normal variables. Thus their remainders commute as well.

The finite field projection is combined with (CE2), rather than subtracting every component equally. For input degrees d_j, set

$$\omega_g = \sum_{j \in g} d_j, \qquad L_g = \max(0, \omega_g + 1 - \nu_g).$$

Resolve the group's field expression into projected components of degrees $h = 0, \dots, L_g - 1$ and its remaining component. The remaining component is paired with the original test; the degree-h projected component is paired with the normal Taylor remainder through order $L_g - 1 - h$. Denote these finitely many choices by v, with gain $G_v = L_g$ for the remaining component and $G_v = h$ for a projected component. After normal integration the effective scale exponent is

$$e_v = \nu_g + G_v - \omega_g. \tag{CE3}$$

For the remaining component $e_v \geq 1$, including when $L_g = 0$. For a projected component, $h < L_g$ implies

$$e_v + (L_g - 1 - h) = 0.$$

At a relative shell radius $a \leq 1$, Taylor's remainder contributes $a^{L_g - h}$, leaving one full power:

$$a^{e_v} a^{L_g - h} = a. \tag{CE4}$$

For $a > 1$ the cutoff estimate needs only a^{L_g-1-h}, so the projected product is exactly 1. The remaining component can meet such a shell only up to the fixed geometric reserve, where a^{e_v} is bounded. This near/far distinction is the central integrability gain.

To make this argument uniform under further collisions, a local chart records a partition Q of the current label group into disjoint nonempty parts, a positive length s, a collapsed offset, a linear resizing frame A, and a closed stage inside a fixed bounded set in resized coordinates. Each part has an assigned radius $R_i \in (0, 1]$; the frame rescales its normal directions by that radius. The whole normal extent and its subtraction halo fit a common detector collar $B \in (0, 1]$, with $2s \le B$. For every nonsingleton part, R_i is at most a fixed multiple of s. The detector requirement is

$$B \le \text{clearance}\left(\sum_j d_j\right).$$

The clearance is chosen decreasing with degree, so it also covers every projection degree used below. Other groups, their centers and their field expressions are read by a continuation constant along each current part's normal directions; this permits use of ancestor-subtracted expressions as the exterior continuation, with their cancellations retained.

Refine one nonsingleton part i into a proper partition U on an angular annular patch. Assign radius a to the new parts and leave the others unchanged. Replace the frame by its composition with the map multiplying N_i by a/R_i and fixing $N_i^{\perp}$. The determinant changes by

$$|\det A'| = |\det A|(a/R_i)^{\nu_i}, \tag{CE5}$$

and the engineering field factor changes by $(a/R_i)^{-\omega_i}$. The gain in the continuation contributes $(a/R_i)^{G_v}$, so their product is precisely $(a/R_i)^{e_v}$. Supports in the Taylor branch may enter its subtraction halo; this is why the refined route retains both the old stage and its normally collapsed, cutoff-bounded part. Orthogonality keeps the other groups' normal data unchanged. Their continuation remains locally continuous because the Taylor operation is applied within its detector collar and the refined angular patch keeps the distinct part palettes disjoint.

Here is the uniform integrated estimate. Let

$$\tau_Q = 3^{\,|g|-|Q|}, \quad \varepsilon_Q = \tau_Q\gamma, \quad \gamma > 0, \quad (\tau_Q + 1)\gamma \le 1.$$

A proper refinement increases the number of parts, so

$$2\varepsilon_{Q'} + \gamma \le \varepsilon_Q. \tag{CE6}$$

The local base budget for degree list d, continuation bound B_w, frame A_w, and radii $R_j(w)$ is

$$\mathcal{B}_w(\varepsilon) = |\det A_w|\, B_w \prod_{i\in Q} R_i(w)^{-\omega_i} \prod_j R_j(w)^{-\varepsilon}. \tag{CE7}$$

The last product includes all recorded radius registers, including unchanged ones. If selected parts are restricted to dyadic shell numbers n_i, multiply this budget by $\prod_{i\in Q,\ |i|>1} 2^{-n_i\gamma}$; an unrestricted part contributes 1. The number of required test derivatives is

$$N_Q = (2 + \sum_j d_j)(|g| - |Q|).$$

For every permitted choice of projected or remaining components and shell restrictions, the complete resulting integrand F is integrable and

$$\int_E \|F(x)\|\, dx \le C\, \|f \circ A_w\|_{C^{N_Q}}\, \mathcal{B}_w(\varepsilon_Q) \prod_{\substack{i\in Q\\ |i|>1}} 2^{-n_i\gamma}. \tag{CE8}$$

The constant is uniform in the permitted charts and parameters, after fixing the geometric reserves, degree bounds, and continuation estimates. In (CE8) the absolute value is taken after the finite subtraction expressions have been combined.

The induction proving (CE8) uses (CE3)–(CE6). Under refinement the child's budget is the parent's budget at $\varepsilon_{Q'}$, times $(a/R_i)^{e_v}a^{-\varepsilon_{Q'}}$. The Taylor estimate supplies the extra factor from (CE4). For every $0 \le t \le 1$ their product is at most a fixed reserve constant times $(a/R_i)^t$: use (CE4) when $a/R_i \le 1$, and the bounded far-shell estimate otherwise. Choose $t = \gamma + \varepsilon_{Q'} \le 1$. Then

$$(a/R_i)^{\gamma+\varepsilon_{Q'}}a^{-\varepsilon_{Q'}} = a^{\gamma}R_i^{-(\gamma+\varepsilon_{Q'})}.$$

The last R_i loss is absorbed by the parent's radius surcharge using (CE6). The other shell factors are unchanged. Thus each new dyadic annulus supplies a^{γ}, and summation over its annulus index converges. A proper partition strictly decreases $|g| - |Q|$; this supplies the finite induction and the stated derivative budget.

Here are the shell and angular constructions used in the induction. Put

$$\psi_{g,a} = \chi_a(\pi_g x) - \chi_{a/2}(\pi_g x), \qquad a_n = 2^{-n}.$$

The support of $\psi_{g,a}$ lies in $a/2 \le \|\pi_g x\| \le 2a$, and

$$\sum_{n<N} \psi_{g,a_n} = \chi_1 - \chi_{a_N}. \tag{CE9}$$

For a test supported where $\|\pi_g x\| \le 1$, at every point with $\pi_g x \neq 0$ the right side times the test equals the test in a neighborhood of that point for all sufficiently large N. This is an eventual equality of germs, including all derivatives that enter the finite Taylor operations.

On the compact unit normal annulus, partition the labels by equality of their positions. Total collapse is absent, so this is a proper partition. Within each part all subcentroids coincide, and subcentroids belonging to different parts are distinct. The latter condition is open and there are only finitely many proper partitions and subgroups. These open sets therefore cover the annulus. Choose a smooth finite partition of unity subordinate to this cover. Multiplying its rescaled angular factors by $\psi_{g,a}$ gives masks $\psi_{g,a,U}$ whose sum over proper partitions U is $\psi_{g,a}$. On each mask support, all subcentroids in different parts have distance at least $c_U a$, with $c_U > 0$: at radius one take the positive minimum of their continuous distances on the compact support, then rescale. This proves all the reader separations needed after the refinement, even for subgroups used later.

The normal Taylor bound in these coordinates is uniform in the subtraction length. On a fixed compact unit annulus, if f has bounded derivatives through $n+m+1$, normal rescaling by a gives

$$\|\mathsf{R}_{g,s,m}f\|_{C^n \text{ after rescaling}} \le C\|f\|_{C^{n+m+1}} \begin{cases} a^{m+1}, & 0 < a \le 1, \\ a^m, & a > 1. \end{cases} \tag{CE10}$$

Here the parent frame has normal scale R, and the allowed parameters satisfy $R \le Bs$, $Ra \le Bs$, with the fixed normalized support bound. To prove (CE10), bound the family of cutoff functions $\chi_1(u\pi_g x)$, $0 \le u \le B$, and all derivatives through n on the compact annulus; smoothness on the compact product gives one constant. The normal Taylor polynomial after rescaling is bounded by a constant times $\max(1,a)^m$. The rescaled original function has a uniform derivative bound when it meets its compact support, and is zero as a germ when a is beyond that support. Leibniz's rule gives the a^m bound for $a > 1$. For $a \le 1$, if $3Ra/s \le 1$, the cutoff is identically one near the annulus and the integral Taylor remainder gives the factor a^{m+1}. Otherwise $1 \le 3Ba$; multiplying the rough

uniform bound by $(3Ba)^{m+1}$ gives the same a^{m+1} estimate with a larger constant. The angular masks have fixed normalized derivative bounds, so their product costs another fixed constant. This proves (CE10).

We may now make the induction in (CE8) completely explicit. If every current group is a singleton there is no unresolved normal direction; use the permitted continuation bound, its local continuity and the bounded support in resized coordinates. Otherwise choose a nonsingleton part i, and decompose its election into the finite choices of remaining component and projected degrees. For a fixed such choice v, apply (CE9) to its test after the normal Taylor operation and then the finite angular partition. The resulting annular term equals the refined forest expression, read with the continuation that first regroups the children into i and applies its fixed component projection or remainder. This is an identity of the underlying linear expression before evaluating any hierarchy or integrating.

The refined partition Q' has more parts, so its value of $|g| - |Q'|$ is strictly smaller. Its continuation retains the original offsets, length and normal constancy. Its component fields have the prescribed palette of subcentroids; the angular separation and common detector collar ensure the projection is admissible. The inductive bound applies. Formula (CE10) pays its test derivatives, while (CE3)–(CE7) give precisely the budget

$$C_v \|f \circ A_w\|_{C^{N_Q}} \mathcal{B}_w(\varepsilon_Q) a_n^{\gamma} \prod_{\substack{j \in Q \\ j \neq i}} 2^{-n_j \gamma}. \tag{CE11}$$

The number of angular sectors and component choices is finite; summing their constants preserves uniformity. If the specified annular stage is empty, the term is zero and the same estimate holds. Otherwise its geometry is exactly the refined route already described.

Since $\sum_n a_n^{\gamma} = (1 - 2^{-\gamma})^{-1} < \infty$, (CE11) proves summability of the integral norms of the complete annular terms. The eventual germ equality from (CE9), together with the proper-normal exceptional sets of Haar measure zero, identifies the pointwise sum almost everywhere with the unsliced component expression. To pass to its integral without any unproved domination, regard each annular term as an element of the Banach space $L^1(E; X)$. The sum of their L^1 norms is finite, so they sum in L^1 to an integrable function. Absolute integral-norm summability also gives their almost everywhere pointwise sum; uniqueness and the eventual identity identify that function with the prescribed unsliced expression. Continuous integration then gives

$$\int F_v = \sum_{n=0}^{\infty} \int F_{v,n}, \qquad \sum_n \int \|F_{v,n}\| < \infty. \tag{CE12}$$

This proves the bound for the unsliced choice. For a specified shell index use (CE11) directly. Summing the finite component choices establishes (CE8), completing the finite induction.

If the election at i is the whole expression, all its choices occur, so (CE12) yields the exact root formula

$$\int F = \sum_v \sum_{n=0}^{\infty} \int F_{v,n}. \tag{CE13}$$

For any ordered list of distinct nonsingleton current groups with no preassigned shells, one can resolve their shells successively. The common bound becomes a constant times $\prod_j 2^{-n_j \gamma}$. This is summable over the Cartesian product of all shell indices. Repeated use of (CE12) and absolute Fubini therefore gives the full iterated sum, equal to the original integral. In particular every permitted resolution order gives that same pairing. This does not assert absolute convergence of

the individual uncombined Taylor monomials; the terms in (CE12) and (CE13) are the already combined expressions.

For any requested loss $\epsilon > 0$, choose $\gamma = \min(1, \epsilon)/(\tau_Q + 1)$. Then the economy condition holds and $\varepsilon_Q \le \epsilon$. On unrestricted shells the shell factor is one, and monotonicity of $R^{-\epsilon}$ for $0 < R \le 1$ upgrades (CE8) to the requested arbitrary positive loss. Thus all shell identities and bounds apply with the same finite projection choices and devices.

We verify the continuation inheritance required in that induction. Its hypothesis is a tested rate inequality on every logarithmic path through the allowed chart: if the scalar continuation price has rate b, and each current group expression is affordable at its assigned radius with rate c_i, the reading of their tagged product has rate at most $b + \sum_i c_i$. The continuation is also locally continuous on the specified palettes and constant in the normal directions already assigned to each group. These properties hold for the actual hierarchy with separated exterior fields by the signed expression estimates described below.

For a refinement of i at radius a, suppose a has rate $-s$ and its parent's recorded radius has rate $-t$. The angular separation gives positive lower bounds of rate $-s$ between all palettes of different child parts. The finite product rule for affordable expressions therefore gives the product of the children rate $\sum c_j$ on the parent palette at exponent s. Applying the remaining component or degree-h projection at the parent's center changes this to

$$\sum_j c_j + (t - s)G_v. \tag{CE14}$$

For the remaining component the chart's support forces $t \le s$; use the full finite projection remainder estimate through $L_i - 1$, giving $G_v = L_i$. If $L_i = 0$, there is no projection and the same estimate is the coarsening inequality with zero gain. For a projected component, the degree-h readout estimate supplies $G_v = h$. The common detector collar justifies these operations on the actual moving palettes.

The new scalar price is exactly

$$B'_w = B_w(a/R_i)^{G_v}.$$

If B_w has rate b, this has rate $b + (t - s)G_v$, matching (CE14). On the subtraction halo, choose a point of the original stage differing from the present point only in N_i. Such a point exists by the definition of the halo; for the remaining branch it is the point itself. Constancy of the original continuation in N_i permits evaluating there. Other current parts are orthogonal to N_i, so their palettes and expressions are unaffected. Substituting the child product with its fixed component treatment into the old tagged product is exactly the refined reading. The original path inequality now yields

$$b + (t - s)G_v + \sum_{\text{refined parts}} c_j,$$

the required refined continuation bound. This proves the inheritance used in (CE11). The local continuity assertion follows by the same substitution on two closed pieces, the original-stage piece and the normally collapsed-stage piece of the halo, and gluing their agreeing restrictions. Thus the auxiliary continuation is justified at each step, including expressions retaining earlier Taylor cancellations.

We next obtain a distribution on the whole allowed position space. At each $x \in E$, partition the labels by equality of the actual positions $Px+o$. Each part is collapsed at x; its normal displacement is zero. All subcentroids of different parts are distinct. These conditions persist on a sufficiently small closed ball about x: all part normal displacements remain inside the common collar B, and all distinct-part subcentroid distances have a positive lower bound. Crucially the same positive

subtraction length s, with $2s \le B$, works on each of these balls. Applying the local forest estimate to the equality partition gives its integrable test pairing. The exact expression overlap identities identify that pairing with the single root prescription on tests supported in the ball, so no new extension values are assigned on chart overlaps.

For a compact test support K, choose finitely many such balls covering it and a smooth partition of unity $\{\theta_j\}$ subordinate to their interiors. On every chart the test estimate uses at most

$$N = (2 + \sum_i d_i)|g|$$

derivatives, since $N_Q \le N$. Leibniz's rule gives $\|\theta_j f\|_{C^N} \le C_{K,j}\|f\|_{C^N}$. The root prescription is linear in its test, and $\sum_j \theta_j f = f$. Add the finitely many local integral-norm bounds to obtain

$$\int_E |K_s(f)(x)|\,dx \le C_K\|f\|_{C^N} \quad (\operatorname{supp} f \subset K). \tag{CE15}$$

The linear functional

$$\mathfrak{S}_s(f) = \int_E K_s(f)(x)\,dx$$

is therefore continuous on each fixed-compact-support smooth test space, and hence on the usual compactly supported test-function space. This is the asserted Euclidean distribution of finite order N. Its restriction to a local forest is the local pairing, because both are readings of the same test expression. The construction works for every unselected or character reader. The extension and its bounds retain the chosen subtraction length, finite readouts and normal Taylor convention.

We specify the finite expression prescription, including its recursion. A score is a finite sum of field expressions tensored with linear operations on smooth tests; its action at x multiplies the field expression at x by the operated test at x. Products compose test instructions in their prescribed order. For disjoint label groups those instructions commute, by orthogonality of their normal spaces.

Let C_g denote the recursively constructed local correction, and let $\mathcal{P}_g$ be the punctured score formed from proper partitions:

$$\mathcal{P}_g = \sum_{\substack{Q \text{ partition of } g \\ Q \neq \{g\}}} \prod_{i \in Q} C_i.$$

For a singleton, $C_{\{j\}}$ is its actual field expression tensored with the identity test operator. For $|g| \ge 2$, set

$$C_g = -\sum_{h=0}^{L_g - 1} \operatorname{reshape}(\operatorname{reader}_{g,h})\mathcal{P}_g\ \mathsf{T}_{g,s,L_g-1-h}. \tag{CE16}$$

The full score is the sum over all partitions, or equivalently $\mathcal{S}_g = \mathcal{P}_g + C_g$. Every group appearing in a proper partition has strictly smaller cardinality, so (CE16) defines a finite recursion. The reader in this formula is the fixed physical projection at the group centroid; it acts on the field expression, while the right Taylor factor acts on the test. For $L_g = 0$ the sum is empty and there is no local correction. Formula (CE16) is the exact prescription whose finite rearrangement gave the component decomposition preceding (CE3).

Every C_g depends on the test only through its germ at the normally collapsed point $x - \pi_g x$. This follows from (CE2) in each summand of (CE16); for a singleton its normal space is zero. Now compare two partitions Q, R which are both separated on a test support. Expanding $\prod_{i \in Q} \mathcal{S}_i$ gives precisely the terms $\prod_{j \in T} C_j$ over all partitions T refining Q: choosing a partition separately inside every part of Q is bijective with choosing a refinement. The same holds for R. A term refining Q

but not R has a group j containing labels in two different parts of R. At its normal collapse those labels coincide, whereas they are separated everywhere on the test support. Thus the test vanishes in a neighborhood of this collapsed locus, and its germ-dependent factor C_j annihilates it. All such crossing terms vanish as scores after multiplication by the test mask. Both expansions therefore reduce to the common refinements of Q and R. Acting on the test proves their equality before hierarchy evaluation and integration. This proves the overlap identity used in (CE15), including partitions not subordinate to one another.

The signed expression bounds supplying the initial continuation estimates follow a finite proof tree. Its leaves are bounded labels of degree d at one position, with price R^{-d}, and the unit with price 1. One may multiply two already priced expressions when their position palettes are separated by a positive multiple of their common radius; their prices multiply. One may carry an expression from radius R to R' when its palette has diameter $O(R)$ about its chosen center and $R = O(R')$. Subtracting its full projection through degree h multiplies its price by $(R/R')^{h+1}$, with the same size condition and the fixed reader clearance. Projecting onto just degree h multiplies its price by $(R/R')^h$ and replaces its palette by the center alone; no ratio bound is required for this projected operation. Palette enlargement is also permitted.

Along any logarithmic path where the finitely many radii have rates, induction on this tree proves the corresponding affordable-expression rate. The leaf uses the full field bound, multiplication uses separated products, carrying uses coarsening, and the two projection steps use the full tested remainder and output estimates with the displayed gains. All use the same hierarchy and reader, and the remaining component is kept as its signed difference. To obtain a uniform bound in the radii, normalize their logarithms as in the nested expansion proof and pass to a compact refining limit. The product surcharge $\prod R_j^{-\epsilon}$ supplies a strict positive rate margin against any proposed violation. Hence an expression of price p satisfies

$$|\langle E\rangle_\chi| \le C_\epsilon p \prod_j R_j^{-\epsilon}$$

for all admissible geometries and all viewers. Testing its output at degree h gives the additional factor R^h. This proves the signed moment and projection bounds for every finite expression tree. It also supplies the initial continuation bound with separated exterior words: their palettes have fixed positive clearance, their label norms are bounded, and their radius can be set to one. Their full expectations and local continuity are those of the previously constructed hierarchy.

The same construction permits an extra projection or remainder at the root whenever its full palette lies in the required root detector collar. Its scalar continuation price is R^h for degree h output and R^{h+1} for the remainder through degree h. These operations are applied after the root's field expression has been formed, and before the lower forest is tested; (CE14) shows that the subsequent shell induction keeps this root price intact.

Consider now normalized normal tests

$$f_R(x) = R^{-\nu_g} f\big(x - \pi_g x + R^{-1}\pi_g x\big). \tag{CE17}$$

The normal frame is multiplication by R on N_g and identity on its complement. Its determinant is R^{ν_g}; composing f_R with that frame gives $R^{-\nu_g} f$, so the determinant and test normalization cancel exactly. In a one-root chart, (CE7) consequently becomes

$$R^{u-\omega_g-\epsilon},$$

where $u = 0$ for the full reading, $u = h$ for degree-h root output, and $u = h+1$ for its root remainder. Thus each complete contact pairing obeys

$$\int \|K_s(f_R)(x)\|\,dx \le C\|f\|_{C^{N_{\{g\}}}} R^{u-\omega_g-\epsilon} \tag{CE18}$$

on the permitted normalized supports, uniformly over the chart parameters. The subtraction length may be fixed or varying, subject to the same support, halo and clearance conditions; it need not equal R.

If f vanishes near the total diagonal $N_g^{\perp}$, every term C_g in (CE16) annihilates it, by the normal jet formula. Therefore the full score agrees with the punctured score on such tests, even when they meet proper-subcluster collision sets. A test supported in any positive normal annulus has this property. Normal rescaling preserves it, so (CE18) holds for the punctured pairing with the same constants. No uniform positive lower bound on the annulus hole is required. These are the full, output and remainder scaling assertions and their punctured counterparts.

A real basis can be chosen in each homogeneous contact grade simultaneously. This is a coordinate basis for the strata of the already fixed splitting; it does not replace the OPE prescription or its projection functionals. The splitting's stratum is invariant under conjugation and rotations, because its lower-grade map intertwines them. Every vector is a complex combination of conjugation-fixed vectors by the real-and-imaginary decomposition already used for H_m. Extend a linearly independent subset of these real vectors to a maximal such subset; it spans the complex stratum, so finite-dimensionality makes it a finite real basis over $\mathbb{C}$. At degree zero start with the nonzero unit, which is fixed by conjugation, and extend it. This gives a homogeneous system basis containing the unit. Expand a general input field into its finitely many homogeneous parts and then into this basis. Multilinearity of the score and reader gives the general contact expression as the finite sum of its basis expressions. Each term retains its actual assigned homogeneous degree in (CE16). Raising the ambient filtration horizon changes none of these parts, so the resulting linear extension is well defined.

The compatible radius-dependent prescription is also explicit. Given any positive requested bounds f_n, choose

$$g_n = \min\Big(f_n,\, 2^{-n},\, \frac{\mathrm{clearance}(n)}{2(n+1)}\Big), \qquad t_0 = g_0, \quad t_{n+1} = \min(t_n/8, g_{n+1}).$$

Then $t_n > 0$, $4t_{n+1} < t_n$, $t_n \to 0$, and $2(n+1)t_n \le \mathrm{clearance}(n)$. Let $b_0(s) = 1$ and, for $n > 0$, let $b_n(s)$ be the prescribed smooth cutoff at t_n. It is one for $0 \le s \le t_n$ and zero for $s \ge 2t_n$. For every $s > 0$ these cutoffs are locally finite in n. Moreover, a nonzero $b_k(s)$, $k > n$, implies $b_n(s) = 1$, by the strict threshold spacing.

Let $N(s)$ be the last nonzero cutoff index. It exists, and $N(s) \to \infty$ as $s \downarrow 0$. For every active $k > 0$,

$$ks \le \mathrm{clearance}(k).$$

Indeed $s < 2t_k$ and the displayed threshold bound is stronger. Consequently for any fixed C, eventually $Cs \le \mathrm{clearance}(N(s))$. Define the finite smooth projection

$$\mathcal{D}_s(x)E = \sum_{j=0}^{N(s)} b_j(s)\ \mathrm{insertRow}_{x,j}\, E. \tag{CE19}$$

Local finiteness makes this smooth for positive s, even where the last active index changes. Biorthogonality of the homogeneous rows gives exactly

$$\mathrm{row}_{x,n}\, \mathcal{D}_s(x)E = b_n(s)\, \mathrm{row}_{x,n}\, E.$$

For $s \le t_n$, all rows through n equal one, and projecting (CE19) through degree n is exactly the original degree-$\le n$ projection. Thus no limiting approximation is substituted for a finite reader in

the contact estimates. The construction may be slowed to meet all additional positive f_n bounds while preserving these identities. It asserts no convergence as the growing cutoff horizon itself varies.

This completes the contact endpoint's real basis and threshold existence clauses. The affordable-expression rules, the uniform logarithmic compactness argument, and the mixed-label inclusion identities are proved in (UC1)–(UC3) below.

14.1 Uniform subpower bounds on contact layers

The layer estimates used in the contact construction are uniform in all radii. Their deduction from logarithmic-path estimates is the following compactness argument. Let $R_1, \dots, R_m$ be radii in $(0, 1]$ on a configuration set U, let $q > 0$ be the bill without the small losses, and let F be the normed quantity to estimate. Suppose that along every nontrivial filtered sequence with an increasing clock $T \to \infty$ and

$$\frac{\log R_a}{T} \longrightarrow -s_a, \qquad 0 \le s_a \le 1,$$

the bill has some exact exponential rate c, and F has upper rate at most c. Then for every $\varepsilon > 0$ there is a single finite $C > 0$ such that

$$\|F(w)\| \le C q(w) \prod_{a=1}^{m} R_a(w)^{-\varepsilon} \quad (w \in U). \tag{UC1}$$

If this failed, choose w_n violating the inequality with $C = e^n$. Put

$$p_{n,a} = -\log R_a(w_n) \ge 0, \quad S_n = \sum_a p_{n,a}, \quad T_n = n + S_n.$$

Take an ultrafilter refining the sequence tail. Compactness of $[0, 1]$, for finitely many coordinates, gives limits $p_{n,a}/T_n \to s_a \in [0, 1]$ and $n/T_n \to k \in [0, 1]$, with $k + \sum_a s_a = 1$. The path hypothesis applies. The supposed upper bound for F has rate c, while the violating comparison quantity

$$e^n q(w_n) \prod_a R_a(w_n)^{-\varepsilon}$$

has exact rate $c + k + \varepsilon \sum_a s_a > c$. Positivity of the increment follows because k and all s_a are nonnegative and their sum is one. A quantity of strictly larger exact rate eventually exceeds one of smaller upper rate, contradicting every chosen violation. The added term n in T_n handles sequences whose radii remain away from zero.

Every bill used here has the required exact rate: constants, positive radii, products and fixed real powers have exact rates, and these operations generate the finite radius-power bills. Thus the argument does not require a common OPE threshold for all multiparameter paths or a common intermediate output horizon. It supplies the ordinary uniform layer inequality actually used in the contact integral.

14.2 Extending contact expressions from homogeneous to arbitrary labels

Let $M_a^{(d)} : E_{\le d} \to E_{\le a}$, $0 \le a \le d$, be the fixed homogeneous splitting maps, with their inclusion maps between grade caps. The contact route construction for homogeneous labels is multilinear in each input. To see this before integration, fix all other inputs; its grade-transport identity expresses the remaining dependence as a linear map. Hence addition and scalar multiplication hold in that slot for the full expression-valued score.

For labels $f_i \in E_{\le d_i}$, define the mixed score by the finite sum

$$\mathcal{B}_d(f_1,\dots,f_k) = \sum_{0\le a_i\le d_i} \mathcal{A}_{(a_i)}(M_{a_1}^{(d_1)}f_1,\dots,M_{a_k}^{(d_k)}f_k), \tag{UC2}$$

where $\mathcal{A}$ is the homogeneous score for the fixed chart, partition and route parameters. If all inputs are homogeneous of their stated grades, every summand except $a_i = d_i$ for all i vanishes, so (UC2) is the original homogeneous score.

Raising the grade caps leaves (UC2) unchanged. Embed every original tuple (a_i) in the enlarged finite index set. The splitting maps commute with inclusion, so those summands agree exactly. Any additional tuple has some $a_i > d_i$; the corresponding split component of the included input is zero, and multilinearity makes that entire summand zero. This proves cap independence at the expression level.

On a test support containing only distinct configurations, compatibility of contact charts identifies each homogeneous local image with the ordinary field product times the test function. At each insertion point the splitting identity and coherent inclusion give

$$\sum_{a=0}^{d} \mathrm{insert}_a(M_a^{(d)}f) = \mathrm{insert}_d(f).$$

Expanding the finite product of these sums recovers the ordinary mixed-label field word. Evaluation by any allowed spectator then yields exactly the original separated correlation times the test function. Every sum is finite before integration; the previously proved per-grade integral bounds justify taking its integrals term by term. Thus the mixed extension has the prescribed separated restriction, cap independence and multilinearity on the full field space.

The projection gain underlying (UC1) can now be made precise. Along a logarithmic clock T, an upper rate c means a bound by $e^{(c+\delta)T}$ eventually for every $\delta > 0$. A finite expression is affordable at layer s and cost c when it is a finite sum of separated field words whose coefficient for a word of total field weight D has rate at most $c - sD$. This class is closed under finite sums, scalar coefficients with their added rate, and eventual equality. Its supports follow its finite placement palette.

A tested error of precision q at a palette means that after multiplication by any separated affordable spectator word, at terminal layer one and of cost b, its expectation has upper rate $q + b$. This quantifies over the unselected hierarchy and every character, including varying choices of character. An expression has funds at layer s, cost c, if it has some finite terminal affordability bound and, for every precision q, differs by such a tested error from an affordable expression at layer s, cost c.

When a palette lies at radius rate e^{-sT} about a center and $0 \le u < s \le 1$, the one-scale OPE replaces an affordable word by a finite sum of fields at that center to arbitrary tested precision at layer u. Indeed choose the output horizon n so large that the gain $n(s-u)$ dominates the prescribed precision deficit. The coefficient estimate assigns every retained output its rate $c - ds$, where d is its output weight. The tested remainder estimate applies to each allowed spectator geometry on this path. The choice of horizon is allowed to depend on that path and precision; it is not asserted to be uniform before (UC1). Finite sums and scalar coefficients give the same replacement for every affordable expression, and adding its existing tested error gives the same replacement for every expression with funds.

For a centered finite sum of fields with coefficient bounds $c - ds$, the retained coefficient row of weight w has rate at most $c - ws$. A field of weight below w gives zero, while splitting a field of weight at least w applies a fixed linear map and cannot worsen its coefficient rate. To transfer this

statement from a finite approximation to the original expression, test the error against the finite detector family for that row, including its rotations and conjugates. All these tests lie outside the fixed row aperture. Their placement distances remain bounded below independently of T, so their mutual separation has layer zero. If probeFee(w) is the sum of their fixed weights, a precision-q error gives row rate at most $q+\mathrm{probeFee}(w)$. Choose q sufficiently small to absorb this cost. Hence an expression with funds (s,c) has row rate $c-ws$ and its projection through any finite horizon still has funds (s,c).

This row estimate includes the full rotation average. For any varying rotation, finite-dimensional covariance and compactness of the orthogonal group preserve the rate of a coefficient vector. The rate bound for every varying choice is uniform over rotations eventually: otherwise selecting a violating rotation at each bad index contradicts that bound. The normalized Haar integral is then bounded by the same exponential bound. Conjugation preserves norms, and the fixed finite decoding maps preserve rates. This proves the detector-row estimate without exchanging a nonuniform family of limits and an integral.

Let $\Pi_{\le w}$ be the prescribed finite projection. Subtracting it and passing to a coarser layer $u \le s$ gives

$$\mathrm{Funds}(s,c;e) \implies \mathrm{Funds}\,(u, c-(w+1)(s-u);\, e-\Pi_{\le w}e). \tag{UC3}$$

For a centered finite field sum this is immediate: all outputs of weight at most w cancel, and any surviving weight $d \ge w+1$ satisfies

$$c-ds \le c-(w+1)(s-u)-du.$$

For a general expression choose its centered approximation at precision $\min(q, q-\mathrm{maxFee}(w)-wu)$. Projection of its error costs at most $\mathrm{maxFee}(w)+wu$, where $\mathrm{maxFee}(w)=\sum_{j=0}^{w}\mathrm{probeFee}(j)$. Both the original error and its projection consequently fit precision q. Subtract them and use the finite-sum result for the approximation. Since q was arbitrary, (UC3) follows. At $u=s$, simply subtract the projected affordable expression, with no gain. Repetition of (UC3) supplies the full $w+1$ gain at every signed ancestor move of the contact forest.

Funds multiply across separated palettes with addition of costs. Suppose e has funds (s,c), f has funds (s,d), their palettes are separated at layer s, and $s \le 1$. Fix terminal affordability costs a,b for e,f. For any desired tested precision q, first approximate e by an affordable g with error precision $q-b$. After choosing g, enlarge its layer to one and record its finite terminal cost c'. Then approximate f by an affordable h with error precision $q-c'$. The identity

$$ef-gh=(e-g)f+(f-h)g$$

places both terms at precision q, by the definition of tested precision and the two recorded terminal costs. The product gh is affordable at layer s, cost $c+d$: concatenate its separated words and multiply their coefficients. This proves the multiplication rule for funds. Induction proves it for any finite family of pairwise separated palettes.

The order of these approximation choices matters. Although g has a controlled outgoing layer-s cost, its terminal cost at layer one can be much larger. Choosing the second approximation only after recording that finite cost pays for it without imposing a false uniform bound on all OPE horizons. The arbitrary-precision definition is exactly what permits this successive choice.

Funds also imply the stated expectation bounds. To test e against an affordable spectator g of layer v, cost q, with $s \le v \le 1$, first enlarge g to terminal layer one and record its finite cost b. Choose an affordable layer-s approximation to e whose error precision is $c+q-b$. The direct separated correlation estimate bounds the approximation times g at rate $c+q$, and the definition

of precision bounds its error times g at that same rate. Addition proves the desired tested bound. Taking the spectator to be the empty word gives expectation upper rate c, uniformly over all character choices. This closes the expectation and multiplication rules used with (UC3).

The avoidance geometry used for the equation lift can also be given explicitly. A lattice step that would traverse the forbidden positive edge in either direction can be replaced by three steps: move one unit in a different coordinate, make the intended step, and move back in that coordinate. The replacement has the same endpoints and avoids the forbidden stored edge. Replacing each offending step of any fixed coordinate path constructs a path between any two cover sites avoiding that edge.

For a plaquette of axes $i < j$, its boundary can use a fixed forbidden edge based at b only if its own base belongs to

$$\{b,\ b - e_i,\ b - e_j\}.$$

This follows by listing its four stored positive edges, including the bases of its two reversed steps. Outside this finite set, take an avoiding transport path to the plaquette base as above; both the transport and the boundary then avoid the edge. These are the safe sampling probes used in the finite-dimensional span argument.

Avoidance on the integer cover yields exact independence on all sufficiently large tori, uniformly over translations of the stencil. For each of the finitely many used cover edges choose the side larger than twice the largest absolute coordinate of its base and of the forbidden base, plus one. Their centered wrapped lifts then recover those bases exactly. Distinct stored cover edges stay distinct after wrapping. The link variation therefore changes none of the path factors, so their holonomies, curvatures, adjoint transports and all finite products are unchanged. Taking the maximum over the finite used-edge set gives the common side threshold in the marking identity.

15 The Gaussian short-distance comparison

The short-distance comparison has two normalizations. A weak scaling sequence compares the actual hierarchy with the free Gaussian hierarchy. A finite determining collection of position tests then adjusts the coefficient frame by a small invertible map. Those finitely prescribed columns are matched exactly. General interacting moments still satisfy the purchased error bounds; they are not all equal to their Gaussian counterparts at a finite station.

15.1 Leading perturbative comparison on a common weak scaling sequence

There is a further conclusion about the mean Euclidean hierarchy. Its scale parameter is the same weak sequence $r_k \to 0$, with ultraviolet inverse scale r_k^{-1}; it does not assert interpolation to every small positive real scale. Let E_d^0 be the free Maxwell field spaces and $U_{k,d} : E_d^0 \to E_d$ the simultaneous scaling frames. Their unit, conjugation, signed rotation and filtration compatibilities hold eventually in each fixed degree. The transported first derivatives satisfy

$$r_k U_{k,d+1}^{-1} \partial_a U_{k,d} \longrightarrow \partial_a^0, \tag{PT1}$$

and the transported infinitesimal rotations converge to the free ones. The real conserved stress space has dimension one.

The free stress comparator is translation-normalized: it is the negative of the source's original Wick Maxwell tensor. This sign is fixed by a Ward identity, rather than by conservation alone. If a nonzero first-derivative family D and one tensor T satisfy the translation shell identities both with cD and with $c'D$, subtract those identities, test against arbitrary separated exterior fields and use

field separation. The resulting equality $(c - c')D = 0$ gives $c = c'$. Apply this to the constructed free tensor and the explicit Maxwell tensor, whose translation convention is $-D^0$. It forces their comparison constant to be -1. The explicit curvature-label calculation identifies the underlying tensor with the original Maxwell tensor, giving the claimed negative comparator, which we denote $T^{0,\mathrm{tr}}$.

Fix any finite list of free homogeneous labels q_i, any subset of its slots to be replaced by stress components, and a pair of tensor indices in each such slot. In a regular slot use $U_{k,d_i} q_i$; in a stress slot use $r_k^4 T_{a_i b_i}$. Then, for every compact set K of jointly distinct positions and every position-derivative order j, their mean moments at $r_k x$ converge uniformly on K, with those derivatives, to the free Wick moments with stress slots $T^{0,\mathrm{tr}}_{a_i b_i}$. All arities, finite label choices and derivative orders use the same sequence and frames.

To prove this passage, pull every slot back by U^{-1}_{k,d_i}. Regular slots become constant labels; stress slots converge to $T^{0,\mathrm{tr}}_{a_i b_i}$ by the physical stress comparison and its Ward normalization. Expand the resulting finite-dimensional labels in a fixed basis. The common compact jet comparison handles the finitely many basis moments, while coordinate convergence and compact boundedness of the free derivatives handle the varying coefficients. This proves the mixed statement with any number of stress slots. The construction of the sequence supplies the mean hierarchy, derivatives, rotations and physical stress on the same table as the selected and reconstructed theory; no change of continuum witness is made here.

The position normalization through degree four is a second finite-dimensional construction. At a general finite degree horizon d, let $\mathcal{B}_d$ be the real vector space of allowable small endomorphisms b of the free field space E^0_d. They annihilate the unit, preserve every lower filtration subspace, commute with conjugation and signed symmetries, and commute with multiplication by i. The last property makes $I + b$ complex linear. Put

$$M_b = I + b. \tag{PT2}$$

For b near zero it is invertible; the inverse follows from a geometric series in the endomorphism norm. It fixes the unit and preserves the listed structures. The determining finite collection of position tests will choose $b = b_k \to 0$, and the adjusted frame is $U_{k,d} M_{b_k}$.

At horizon four the adjustment preserves all lower fields exactly. The free degree-d field space for $d < 4$ consists of unit multiples for a simple gauge algebra. The degree-raising maps take that unit to the unit. Since b annihilates the latter, M_b fixes every raised lower field. Thus the degree-four adjustment is compatible with the already chosen frames in degrees zero through three.

A determining position test consists of a finite arity n_i, labels $v_{i1}, \dots, v_{in_i} \in E^0_d$, and a smooth compactly supported complex profile ψ_i whose support lies in the open set of distinct configurations. A finite real linear map Λ combines the real parts of these readings. Choose a complex basis e_a of E^0_d. A finite array w records the profile-integrated basis moments, for each test and every basis tuple. Multilinearity reconstructs the reading of any tuple:

$$D_i(w; v_1, \dots, v_{n_i}) = \sum_{a_1, \dots, a_{n_i}} \prod_j (v_j)_{a_j} \, w_{i, a_1, \dots, a_{n_i}}.$$

Define the finite position map

$$\mathcal{C}_Q(w, b) = \Lambda\big(\operatorname{Re} D_i(w; M_b v_{i1}, \dots, M_b v_{in_i})\big)_i. \tag{PT3}$$

It is a polynomial in its finite real coordinates. Let w^0_Q be the array obtained by integrating the free Wick moments against the profiles. A square determining collection means exactly that

$$\partial_b \mathcal{C}_Q(w^0_Q, 0) : \mathcal{B}_d \longrightarrow \mathbb{R}^{\dim \mathcal{B}_d} \quad \text{is invertible.} \tag{PT4}$$

This condition is proved for a suitable collection at degree four.

For this existence proof, a single test's derivative in direction b is the real part of the integral of

$$\mathcal{J}_b(x; v_1, \dots, v_n) = \sum_j \langle v_1(x_1) \cdots (bv_j)(x_j) \cdots v_n(x_n) \rangle_0. \tag{PT5}$$

If all these real integrated derivatives vanish, replace each profile by $i\psi$ to see that their imaginary parts vanish as well. Continuity on distinct configurations and smooth compact test separation then show that (PT5) vanishes pointwise there. The degree-four rigidity calculation below proves that this forces $b = 0$. The resulting real linear functionals therefore separate $\mathcal{B}_4$. In a finite-dimensional space finitely many suffice: choose a finite spanning subset of their span in the dual. Their combined map has zero kernel, hence a linear left inverse. Compose that left inverse with coordinates on $\mathcal{B}_4$ to choose Λ. By differentiation of (PT3), its b-derivative is then precisely that coordinate isomorphism. This proves (PT4), with profiles supported away from every collision.

The finite-dimensional implicit-function theorem now provides a function ϕ on a neighborhood of w_Q^0, with

$$\phi(w_Q^0) = 0, \qquad \mathcal{C}_Q(w, \phi(w)) = \mathcal{C}_Q(w_Q^0, 0), \tag{PT6}$$

and local uniqueness: for (w, b) near $(w_Q^0, 0)$, the equality in (PT6) holds if and only if $b = \phi(w)$. The map is smooth to every finite order at the base point. The proof uses only the polynomial map and the invertible derivative (PT4).

Let $w_Q(k)$ be the same array measured in the actual mean hierarchy with labels $U_{k,d}e_a$ at positions $r_k x$. Compact separated convergence implies $w_Q(k) \to w_Q^0$. To justify the integrals, the smooth profile has a fixed compact separated support, so the uniform moment error is bounded by its supremum there times $\|\psi_i\|_{L^1}$. Set

$$b_k = \phi(w_Q(k)). \tag{PT7}$$

Then $b_k \to 0$, M_{b_k} is invertible eventually, and (PT6) gives the exact determining equations in the interacting mean hierarchy. Equality of (PT3) with those measured equations follows by expanding the inserted labels in the finite basis and exchanging a finite sum with the integrable profile. For any other finite collection R of such separated profiles, continuity of $\mathcal{C}_R$ and convergence $(w_R(k), b_k) \to (w_R^0, 0)$ give

$$\mathcal{C}_R(w_R(k), b_k) \longrightarrow \mathcal{C}_R(w_R^0, 0). \tag{PT8}$$

Thus the exact normalization does not spoil any other leading position test.

The endpoint also compares this construction with numerical jets of finite Wilson data. For a fixed finite torus with side greater than one, let $w_{Q,\ell}(m)$ denote the coefficient of order ℓ of its regulated profile series at mesh parameter m. Its order-zero coefficients tend to w_Q^0 as $m \to \infty$. For a fixed truncation order p, form the finite polynomial curve $c_{Q,p,m}(t) = \sum_{\ell=0}^{p} t^\ell w_{Q,\ell}(m)$. This curve is used only to compute numerical derivatives. It requires no convergence of the full formal series. Equation (PT6) holds along the curve near $t = 0$, for all sufficiently large m, because its base point tends to w_Q^0. Therefore every positive derivative at zero of $\mathcal{C}_Q(c_{Q,p,m}(t), \phi(c_{Q,p,m}(t)))$ vanishes exactly, and its zeroth derivative equals the prescribed constant.

For another query R, the order-zero numerical reading is

$$\mathcal{C}_R(w_{R,0}(m), \phi(w_{Q,0}(m))) \longrightarrow \mathcal{C}_R(w_R^0, 0). \tag{PT9}$$

Subtract (PT9), at $m = k$, from (PT8). Their difference tends to zero on the same actual-theory sequence. The order-zero numerical reading is independent of the chosen polynomial truncation

p. These are exactly the numerical-jet calibration and leading-comparison assertions: calibration of the determining equations at each finite formal horizon, and order-zero comparison for other profiles. They do not assert comparison of every higher-order coefficient with the actual interacting hierarchy.

Any other small correction satisfying the same exact determining equations equals (PT7) eventually, by the local uniqueness in (PT6). This identifies the correction chosen by the local map with the correction appearing in the endpoint. Combined with (PT2), the unit and symmetry identities and the exact preservation of raised fields of degree below four, it proves all the frame properties claimed for the degree-four position normalization.

Here is the degree-four rigidity used in (PT5). Its free field basis consists of the unit and the color-contracted curvature products

$$e_c = \sum_{\alpha} F^{\alpha}_{ab} F^{\alpha}_{uv}, \tag{PT10}$$

indexed by unordered pairs c of oriented coordinate planes, each plane put in its increasing-axis orientation. Reversing a plane changes the sign and interchanging the planes preserves the product. The free quotient calculation proves that these labels, together with the unit, form a basis; its dual coordinates will be denoted ℓ_* and ℓ_c.

At the separation $(e_0, 0)$, the free curvature two-point form is diagonal on the six coordinate planes, with nonzero coefficients γ_p. Wick contraction therefore gives

$$\langle v(e_0) e_c(0) \rangle_0 = a_c \ell_c(v), \qquad a_c \neq 0. \tag{PT11}$$

For a repeated plane the contraction has its two equal pairings; for two different planes it has the single matching pairing. Accordingly a_c is the product of the two nonzero curvature coefficients and $\dim \mathfrak{g}$, multiplied by the corresponding nonzero Wick multiplicity. The dual-coordinate normalization absorbs that multiplicity. The unit coordinate is expectation. If $\ell_*(v) = 0$, the reversed pairing $\langle e_c(e_0) v(0) \rangle_0$ has the same formula. These facts follow by expanding v in the basis and applying the diagonal contractions, so no positive-definiteness of this particular bilinear form is required.

For each plane pair $c = (p, q)$, let its parity set be the symmetric difference $p \triangle q$ of their axis sets. Reflection in axis a multiplies e_c by -1 exactly when a belongs to this parity set. If b commutes with all such reflections, its matrix coefficient $b_{cd} = \ell_c(b e_d)$ vanishes whenever c, d have different parity sets.

For equal parity sets there is an involutive permutation of the four axes exchanging the two plane pairs, up to reversing individual plane orientations and swapping the planes. This finite assertion has three cases. If the parity set is empty, both pairs repeat a plane, and a transposition or a product of two disjoint transpositions exchanges the two planes. If it has two elements, each pair consists of two planes sharing one of the other two axes; exchange those common axes when they differ. If it has four elements, each pair is a partition into two complementary planes; a transposition exchanges any two different such partitions. The identical-pair case uses the identity. These permutations belong to the list consisting of the identity, six transpositions and three products of disjoint transpositions.

The induced action sends e_c to ϵe_d and e_d to ϵe_c, with the same sign $\epsilon \in \{1, -1\}$: applying the involution twice gives product of signs one and the basis vectors are nonzero. Commutation of b with this action gives $b_{cd} = b_{dc}$. Moreover $a_c = a_d$ for equal parity sets. Their Wick multiplicities agree, and the product of their two curvature signs at the axis-zero separation depends only on whether zero belongs to the parity set; their absolute curvature factors are the same. This proves the equality of the coefficients in (PT11).

Now suppose (PT5) vanishes for every separated configuration and tuple. Arity one gives $\ell_*(bv) = 0$ for every v. Arity two at $(e_0, 0)$ gives

$$\langle (bv)(e_0)w(0)\rangle_0 + \langle v(e_0)(bw)(0)\rangle_0 = 0. \tag{PT12}$$

Take $v = e_c, w = e_d$. Equation (PT11) and the zero unit coordinate give $a_d b_{dc} + a_c b_{cd} = 0$. For different parity sets both coefficients already vanish. For equal parity sets, symmetry and $a_c = a_d \neq 0$ give $2a_c b_{cd} = 0$, hence again both vanish. The unit column vanishes because $b1 = 0$; the unit row vanishes by arity one. Every coefficient of b is zero, proving the rigidity assertion. This closes the existence proof of the square position tests through degree four.

The physical stress comparison used before (PT2) can be proved directly on this weak sequence. Let $J_{\mathrm{raw}} \neq 0$ be its constructed raw conserved centered tensor, and let $\xi_k \neq 0$ be the scale used in its preliminary stress comparison. The one-dimensional divergence kernel implies

$$J_{\mathrm{raw}} = c\ \mathrm{center}(T), \qquad c \neq 0.$$

The centered tensor is nonzero because its translation Ward charge differentiates a nonzero field. Compare its charges with those of J_{raw}. Field separation applied to every translation shell identity gives equality of their derivative families, $D_{\mathrm{raw}} = cD$. The weak derivative comparison then yields $c/\xi_k \to \kappa_0$, where $\kappa_0 \neq 0$ is the free stress Ward amplitude. Multiplying the preliminary raw stress limit by $(c/\xi_k)^{-1}$ gives

$$U_{k,4}^{-1}(r_k^4\,\mathrm{center}(T)) \longrightarrow \kappa_0^{-1} T^0. \tag{PT13}$$

The difference from T is a fixed constant unit tensor. Its inverse-scaled contribution is r_k^4 times that unit, tending to zero; the unit identifications are exact eventually. The Ward argument preceding (PT2) gives $\kappa_0 = -1$, proving the stress-slot limit used there. Similarly, the preliminary angular-generator limit is normalized by the factor $(c/\xi_k)^{-1}$, giving the transported rotation limit in the final physical normalization.

The compact-profile convergence needed in (PT7) requires no distribution extension through collisions. If ψ is one of the allowed compact separated profiles and $f_k \to f_0$ uniformly on its support, take $\delta = \varepsilon/(1 + \|\psi\|_{L^1})$. Eventually

$$\left| \int \psi(f_k - f_0) \right| \leq \|\psi\|_{L^1} \sup_{\operatorname{supp}\psi} |f_k - f_0| \leq \|\psi\|_{L^1}\delta < \varepsilon.$$

Smoothness of the separated moments makes both integrals ordinary Bochner integrals. This supplies every component of $w_Q(k) \to w_Q^0$.

For the regulator arrays in (PT9), the order-zero Wilson coefficient is the Gaussian Wick contraction of the reduced local field polynomials. The common field-address identity identifies these polynomials with the free field labels. At fixed lattice refinement base $L > 1$, the denominator is L^m, the mesh spacing is L^{-m}, and the chosen period tends to infinity fast enough for every fixed polynomial margin. The Gaussian row limit holds for every refining subsequence and every sequence of grid tuples tending to a distinct continuum tuple. The grid-sum limit for a compact separated profile then gives

$$w_{Q,0}(m) \longrightarrow w_Q^0. \tag{PT14}$$

The Gaussian kernel, Wick row and grid-sum arguments underlying this assertion are developed next, including the finite-regulator subtraction and its limiting cancellation.

15.1.1 From moving Gaussian rows to profile integrals

Here is the grid-sum passage in (PT14). Let $h_m > 0$ tend to zero, $h_m = D_m^{-1}$ with integer D_m, and let the torus side N_m satisfy $h_m N_m \to \infty$. For an arity n, round every coordinate of a continuum configuration x to the grid. Denote the rounded point by $\rho_m x$, so

$$|\rho_m x - x| \le 4nh_m. \tag{PT15}$$

Every fixed bounded set eventually embeds in the torus without wrapping, because $h_m N_m \to \infty$.

Suppose lattice kernels K_m have the following joint moving limit: for every subsequence $m_j \to \infty$, every x in the allowed open region, and every grid tuple o_j with embedded positions tending to x, one has $K_{m_j}(o_j) \to K_0(x)$. This is stronger than convergence only at a fixed rounded point and is exactly what the Gaussian row calculation supplies. It implies a uniform eventual bound for K_m on rounded points in any fixed compact allowed set. Otherwise choose $m_j \ge j$ and points x_j in that compact set with $|K_{m_j}(\rho_{m_j} x_j)| > j$. A convergent subsequence $x_j \to x$, (PT15) and eventual nonwrapping give a sequence of embedded grid tuples tending to x, contradicting the moving limit.

For a continuous compactly supported profile ψ with support in the allowed region, define the step function

$$f_m(x) = \psi(\rho_m x) K_m(\rho_m x).$$

It converges pointwise to $\psi(x)K_0(x)$. On the support use the moving limit and continuity of ψ; outside its closed support the rounded points are eventually outside it, so the expression vanishes. If ψ is supported in the radius-R ball, the step functions, for $h_m \le 1$, are supported in the radius-$R + 4n$ ball and are bounded by a common constant there. They are integrable step functions. Dominated convergence gives convergence of their integrals.

Each grid cell has exactly volume h_m^{4n}. Summing the cells, with no wrapping on the fixed support, identifies the integral exactly with

$$h_m^{4n} \sum_{o \in (\mathbb{Z}/N_m\mathbb{Z})^{4n}} \psi(h_m o) K_m(o). \tag{PT16}$$

The cell partition and its volume hold also at arity zero, when both the configuration space and the grid consist of one point and $h_m^0 = 1$. This proves convergence of (PT16) to the desired continuum integral with its correct normalization.

The Gaussian row limit is reduced to a two-point covariance limit by an exact finite recursion. Give every occurrence of every generator in the finitely many local polynomial labels a distinguishable index a. Let $\tau(a)$ be its insertion slot and let ζ_a be its field-generator label. At order zero, the row is the Gaussian moment $M_m(A)$ of the product of the corresponding centered linear variables, for each finite set A of occurrences. Gaussian integration by parts gives

$$|A| M_m(A) = \sum_{a \in A} \sum_{b \in A \setminus \{a\}} C_m(a,b)\, M_m(A \setminus \{a,b\}). \tag{PT17}$$

For $A = \varnothing$, $M_m(A) = 1$ because the Gaussian law is a probability measure. The continuum Wick moments obey the same recursion with covariance $C_0(a,b)$. Induct on $|A|$: the covariance limit and the finitely many already proved smaller-product limits give convergence of the right side, and division by the positive integer $|A|$ gives convergence of $M_m(A)$.

This occurrence indexing accounts for repeated powers without identifying identical generators prematurely. An arbitrary monomial is the product over its occurrence set; arbitrary finite polynomial labels are finite sums of such monomials with their coefficients. Multilinearity and finite summation extend the preceding limit to every tuple of the allowed local polynomial sections. The

field-address identity sends the reduced polynomial section of a free label to that same label in the free Wick hierarchy. Therefore the limiting row is precisely the continuum free moment, not an independently fitted polynomial. Applying the result along every refining subsequence supplies the moving-limit hypothesis used in (PT16). This proves (PT14), once the corrected covariance limit used in (PT17) is established.

15.1.2 Private Gaussian marks and the subtraction covariance

Private Gaussian marks distinguish contraction occurrences without changing the physical covariance. Their subtraction covariance must be computed with the stated Fourier and opposite-momentum convention. Only after the corrected bilinear covariance is identified can the occurrence recursion and scalar return estimate be used for higher products.

The covariance used in (PT17) has a finite-regulator subtraction that must be retained before taking the limit. Fix the largest jet grade d occurring in a polynomial row. Its finitely many linear coordinates are indexed by keys a. Let Z be the centered real Gaussian coordinate vector of the regulated potential, with precision matrix P and covariance P^{-1}. Write $k_{h,x,a}$ for the coefficient vector of the linearized, symmetrically averaged jet coordinate at x. Thus its value on Z is $k_{h,x,a}^{T}Z$. These coordinates use the same Gaussian density as the regulated perturbative integral.

The subtraction matrix is defined on the infinite lattice. If c ranges over sine and cosine waves, vector axes and color indices, and $l_{h,a,c}(p)$ is the jet coordinate evaluated on that wave, set

$$S_h(a,b) = (2\pi)^{-4} \int_{[-\pi,\pi]^4} \delta(p)^{-1} \sum_c l_{h,a,c}(p) l_{h,b,c}(p)\, dp. \tag{PT18}$$

Here δ is the lattice Laplacian symbol. The wave responses are continuous on the cube, hence bounded; δ^{-1} is integrable in four dimensions, so every entry exists. Symmetry is immediate. For any real vector v,

$$v^T S_h v = (2\pi)^{-4} \int_{[-\pi,\pi]^4} \delta(p)^{-1} \sum_c \Big(\sum_a v_a l_{h,a,c}(p) \Big)^2 dp \geq 0.$$

Consequently S_h is the covariance of a possibly degenerate centered real Gaussian vector. No inverse of this matrix is needed.

For insertion indices $i = 1, \ldots, n$, take independent copies Y_i of this Gaussian vector, independent also of Z, and put

$$X_{i,a} = k_{h,x_i,a}^{T} Z + \mathrm{i}\, Y_{i,a}.$$

All these variables are complex linear functions of one finite-dimensional real Gaussian vector. Their means vanish and their **bilinear** covariance is

$$C_h((i,a),(j,b)) = k_{h,x_i,a}^{T} P^{-1} k_{h,x_j,b} - \mathbf{1}_{i=j} S_h(a,b). \tag{PT19}$$

Indeed the mixed terms vanish by independence and centering; the two private marks have covariance zero for different insertions, and multiplication of their imaginary coefficients gives the minus sign for equal insertions. This is a bilinear covariance, with no complex conjugation. At finite regulator its same-insertion entries need not vanish.

The constant term of the normalized perturbative integral is exactly the expectation of the product of the reduced insertion polynomials evaluated at these $X_{i,a}$. To see this directly, the weight and vacuum denominator both have constant coefficient one. The constant coefficient of each linear insertion generator is its physical Gaussian coordinate plus its imaginary private mark.

Expanding a polynomial is a finite operation, so taking constant coefficients commutes with that expansion. Integration over each private mark produces the prescribed subtraction within its own insertion. The product law then combines these independent integrals with the physical Gaussian integral. Every integrand is a polynomial in a finite Gaussian vector and is integrable: a polynomial is bounded by $C(1 + \|z\|)^q$, and all Gaussian moments are finite. This also justifies each finite expansion and use of Fubini.

For completeness the moment recursion does not require a nonsingular Gaussian law. A centered real linear Gaussian variable L, of variance v, has moment generating function $\exp(vt^2/2)$. Differentiating at zero gives

$$\mathbb{E}L^{r+1} = rv\,\mathbb{E}L^{r-1}.$$

Apply this identity to arbitrary real linear combinations of a finite family of centered real linear probes and compare coefficients. Polarization gives

$$|A|\,\mathbb{E}\prod_{\alpha\in A} X_\alpha = \sum_{\alpha\in A}\sum_{\beta\in A\setminus\{\alpha\}} \mathbb{E}(X_\alpha X_\beta)\,\mathbb{E}\prod_{\gamma\in A\setminus\{\alpha,\beta\}} X_\gamma. \tag{PT20}$$

Initially the probes can be taken real. Expanding each complex probe into its real and imaginary parts extends the identity complex multilinearly. Occurrences, including repeated occurrences of the same generator, remain distinct elements of A. The empty product has expectation one. These facts prove the finite Gaussian recursion used above.

15.1.3 The limit of the corrected covariance

Same-insertion subtraction The corrected covariance removes contractions within one prescribed insertion while retaining contractions between distinct insertion positions. This is an occurrence-level subtraction: repeated generator names do not merge their occurrence indices. Its two limits have different purposes. At different positions it recovers the Maxwell two-point kernel; at the same insertion it tends to zero. The latter is a regulator-removal conclusion, not an identity imposed on every finite torus covariance.

The subtraction in (PT19) has the following two different limiting roles. For distinct physical limiting insertion positions $z_i \neq z_j$,

$$C_h((i,a),(j,b)) \longrightarrow K(k_a, k_b; z_i - z_j),$$

where k_a is the corresponding coordinate functional on free solution jets and K is their Maxwell two-point kernel. For $i = j$,

$$C_h((i,a),(i,b)) \longrightarrow 0. \tag{PT21}$$

Thus zero same-insertion covariance is a conclusion after removing the regulator. It is not an identity for a finite torus.

Here is the averaging step that identifies both limits. Let $\mathcal{S}$ be the finite group of signed coordinate transformations and $c_{\mathcal{S}} = |\mathcal{S}|^{-1}$. For a jet functional A, its averaged lattice stencil has the exact form

$$A(\mathrm{wideLinear}_{h,x}) = c_{\mathcal{S}} \sum_{s\in\mathcal{S}} A \circ \mathrm{turn}_s \circ \mathrm{radialMap}(\mathrm{chart}_{s,h,x}). \tag{PT22}$$

Denote the functional inside this sum by A_s. On free equation solution jets the radial reconstruction is an inverse to sampling; equivariance therefore gives

$$\mathrm{plain}_s(A_s) = A \circ \mathrm{gaugeJet}\,.$$

In particular pairing independently transformed stencils with their inverse turns gives the same continuum two-point form for every pair s, t. The relative endpoint positions are kept fixed. Bilinearity expresses the finite-torus covariance as

$$c_{\mathcal{S}}^2 \sum_{s,t} \text{ordinary}_{s,t}(h; x, y; A_s, B_t).$$

Similarly expanding the two factors in the wave integral (PT18), then interchanging its finite sums and integral, expresses S_h as the same double average of the infinite-lattice shell forms. The integrability already established justifies this interchange. The separated chart limit gives the same free pairing for every summand when $x - y$ tends to a nonzero physical vector. The return chart limit says that, at equal endpoints, each ordinary form minus its shell form tends to zero. Taking the finite double average proves (PT21) and the separated limit.

The reporting schedule meets both chart-limit hypotheses. With an integer side $L > 1$, its spacing and period are

$$h_m = L^{-m}, \qquad N_m = 4L^{m^2+m+1}, \qquad h_m N_m = 4L^{m^2+1}. \tag{PT23}$$

In particular $N_m h_m^q \to \infty$ for every fixed integer q, and $4 \mid N_m$. Center the observed positions at the quarter-period point. If their centered physical coordinates converge, their centered integer coordinates are eventually less than $N_m/4$ in absolute value, since $h_m N_m \to \infty$. Hence every fixed stencil shift of their difference is represented without wrapping and has the required physical limit. A sum of the two positions lies at twice the quarter-period center plus a bounded physical displacement. Its lift has at least one coordinate of absolute value at least $N_m/4$. Its physical distance is consequently at least $h_m N_m/4 \to \infty$. This puts the inversion image outside every bounded physical region, including when the two observed positions coincide. It supplies the return limit used for the subtraction.

Finally apply (PT20) by induction on the number of occurrences. Covariance convergence (PT21) gives convergence of each moment to the recursion with zero within-insertion pairings and Maxwell pairings between distinct insertions. This is precisely the free Wick rule for reduced insertion polynomials. Finite polynomial expansion proves convergence of every moving row. The compactness and cell-integration argument in (PT15)–(PT16) then proves convergence of their profile integrals. This completes the Gaussian constant-coefficient step needed in the leading perturbative comparison.

The reduction of the chart forms to scalar lattice kernels is also explicit. An elementary stencil a is a finite signed sum of shifted edge readings, with coefficients a_t, offsets u_t, a vector axis, and an ordered derivative list D_a. Its Fourier polynomial factors as a sign, a translation character, and the product of the difference multipliers in D_a. Thus, for two such stencils,

$$\sum_{t,s} a_t b_s \chi_p(z + u_t \pm v_s) = \sigma_{a,b,\pm} \chi_p(z + u_{a,\pm} + u_{b,\pm}) \prod_{\nu \in D_a D_b} q_\nu(p). \tag{PT24}$$

This follows by multiplying their two finite Fourier polynomials and using $\chi_p(x + y) = \chi_p(x)\chi_p(y)$. The prescribed offsets and signs are precisely those in the elementary stencil factorization, so no limiting replacement occurs in this identity. Multiplying by the torus inverse-Laplacian Fourier weight and summing over p identifies their covariance with the corresponding differentiated scalar torus kernel, multiplied by $h^{-|D_a|-|D_b|-2}$ and the same sign.

For the infinite-lattice expression, sum the sine and cosine channels first. The elementary identity

$$\cos(p \cdot u)\cos(p \cdot v) + \sin(p \cdot u)\sin(p \cdot v) = \cos(p \cdot (u - v))$$

shows that the sum over channels is zero unless both vector axes and colors match; when they match it is the cosine of the offset difference. Integrating against $(2\pi)^{-4}\delta(p)^{-1}$ gives exactly the infinite-cover scalar Green function at that difference. The bound $|\cos| \le 1$ and integrability of δ^{-1} justify every such integral.

A curvature coordinate is a difference of two elementary edge stencils. Its two-point form consequently has four terms, with signs $+, -, -, +$. Both its finite-torus covariance and its infinite-cover subtraction have exactly this four-term expansion. At separated positions, (PT24) reduces each torus term to the scalar differentiated-kernel limit. At inversion-image positions it reduces each term to the scalar far-distance limit, which is zero. At coincident positions subtract the infinite-cover value before rescaling. The torus-to-cover return estimate makes this difference tend to zero even after multiplication by any fixed inverse power of h, because $N_m h_m^q \to \infty$ for every fixed q. These are respectively the separated, far-image, and return scalar estimates needed by the chart argument.

Finally elementary coordinate readings span the finite-dimensional dual of every raw jet-response space: these responses are arrays of color components, vector-index pairs and derivative-index words, and coordinate projections span the dual of a finite array space. Expand each of two arbitrary dual vectors as a finite linear combination of these readings. Bilinearity reduces their covariance convergence to the elementary curvature calculations. This establishes the passage from elementary stencils to arbitrary chart functionals used in (PT22).

15.1.4 Scalar return estimate with its rescaling rate

The scalar subtraction at a fixed integer displacement has a useful quantitative proof. Rescale the Fourier cube to $Q = [-\frac{1}{2}, \frac{1}{2}]^4$, put $d(x) = \delta(2\pi x)$, and let

$$f_a(x) = d(x)^{-1}\cos(2\pi a \cdot x), \qquad f_a(0) = 0.$$

The value assigned at zero agrees with the zero-mode convention for the inverse Laplacian and does not affect its integral. The elementary lower bound for the lattice symbol implies

$$|f_a(x)| \le \|x\|^{-2} \quad (x \in Q).$$

Let $v_N(x)$ be the cell representative in the torus Fourier quadrature. Outside cell boundaries, which have measure zero, $v_N(x) \in Q$, $\|v_N(x) - x\| \le 4/N$, and, when $v_N(x) \ne 0$, $\|v_N(x)\| \ge 1/N$. These imply $\|x\| \le 5\|v_N(x)\|$, hence

$$|f_a(v_N(x)) - f_a(x)| \le 26\|x\|^{-2}.$$

If $0 < r \le 1$, $8/N \le r$, and $\|x\| > r$, then $\|v_N(x)\| \ge r/2$. The numerator cosine and the symbol d have bounded derivatives on a ball containing Q. Subtracting two quotients therefore gives

$$|f_a(y) - f_a(x)| \le C_a r^{-4}\|y - x\| \quad (\|x\|, \|y\| \ge r),$$

and with $r/2$ in place of r,

$$|f_a(v_N(x)) - f_a(x)| \le A_a N^{-1} r^{-4} \quad (\|x\| > r).$$

For example the quotient subtraction is

$$(\cos_a(y) - \cos_a(x))d(y)^{-1} + (d(x) - d(y))d(y)^{-1}d(x)^{-1}\cos_a(x),$$

which proves the asserted bound directly from the Lipschitz bounds and $d^{-1} \le r^{-2}$.

The singular part is integrable. In four dimensions the change of variables $x = ru$ gives

$$\int_{\|x\|<2r} \|x\|^{-2}\, dx = r^2 \int_{\|u\|<2} \|u\|^{-2}\, du.$$

Split the quadrature error at radius r, use the bound $26\|x\|^{-2}$ inside and the uniform estimate outside, and obtain

$$|\mathrm{Re}\, G_N(a) - G_\infty(a)| \le C_a r^2 + A_a N^{-1} r^{-4}. \tag{PT25}$$

Here the torus Fourier sum is exactly the integral of $f_a \circ v_N$ over Q: every representative cell has volume N^{-4}. The infinite-cover value is exactly the integral of f_a over Q, by $p = 2\pi x$. Even N ensures the chosen symmetric representatives tile this same cube.

For any fixed integer $d \ge 0$, choose $r = h^{d+1}$. Along the schedule, eventually $r \le 1$ and $8/N \le r$. Multiplying (PT25) by h^{-d} gives

$$h^{-d}\, |\mathrm{Re}\, G_N(a) - G_\infty(a)| \le C_a h^{d+2} + A_a (N h^{5d+4})^{-1} \longrightarrow 0. \tag{PT26}$$

This proves the return estimate with every fixed rescaling power that occurs in a finite jet calculation.

15.1.5 Recovering the whole separated scalar kernel from annuli

Annular convergence alone determines differences of cutoffs and can leave an additive ambiguity. The small-cutoff bound and the large-cutoff tail bound below fix that ambiguity. The comparison chooses the small cutoff first, then a finite number of dyadic annuli, and only then the lattice index. No uniform convergence in a moving cutoff is assumed.

The separated Fourier limit uses a cutoff comparison that also fixes the otherwise undetermined constant between annular limits. Here is the complete comparison principle. Suppose $F_m \in \mathbb{C}$, cutoff functions $A_m(s)$, a continuum cutoff $A(s)$, and a candidate limit F satisfy, for every $s > 0$,

$$A_m(2s) - A_m(s) \longrightarrow A(2s) - A(s).$$

Suppose there are constants $C, B, E, D \ge 0$, independent of s, such that eventually in m, for each fixed s,

$$|A_m(s)| \le C s^p, \quad |F_m - A_m(s)| \le B/s, \qquad |A(s)| \le E s^p, \quad |F - A(s)| \le D/s, \tag{PT27}$$

where $p \ge 1$. Then $F_m \to F$.

Indeed choose $0 < s \le 1$ so that $(C+E)s^p < \varepsilon/3$, and next an integer q so that $(B+D)/(2^q s) < \varepsilon/3$. Finite telescoping of the annular convergence gives

$$A_m(2^q s) - A_m(s) \longrightarrow A(2^q s) - A(s).$$

For sufficiently large m its error is less than $\varepsilon/3$. Adding the two large-cutoff errors, this difference error, and the two small-cutoff sizes bounds $|F_m - F|$ by ε. The order of choices is s, then q, then m; no uniform convergence in the cutoff parameter is assumed.

Apply this principle to the differentiated torus Green kernel at a nonzero limiting physical displacement, with $A_m(s)$ its smooth Fourier cutoff, $A(s)$ the corresponding continuum cutoff, and $p = |D|+2$. The cutoff kernel is bounded by $Cs^{|D|+2}$. Its complement is bounded by B/s, uniformly

when the physical displacement is bounded away from zero. The Fourier mesh condition required for this estimate,

$$(2\pi/N)(|D|+4) \le hs, \qquad hs \le \tfrac{1}{2},$$

holds eventually for each fixed s, since $Nh \to \infty$ and $h \to 0$. The annular Riemann sum converges to the continuum annular integral. Thus (PT27) proves the separated scalar limit used in (PT24).

The continuum candidate in this argument is the ordinary derivative $\partial_D G(z)$, with $G(z) = (4\pi^2\|z\|^2)^{-1}$. One can identify it while retaining an integrable heat regulator $b > 0$: inverse Fourier transformation of the regulated differentiated inverse-square symbol gives

$$\partial_D G(z) - \partial_D\Big(G(z)e^{-\pi^2\|z\|^2/b}\Big) \quad (z \neq 0).$$

The Fourier convention in this formula has character $e^{2\pi ik\cdot z}$. To remove heat take $b = v^{-2}$, $v \to \infty$. The second term equals

$$v^{|D|+2}\partial_D\Big(Ge^{-\pi^2\|\cdot\|^2}\Big)(vz).$$

Every derivative of the fixed function on the right decays faster than any inverse power at infinity, so this term tends to zero for $z \neq 0$. At a fixed Fourier cutoff, dominated convergence also removes heat under the integral. The annular tail bound survives this limit. Summing its dyadic estimates $C/(2^j s)$ yields at most $2C/s$, supplying the continuum half of (PT27).

For a physical displacement whose norm tends to infinity, the same comparison principle applies with $F = 0$, $A(s) = 0$. The annular summation-by-parts estimate is

$$|\operatorname{annulus}_{h,s,D}(x)| \le \frac{C_D}{s}\left(\frac{2\pi}{\|h\ \operatorname{lift}(x)\|}\right)^{|D|+3}.$$

For each fixed s this tends to zero. The small-cutoff size and B/s tail estimates remain valid once the physical distance is at least one. This proves the far-image scalar limit.

The uniform low-frequency bound in (PT27) can be proved without oscillation. In normalized momentum coordinates $q = p/(hs)$, put $t = hs$. The lattice denominator is

$$d_t(q) = \sum_{\nu=0}^{3} q_\nu^2 \operatorname{sinc}(tq_\nu/2)^2.$$

On $\|q\| \le 1$, $0 \le t \le 2$, the inequality $\operatorname{sinc}(u) \ge 2/\pi$ for $|u| \le 1$ gives

$$d_t(q) \ge (2/\pi)^2\|q\|^2.$$

Each normalized difference numerator has modulus at most $|q_\nu| \le 1$, because $|e^{iu} - 1| \le |u|$. The normalized cutoff symbol for any derivative list is therefore bounded by $(\pi/2)^2\|q\|^{-2}$ on its unit-ball support, independently of the length of that list.

For centered integer momentum lifts, let B_m contain those whose largest coordinate in absolute value is at most m. Injectivity of the lift gives $|B_m| \le (2m+1)^4$. On $B_m \setminus B_{\lfloor m/2\rfloor}$, the inverse squared Euclidean norm is at most $4/m^2$. Strong induction now gives

$$\sum_{p\in B_m} \|\operatorname{lift}(p)\|^{-2} \le 432m^2,$$

where the zero summand is assigned zero. Indeed the inner contribution is at most $432(m/2)^2 = 108m^2$, and the outer contribution is at most $(3m)^4(4/m^2) = 324m^2$; the case $m = 0$ is zero.

Write $\eta = 2\pi/N$, $T = \eta/t$. Rescaling this last inequality bounds the inverse-square sum on $\|q\| \le 1$ by $432T^{-4}$. The physical normalization of the cutoff sum is

$$(Nh)^{-4} s^{|D|-2}.$$

Since $Ts = 2\pi/(Nh)$, multiplication gives

$$|\operatorname{piece}_{h,s,D}(x)| \le \frac{432(\pi/2)^2}{(2\pi)^4} s^{|D|+2}. \tag{PT28}$$

This is the first estimate in (PT27), uniformly also at coincident positions.

For the oscillatory estimate, repeated discrete integration by parts is exact on the finite torus. If Δ_{-e_ν} denotes a momentum difference, translation of the summation index gives

$$\sum_p \chi_p(x) \Delta_{-e_\nu}^{\,v} f(p) = (\chi_{e_\nu}(x) - 1)^v \sum_p \chi_p(x) f(p). \tag{PT29}$$

The normalized annular symbol is smooth, supported where $\|q\| \le 2$, and all its derivatives through any prescribed finite order are bounded uniformly for $0 < t \le 1$. This follows from its cutoff away from the denominator's zero and the smooth sinc extension at $t = 0$. A momentum difference has step size T in q. Repeated use of the fundamental theorem of calculus bounds its v-th difference by $C_{D,v} T^v$. If $vT \le 1$, the difference vanishes for $\|q\| > 3$. There are at most $8^4 T^{-4}$ momentum representatives in this enlarged support: each integer coordinate has absolute value at most $\lfloor 3/T \rfloor$, and $T \le 1$.

These statements remain valid across a representative boundary. The periodic uncut symbol agrees on the two lifts; the annular cutoff is zero on both whenever a shift crosses that boundary, provided $\eta v \le 1/2$. For the complementary high-frequency symbol, which has no compact support in q, one instead uses its uniform derivative bound on the rescaled Fourier box when $t \ge 1/2$. All N^4 representatives then contribute, but $N \le 4\pi/T$, so the same CT^{v-4} bound results. Combining these facts with (PT29) gives

$$|\chi_{e_\nu}(x) - 1|^v \left| \sum_p \chi_p(x) f(p) \right| \le C_{D,v} T^{v-4}.$$

Choose a coordinate whose physical displacement has absolute value at least half its Euclidean norm. For centered lifts, the chord inequality for the circle implies

$$|\chi_{e_\nu}(x) - 1| \ge \|h \operatorname{lift}(x)\|/(Nh).$$

Set $v = |D| + 3$ and restore the physical normalization. Using $Ts = 2\pi/(Nh)$ gives

$$|\operatorname{annulus}_{h,s,D}(x)| \le \frac{C_{D,v}}{(2\pi)^4 s} \left(\frac{2\pi}{\|h \operatorname{lift}(x)\|} \right)^v. \tag{PT30}$$

The same bound holds for the high-frequency complement when $1/2 \le hs \le 1$.

To bound the entire tail choose the first integer q for which $h2^q s > 1/2$. Then $h2^q s \le 1$. Decompose the tail into the annuli at $s, 2s, \dots, 2^{q-1}s$ and the final high-frequency complement. At a physical distance at least $R > 0$, (PT30) bounds the annular sum by $2C_{D,R}/s$ and the complement by $C_{D,R}/s$. Hence

$$|\operatorname{joint}_{h,D}(x) - \operatorname{piece}_{h,s,D}(x)| \le 3C_{D,R}/s.$$

This proves the torus tail estimate used in (PT27) and shows explicitly why an absolute bound on the uncut infinite-volume derivative sum is unnecessary.

The annular quadrature limit used above follows from a fixed-support dominated-convergence argument. Fix $s > 0$ and write $c_m = N_m h_m$, an integer for the reporting schedule. Physical Fourier sample points have spacing c_m^{-1}. The normalized annular integrand is jointly continuous in its lattice parameter $t = h_m s$, its physical displacement, and its Fourier variable, including $t = 0$. It vanishes outside a fixed ball of radius $2s/(2\pi)$ and is uniformly bounded there when $0 \le t \le 1$. Choose an integer R containing that support. Eventually $2Rc_m < N_m$, because this inequality is $2Rh_m < 1$. Thus every integer Fourier address in the supporting box is its own centered torus representative. The annular torus sum is exactly the unrestricted grid quadrature over this box.

Let $r_m(k)$ be the grid representative of k, so $\|r_m(k) - k\| \le 4/c_m$. If F_m is the annular integrand, define the step function $F_m(r_m(k))$. Its integral is the quadrature, since each cell has volume c_m^{-4}. The step functions converge pointwise to the continuum annular integrand by joint continuity, $t \to 0$, convergence of the physical displacement, and $r_m(k) \to k$. They are bounded by a common constant times the indicator of the ball of radius $R + 5$: a cell whose representative lies in the support is within distance four of it. Dominated convergence proves convergence of their integrals. The exact phase identity underlying this representation is

$$e^{2\pi i \langle h_m \operatorname{lift}(x_m), \operatorname{lift}(p)/c_m \rangle} = \chi_p(x_m).$$

The normalized radial variable is also exactly the lattice radial variable, since $c_m = N_m h_m$. This proves the annular Riemann-sum statement with moving displacements.

For clarity the continuum oscillatory bound can be obtained uniformly in the heat parameter. Let $M_D(q)$ be the normalized continuum annular symbol and

$$H_{D,a}(q) = e^{-a\|q\|^2} M_D(q), \qquad a \ge 0.$$

Its support is contained in the compact annulus $1/2 \le \|q\| \le 3$. For each fixed derivative order j, the derivatives through order j are bounded by a common constant times the indicator of this annulus, uniformly in $a \ge 0$. To prove the heat-factor part, put $y = \sqrt{a}\, q$. Its i-th derivative is $a^{i/2}$ times the i-th derivative of $e^{-\|y\|^2}$. When $\|y\| \le 1$, the lower bound $\|q\| \ge 1/2$ implies $\sqrt{a} \le 2$, so compact derivative bounds suffice. When $\|y\| \ge 1$, use $\sqrt{a} \le 2\|y\|$ and the rapid decay of the Gaussian derivatives. Leibniz's rule and bounded derivatives of M_D give the stated common bound.

Integrating by parts in the continuum Fourier integral now gives, for every integer $j \ge 0$,

$$\|x\|^j \left| \int e^{2\pi i k \cdot x} H_{D,a}(k)\, dk \right| \le C_{D,j}, \qquad a \ge 0.$$

For example select a coordinate with $|x_\nu| \ge \|x\|/2$, integrate j times in that coordinate, and bound the resulting derivative in L^1. There are no boundary terms because the integrand is smooth and compactly supported. Rescaling $k = (s/(2\pi))q$ in the continuum annular integral contributes $s^{|D|+2}$; taking $j = |D| + 3$ leaves a factor s^{-1}. The rescaled heat parameter remains nonnegative, so the bound is uniform during heat removal. This proves the continuum annular bound and the summable dyadic tail asserted after (PT27).

16 The centered trace and its non-freeness test

The trace is obtained from actual finite current and stress functions, not from an assumed anomaly formula. Localizing those functions introduces explicit finite residuals. The coefficient detector,

common scaling panels and equation marker control those residuals, and a two-coordinate correction then yields an exactly conserved physical stress tensor. Its centered trace is the field whose positive spectral measure gives the final asymptotic and comparison exclusions.

16.1 The finite trace candidate, its equation identity and its compensation

The local current and stress entering this chapter are finite configuration functions. For a base site o, let

$$C_{ab}(o,U) = \operatorname{curv}(U_a(o)U_b(o+e_a)U_a(o+e_b)^{-1}U_b(o)^{-1}).$$

The curvature map is the real color projection specified in the startup construction, so this expression is defined on every group configuration. Gauge transformations act by the adjoint at o. Define a link-valued current and a scalar tensor by

$$(Y_b(o,U))_{(x,i)} = 1_{\{x=o\}}C_{bi}(o,U), \qquad J_{ab}(o,U) = \tfrac{1}{4}\delta_{ab}\sum_{i,j}\langle C_{ij}, C_{ij}\rangle - \sum_i \langle C_{ai}, C_{bi}\rangle. \tag{TF0a}$$

Color coordinates are converted back to the actual link-coordinate space by the fixed color isometry. At the identity every C vanishes. Its derivative in a logarithmic link direction A is precisely

$$c_{ab}(o,A) = A_a(o) + A_b(o+e_a) - A_a(o+e_b) - A_b(o).$$

Thus $Y(1) = 0$; $J(1) = 0$ and $DJ(1) = 0$. Its Hessian is obtained by replacing a product $\langle C, C'\rangle$ by $\langle c(A), c'(B)\rangle + \langle c(B), c'(A)\rangle$. These formulas supply the actual quadratic jets, rather than an unspecified continuum stress approximation. Initially J is symmetric and has identically zero trace.

The divergence used for the finite Taylor calculation is the fixed high-order antisymmetric stencil

$$\mathscr{D}_a f(o) = \sum_{\ell=1}^{4} d_\ell\{f(o+\ell e_a) - f(o-\ell e_a)\}, \qquad (d_1,d_2,d_3,d_4) = (4/5, -1/5, 4/105, -1/280). \tag{TF0b}$$

It has first moment one and its odd moments three, five and seven vanish. Substitution of the four rational coefficients proves these identities. Taylor expansion therefore reads the ordinary derivative exactly on polynomials of degree at most eight. Physical spacing factors are inserted only when the finite tensor is rescaled.

During a stationary descent retain the full tuple (J, Y, E, Q, Ξ), with exact identities

$$\begin{aligned} \sum_a J_{aa} &= \sum_a \nabla_a Q_a, \qquad J_{ab} - J_{ba} = \sum_i \nabla_i \Xi_{abi}, \\ \sum_a \mathscr{D}_a J_{ab} - E_b &\equiv D_{Y_b}\mathcal{P} \quad \text{through Taylor order four.} \end{aligned} \tag{TF0c}$$

Here ∇ is the unit forward difference, and the last equivalence includes orders zero through four at the identity. Initially $Q = \Xi = 0$, and E is defined as the difference on the last line, so that it is an equality. The Hessian of $\mathcal{P}$ is the retained covariant quadratic form; differentiating $D_Y\mathcal{P}$, using $Y(1) = D\mathcal{P}(1) = 0$, gives its two cross Hessian pairings. The local analytic change of variables preserves these four jets by the chain rule. The scalar rows J, E, Q, Ξ remain analytic, gauge invariant, and zero through order one.

To see why the identities persist under blocking, denote the main averaging stencil by M, the directional flux stencil by F_a, and its correction by T_a. Their finite coefficient identities are

$$F_a = M + \nabla_a T_a, \qquad L^5 M \mathscr{D}_a = \mathscr{D}_a L^4 F_a.$$

Pull each fine scalar function back along the stationary section. Transport J_{ab} with L^4F_a, E_b, Y_b with L^5M, and the difference potentials with their degree-three edge stencils. Add $L^4T_aJ_{aa}$ to the transported Q_a, and add $1_{\{i=a\}}L^4T_aJ_{ab} - 1_{\{i=b\}}L^4T_bJ_{ba}$ to the transported Ξ_{abi}. Substitution proves all three identities in (TF0c); the two added terms cancel precisely the change from the main stencil to the flux stencils. The current is transported by the derivative of the stationary section, so the chain rule also preserves the last identity. No exact symmetry or zero trace of each blocked local representative is presumed.

Localization is performed on the scalar and current functions themselves: gauge fix along the prescribed finite axial paths, discard coordinates beyond the chosen radius, and retain the Taylor polynomial through order four in those local coordinates. Denote the localized tensor by J^{loc}, and localize Q, Ξ by the same operation. Restore its exact trace and its six antisymmetric divergences by defining

$$\begin{aligned} \widetilde{J}_{00} &= \sum_i \nabla_i Q_i^{\mathrm{loc}} - \sum_{i\neq 0} J_{ii}^{\mathrm{loc}}, \\ \widetilde{J}_{ab} &= J_{ab}^{\mathrm{loc}} \quad (a<b), \\ \widetilde{J}_{ba} &= J_{ab}^{\mathrm{loc}} - \sum_i \nabla_i \Xi_{abi}^{\mathrm{loc}} \quad (a<b). \end{aligned} \tag{TF0d}$$

Other diagonal entries are unchanged. The six pairs are $01, 02, 03, 12, 13, 23$; their identities follow individually by subtraction. Summing the four diagonal entries proves the trace identity. The same construction has bounded support, enlarged by one difference step.

The localization error of a differentiated tensor is controlled at the actual stencil sites: (TF0b) uses only shifts of length at most four, so its error is bounded by $\sum_\ell |d_\ell|(A+A)$ if the undifferentiated error is at most A. Summing the four directions gives four times this allowance. Support and anchor weights are enlarged by four at this step. The corresponding derivative and tail allowances are summed with the same absolute coefficients. Exponential path decay pays the discarded tails including their first moments; the stationary four-jet identity pays the local Taylor difference. These are the separate errors retained when the localized divergence is issued as a degree-five insertion.

Finally average the scalar tensor and its divergence over signed permutations and conjugation, with weight $1/|\mathrm{Signed}\times\{0,1\}|$, multiplying each component by its tensor or vector sign. Reindexing this finite average proves reality and the required signed covariance. The vector average kills the identity-configuration constants: the axis reflection reverses the vector sign and pairs identical constants. It does not increase a common rooted bound, because the sum of the absolute averaging weights is one, every transformed chamber has the same rooted seminorm, and the rooted triangle inequality applies. Transport commutes with the same transformations. Thus the issued vector and tensor have the correct exact finite symmetries before any coefficient limit is taken.

We first specify the finite objects whose limits enter the trace argument. This prevents a formal Taylor coefficient from being confused with a continuum insertion. At a fixed finite representative let Y_b be its rooted infinitesimal link-current vector field, $D_Y f$ the directional derivative of a configuration function along its left logarithmic link increments, and $\mathrm{div}_\nu Y$ its divergence for product Haar measure. In exponential coordinates this divergence includes both the ordinary coordinate divergence and the derivative of the logarithm of the Haar density. Write P_A for a finite partial sum of the classical action, L_Q for a finite partial sum of the ordinary logarithmic corrections, and

$$\mathcal{E}_{A,Q} = -\beta P_A - L_Q, \quad \mathcal{R}_b = \beta(B_b - D_{Y_b}P_A). \tag{TF1}$$

Here B_b is the rooted lattice stress expression, with the same mean and curvature rows as in the finite stress construction. The finite sums are taken from the actual source dictionaries, rather than from new representatives of their limiting functions.

The protected configuration change along Y_b, with step $u > 0$, gives a finite difference identity. If $\mathcal{C}_{u,b}$ denotes the resulting density-difference expression, its exact decomposition on the protected healthy region is

$$\frac{\mathcal{C}_{u,b}}{u} = \mathcal{Z}_b - \beta B_b + (\mathrm{div}_\nu\, Y_b)(1), \tag{TF2}$$

$$\mathcal{Z}_b = \mathcal{R}_b + \{\mathrm{div}_\nu\, Y_b - (\mathrm{div}_\nu\, Y_b)(1)\} - D_{Y_b} L_Q + \mathcal{K}_{u,b} + \sum_{\gamma\in\mathcal{F}} \mathcal{H}_{u,b,\gamma}. \tag{TF3}$$

The constant denoted by evaluation at 1 is evaluated at the identity configuration. The fourth term is the finite Taylor remainder of the changed action and Jacobian after subtracting their first directional terms, divided by u. The last sum is the finite correction for the chosen finite shelves of activity terms: each summand is the corresponding changed activity difference with the coefficient and the factor u^{-1} dictated by that difference. Thus (TF3) defines a literal configuration function. To prove (TF2), expand the change of density, distribute $D_Y(-\beta P_A - L_Q)$, and use (TF1); the $\beta D_Y P_A$ terms cancel. Splitting the activity change into the finite retained shelves and their remainder gives exactly the last sum. This is an algebraic equality before any limit or expectation is taken.

The correction shelves can be enlarged to approximate the full changed expression uniformly on configurations. The estimates for rooted activity sums above give nonnegative numbers g_γ with each summand regular on its actual support P_γ, and

$$\sum_{\gamma\in\mathcal{F}} \mathrm{mass}(P_\gamma) g_\gamma \le \varepsilon_{\mathrm{corr}}. \tag{TF4}$$

The finite four-term part of (TF3) is regular on the common neighborhood and has derivative allowance $2c_{\mathrm{small}} + 2c_{\mathrm{loop}}$. The sum in (TF4) is included with its own supports, not paid as if every term had the largest common support.

The conversion of a regular function f into a coefficient insertion is particularly explicit. Let $\Pi_{\le 5} f$ be its Taylor polynomial along the local analytic section at the identity, projected to the degree-five frame, and weight its homogeneous components with the prescribed threshold powers. Call the resulting coefficient $a(f)$. Choose the primary local section $s(a(f))$, and put the masked difference $f - s(a(f))$ into one regular activity slot, with every other slot zero. The resulting formatted insertion equals f on the protected region, has zero extra scalar term, and retains locality and symmetry. If r_{read} is the norm of the finite Taylor reading map and r_{sect} the norm of this section construction, a bound c for the original section gives the two respective bills

$$\mathrm{mass}(P) c (r_{\mathrm{read}} + r_{\mathrm{sect}}), \qquad \mathrm{mass}(P) c ((1 + c_{\mathrm{sect}}) r_{\mathrm{read}} + r_{\mathrm{sect}}). \tag{TF5}$$

These are finite-dimensional operator bounds and the ordinary triangle inequality for $f - s(a(f))$. Add this construction for the four terms and the finite correction sum. Taylor reading is linear, so the actual candidate coefficient is the sum of the four weighted Taylor readings plus the weighted readings of $\mathcal{H}_{u,b,\gamma}$. Equation (TF4) bounds the size of the latter by $(r_{\mathrm{read}} + r_{\mathrm{sect}})\varepsilon_{\mathrm{corr}}$. This supplies the coefficient and its paid remainder together.

The equation identity for this candidate is also literal. Differentiate the finite change of variables in the normalized Haar integral, or integrate its density difference before dividing by u. A spectator supported outside the changed links is unchanged. The resulting expectation is zero. Apply (TF2), subtract its identity-configuration constant, and then let the shelves increase and u follow its admitted small-step schedule. Their uniform allowances justify these passages by dominated convergence. The same identity with a spectator derivative gives the finite lattice Ward contact

term when that spectator meets the support. In particular equation directions have zero separated readings, precisely the equation cancellation already used in constructing the physical quotient.

There is a possible degree-five vector component in this procedure. The raw-to-free vector map has exactly the kernel $\mathbb{R}E$ proved in (FC9)–(FC10). Choose its real signed-equivariant section by the corrected polynomial lift and subtract the unique remaining multiple of E. This is finite-dimensional subtraction before the coefficient limit. It explains why the trace candidate can be compared in fixed vector coordinates despite its original equation component.

The compensation by the quiet logarithmic part is exact at the coefficient level needed here. In logarithmic link coordinates the vector-current diagnostic is a sum, over $d \neq i$, of the quadratic and cubic Taylor coefficients with two equal fixed-color variations in direction i and opposite lattice directions $d, -d$. Its cubic part is the contraction denoted crest_q; denote the quadratic contraction by $\mathrm{suite}_{d,i}\,\mathrm{floft}_{d,i,q}$. Define

$$H_q(X) = \mathrm{crest}_q(X) - \frac{1}{2}\sum_{d\neq i} \mathrm{floft}_{d,i,q}(X_i). \tag{TF6}$$

The quiet logarithmic contribution is constructed with two moment cancellations; gauge invariance alone would not establish the required zero in (TF6). For each fixed ordered pair of plaquette types and color coordinates, let c_γ be the finite two-slot path coefficient and z_γ the displacement of the second slot from the first. Put

$$C_0 = \sum_\gamma c_\gamma, \qquad C_a = \sum_\gamma c_\gamma (z_\gamma)_a. \tag{TF6a}$$

Append a zero-displacement two-slot term with coefficient $-C_0 + \sum_a C_a$, and four terms of displacements e_a with coefficients $-C_a$. Repeat this for every plaquette and color pair. The augmented finite family then satisfies exactly

$$\sum_\gamma c_\gamma = 0, \qquad \sum_\gamma c_\gamma (z_\gamma)_a = 0 \quad (a = 0, 1, 2, 3). \tag{TF6b}$$

Its appended coefficients remain color invariant because that condition is a real linear subspace. The quiet logarithm is the signed and translated sum of this augmented family. The difference between it and the retained ordinary loop sum is precisely the finite appended part. Its coefficient sizes are bounded by the zero and first moments of the omitted path tail; exponential path decay bounds these moments, giving the small remainder already paid in (TF4).

Here is why this construction cancels (TF6). Let $T = L_{\mathrm{quiet}}$. Its Hessian at the identity vanishes when one direction has bounded support and the curl of the other is affine on the enlarged support neighborhood. In fact the Hessian of a two-slot term is the sum of the two products of its linearized curl slots. With the first direction fixed, the other slot evaluated at displacement z_γ is $A + F(z_\gamma)$, with F additive. Summing over the family gives $A\sum_\gamma c_\gamma + \sum_a F(e_a)\sum_\gamma c_\gamma (z_\gamma)_a = 0$ by (TF6b). The same calculation applies to the reversed slot order and each signed translate. The nonwrapping and support hypotheses ensure that these displacements are the actual integer displacements; in the application $2(t + 2 + M) + 6 < n$, where t bounds the current support and M the retained path length.

For $F_b = D_{Y_b} T$, differentiate twice at the identity. The current vanishes there, and the first derivative of a gauge-invariant scalar there is zero. The product rule consequently reduces the Hessian to

$$D^2 F_b(0)[X, Z] = D^2 T(0)[DY_b(0)X, Z] + D^2 T(0)[DY_b(0)Z, X]. \tag{TF6c}$$

The directions $DY_b(0)X$ have the bounded support of the current. Choose X, Z to be the polynomial link directions of degrees at most two that occur in the coefficient test. Their curls are affine on that neighborhood. The preceding Hessian cancellation makes both terms in (TF6c) zero. Terms with a constant pure-gauge direction vanish by the differentiated gauge identity. Expanding the cubic coefficient in the two position variables of the test uses just these Hessian pairings, so its centered-link reading is zero.

One must still convert the centered-link reading into the actual link reading. If the two position arguments are z_0, z_1, the change is multiplication by

$$e^{\,i((z_0)_i+(z_1)_i)/2}.$$

The constant and first coefficients of the original reading vanish. The normalized third coefficient of this product is therefore its original third coefficient plus $i((z_0)_i+(z_1)_i)/2$ times its second coefficient. Apply the finite coefficient functional in the d, b variables. The latter term is zero unless $i = b$; when $i = b$, its value is $-\frac{1}{2}\,\mathrm{floft}_{d,i,q}$. This follows by expanding the quadratic coefficient on $(e_d, -e_d)$ and using $i^2 = -1$. Summing over b and $d \neq i$ gives exactly (TF6). Its centered-link version has just been shown zero, so $H_q(D_Y L_{\mathrm{quiet}}) = 0$. The factor $1/2$ thus comes from this link-centering phase and the normalized Taylor coefficients. Subtracting the paid appended part gives the stated bound for the ordinary loop contribution. No quiet contribution is left undetermined.

The needed vanishing of the quadratic part can be derived without setting a mass regulator equal to zero inside an integral. For $0 < \mu \leq 1$, the local resolvent estimates below give, for the massive comparison expression L_μ,

$$|\,\mathrm{cob}_{d,i,q}(L_\mu)| \leq C\mu^{-1}. \tag{TF7}$$

The proof differentiates the Fourier-resolvent expression twice; its pointwise integrand is bounded by $C(|p| + \mu)^{-5}$. Counting lattice shells and multiplying by the momentum cell volume gives $C\mu^{-1}$ in four dimensions. This applies to both the physical and comparison resolvents, hence to their sum. At each fixed positive μ, the resolvent identity identifies the unregularized current coefficient with $\mu^2\,\mathrm{cob}_{d,i,q}(L_\mu)$, up to an error tending to zero in the retained large-volume limit. Given $\epsilon > 0$, choose $\mu = \min(1, \epsilon/[2(C+1)])$ first. The massive term has magnitude at most $C\mu \leq \epsilon/2$; take the volume large enough to make the remaining error below $\epsilon/2$. The unregularized quadratic coefficient therefore tends to zero. The Taylor remainder bounds give the same limit for its finite-difference reading. Summing the finitely many d, i shows that (TF6) has the same limit as its cubic part.

That cubic limit is $-12\,s_G$, where s_G is the one-loop slope constant in the coupling drift. Here is how the logarithmic slope fixes it. For the regulated determinant expression $A(\log\mu)$, the determinant estimates give

$$|A(x) + (-12 s_G)x| \leq C, \qquad x \leq 0. \tag{TF8}$$

The regulated response is $-A'(x)$, with $x = \log\mu$. The real mean value theorem on an interval $[a, b] \subset (-\infty, 0]$ supplies an intermediate regulator at which this response differs from $-12 s_G$ by at most $2C/(b-a)$. First choose $b - a$ large, then use the local resolvent difference estimate to compare that response with the unregularized cubic current coefficient; this error tends to zero when the upper endpoint b is sufficiently negative and the retained volume is sufficiently large. The intermediate regulators stay in the fixed positive interval $[e^a, e^b]$, so the large-volume estimates apply uniformly there. The triangle inequality proves the limit. Together with (TF7), it proves $H_q \to -12 s_G$ for every color q. These are finite rooted Haar-Jacobian coefficient tests, not an assumption about a mollified continuum field.

Finally the passage through many adjacent coefficient transports retains a relative error. Put

$$\omega_i = \exp\{-\tfrac{2}{5}\ \mathrm{RLog}(\widehat{\beta}_i)\} \longrightarrow 0, \qquad k = k(i) > i. \tag{TF9}$$

The chosen calendar ensures every adjacent transport error between i and k is at most $\omega_i/[2(k-i)]$ times the next coefficient size. If v_j denotes that coefficient, backward induction proves

$$\|v_j - v_k\| \le \frac{\omega_i(k-j)}{k-i}\|v_k\| \le \omega_i\|v_k\|, \qquad i \le j \le k. \tag{TF10}$$

Indeed the induction hypothesis and $\omega_i \le 1$ imply $\|v_{j+1}\| \le 2\|v_k\|$; adding the next adjacent error gives the asserted bound at j. The calendar inequality follows from the positive reserve $-2/5 + 1/2 = 1/10$: logarithms and the number of intervening steps are eventually dominated by the exponential with exponent $1/10$. Thus this accumulated relative error vanishes. The finite candidate, its equation identity, the quiet compensation and its scalar diagnostic consequently survive on the same retained field and enter the following trace comparison.

16.2 The interacting-to-Gaussian comparison used for the trace

The radius selection is exact as a geometric statement: the requested radius s equals a purchased station radius $\rho_k(b+u)$ for some $0 \le u \le 1$. This retains every comparison estimate bought on that interval. For the actual moments the conclusion remains approximate, for example $|u_{01} - (m_0m_1 + z)| \le (1 + |m_0| + |m_1|)c$. It becomes exact in the specified limit as $c \to 0$. This distinction is needed when arbitrary sequences of radii are brought back to the same physical table and the same selected phase.

This comparison concerns the interacting Wilson family and its selected continuum phase. It is distinct from the calculation of a Gaussian constant coefficient in the preceding chapter. We derive the passage between them.

At a fixed retained comparison station take two physical lattice insertions A_0, A_1, supported in its strict chamber, with their coefficient representatives converging along the common comparison filter. Let m_0, m_1 be the corresponding one-point Gaussian readings and let z be the centered Gaussian two-point reading. The corrected singleton and pair polynomials are

$$D_j = A_j - m_j, \qquad D_{01} = (A_0 - m_0)(A_1 - m_1) - z. \tag{IG1}$$

We next derive the finite interacting bound for these functions. This also specifies which Gaussian quantities are subtracted. For n marked insertions let f be the cumulant polynomial of the Gaussian profile at the rounded positions. The corrected cylinder is the coefficient polynomial obtained by multiplying the insertion generating polynomial by $\exp(-f)$, in the finite algebra in which each marker has square zero. Its singleton coefficients are $A_j - m_j$, and its pair coefficient is exactly D_{01} in (IG1). Thus the subtraction includes the Gaussian means as well as its connected pair coefficient.

Let k be a comparison station, $M \ge n$ its fixed maximum number of marks per chamber, and H_k its monitor. Define

$$I_k = \frac{\log H_k}{16M}, \qquad R_k = e^{2I_k}, \qquad \varepsilon_k = e^{-I_k}. \tag{IG4}$$

Here $H_k \to \infty$, so $\varepsilon_k \to 0$. The initial marked density is the Wilson density times the product of the reflected copies of the corrected cylinder over all chambers of the tiling. Push this density along the same exact renormalization course as its unmarked Wilson density. In particular its vacuum integral is the original Wilson partition function. The comparison uses this density identity, not a replacement of the interacting law by a Gaussian law.

We describe the bounds on this push. At the middle stop its regular marked sources have values $W_\alpha(U)$, with specified fired marker sets and support cubes. They occur together with the forced

sources and the unmarked background in the finite density expansion described in the marked-integration chapter. If C_α is the corresponding coefficient of the negative Gaussian cumulant, then on the healthy local chart

$$|W_\alpha(U) + C_\alpha| \le b_{-31/32}(t), \qquad |C_\alpha| \le b_{1/4}(t). \tag{IG5}$$

Here t is the station coupling, $b_a(t) = \exp(a\,\mathrm{boil}(t))$, and $\mathrm{boil}(t) \to \infty$. To obtain the first bound, expand the actual insertion product by its finite cumulant expansion. The local Gaussian comparison of its calibrated chart coefficients gives an error at most $b_{-3/2}(t)$ for each coefficient; the finite exponential map from cumulants to moments is continuous on this fixed finite coefficient space. Its chosen neighborhood makes the cumulant-polynomial discrepancy plus that error at most $b_{-31/32}(t)$. The exact negative cumulant coefficient cancels the corresponding free coefficient. The second bound follows because the rounded Gaussian profile stays in a compact subset of the generic position domain, where its finitely many coefficients are bounded, while $b_{1/4}(t) \to \infty$.

Here is the polynomial estimate behind the local Gaussian comparison just used. Denote coefficient ℓ^1 size by coin, and let k_0 bound the degree of each insertion polynomial. After the actual linear substitution into the common key and residual coordinates, write these polynomials as P_i; write F_i for their calibrated free polynomials, included in the same coordinates. The calibration is the exact identity that their weighted truncations through the assigned degree d_i agree. If the substitution has weighted size at most b, the original polynomial has size at most a_i, and the weight radius is $\rho \ge 1$, every remaining monomial has weight at least ρ^{d_i+1}. Hence

$$\mathrm{coin}(P_i - F_i) \le a_i b^{k_0} \rho^{-(d_i+1)} \le \xi. \tag{IG8}$$

This is the high-degree tail estimate; it uses the actual solved calibration maps. If $\mathrm{coin}(F_i) \le B_0$, $B_0 \ge 1$, telescoping a product over any subset J of the n insertions gives

$$\mathrm{coin}\left(\prod_{i\in J} P_i - \prod_{i\in J} F_i\right) \le |J|\xi(B_0+\xi)^{|J|} \le n\xi(B_0+\xi)^n. \tag{IG9}$$

Substitution of the Gaussian variables commutes exactly with their Wick functional: the substituted covariance is the pullback of the original covariance and its mean is the pullback of the original mean. These are equalities of finite Gaussian moments, including the residual coordinates, not limiting approximations.

Let all means and covariances have absolute value at most $B \ge 1$, and let E bound their discrepancies from the zero-mean continuum free data in the key coordinates. Put $N = nk_0$ and $C_{B,N} = [2B(N+1)]^N$. Wick's recursion bounds the functional on a degree-N polynomial P by $C_{B,N}\,\mathrm{coin}(P)$, and its change under these mean and covariance discrepancies by $EC_{B,N}\,\mathrm{coin}(P)$. One proves both assertions by induction on the monomial degree: remove one variable, pair it with each remaining variable or replace it by its mean, and, for the difference, telescope the resulting products. The at most $N+1$ choices at each step and the bound B fit the displayed constant.

Combining this bound with (IG9), and then comparing the free-key polynomials under the two Gaussian data, gives total error at most

$$C_{B,N}\big[n\xi(B_0+\xi)^n + B_0^n E\big]. \tag{IG10}$$

The common covariance estimates and small-mean estimates in the Gaussian comparison chapter give precisely these bounds on the pulled-back data, uniformly over the allowed small middle-stop means. The station choices make (IG10) at most $b_{-3/2}(t)$. Thus the estimate feeding (IG5) follows

from the actual covariance comparison, exact calibration and a finite polynomial calculation. It is not an additional interacting-to-free hypothesis.

These coefficient estimates are weighted in the density, as follows. A regular source in a support cube of radius QB has support charge at most $\exp(e(2Q+1)^4)$, since that cube occupies at most $(2Q+1)^4$ budget blocks. Its marker weight at radius R_k is at most $R_k^M = b_{1/8}(t)$, by the identity between the monitor and the station coupling. Consequently multiplication of (IG5) by the support charge and marker weight fits the prescribed regular-source bound. The same calculation with the additional unit of support charge fits the Gaussian compensation source bound. The forced sources already satisfy their full-configuration bounds. The one-step marked integration estimate therefore propagates this compensated density from the middle stop down to the width stop, preserving the unmarked background and bounding the total source weight per marker and the total forced weight per site by 1. This is finite induction over the stops: the initial insertion estimates give the terminal case, and at each earlier stop the chart transport estimate and its forced-source estimate give the next case. The delay conditions put every coupling in this finite induction in the same interval $[t/2, t]$, so the chosen budgets apply throughout.

For completeness, the final conversion of these density bounds to a rooted Wilson bound can be written explicitly. Let Z be the number of width-stop sites, which is also the number of reflected chambers. If $q \neq \varnothing$ is a fired subset of the n markers, the weighted coefficient estimate for the pushed density and the positive partition-function lower bound give

$$\|\langle \prod_{\text{chambers}} \text{copy } D_q \rangle_{\text{Wilson}}\| R_k^{Z|q|} \leq \exp\{ZP - 4\dim \mathfrak{g} \log g\}. \tag{IG6}$$

The parameter $g > 0$ here is the width-stop small-field parameter, with $\beta g^2 = 1$; P is its per-site protection bound with source count at most $1 + n \leq 1 + M$. The partition estimate applies to the actual positive course density. Its comparison with the reference density uses the two-sided exponent error bounded by $Z\delta$: one side bounds the numerator and the other side bounds the positive denominator, so their errors are both included in P.

Taking the Z-th root in (IG6) is legitimate because $Z > 0$. The width-stop allocation was chosen so that

$$P - \frac{4\dim \mathfrak{g} \log g}{Z} \leq I_k.$$

The rooted chamber seminorm therefore obeys

$$\|D_q\|_{\text{root}} \leq \frac{e^{I_k}}{R_k^{|q|}} \leq \frac{e^{I_k}}{R_k} = e^{-I_k} = \varepsilon_k. \tag{IG7}$$

The middle inequality uses $|q| \geq 1$ and $R_k \geq 1$. The allocation is eventually available because the finite source-count ticket at the width stop is bounded by I_k; this is the delayed station choice, made simultaneously with the finitely many chart, insertion and transport inequalities. Applied to the two singleton modes and the pair mode, (IG7) is the asserted common bound $c = \varepsilon_k$ for (IG1). These are bounds for the actual Wilson law, before any character is selected.

Here is why this gives the same bound for the character reading of each corresponding continuum polynomial. Let D be any finite linear combination of distinct-point words with the rooted bound c, and let B be its compressed whole-word operator. For every integer $n \geq 1$, place $2n$ copies of D, with alternating reflected conjugation, in chambers separated by a distance parameter t. At sufficiently large volume these are distinct chambers. The finite chessboard inequality bounds the absolute expectation of their product by c^{2n}. At fixed n, t, the ordered lattice limit preserves this inequality; the comparison filter lies below that same ordered-limit filter and the coefficient

representatives converge there. Then let $t \to \infty$. The long-word compression limit, expanding the finite linear combinations without commuting their operator factors, gives

$$\langle \Omega_0, (B^*B)^n \Omega_0 \rangle \le c^{2n}.$$

The character decomposition therefore gives

$$\int |\chi(B)|^{2n} \, d\mu(\chi) \le c^{2n}. \tag{IG2}$$

It follows that $|\chi(B)| \le c$ for every character. Indeed, if $|\chi_0(B)| > c$, continuity provides an open neighborhood of χ_0, a number $a > c$, and a positive measure h of that neighborhood, on which $|\chi(B)| \ge a$. Full support gives $h > 0$. Then (IG2) implies $ha^{2n} \le c^{2n}$ for every n, an impossibility as $n \to \infty$. This includes $c = 0$. The two limits here have their stated order: first the lattice limit at fixed copy count and separation, then separation to infinity, then the integer $n \to \infty$.

Let u_j and u_{01} denote the continuum one- and two-point readings in any one character. Applying this argument to (IG1), and using the unit reading 1, gives

$$|u_j - m_j| \le c, \qquad |u_{01} - m_0 u_1 - m_1 u_0 + m_0 m_1 - z| \le c.$$

The exact decomposition

$$u_{01} - (m_0 m_1 + z) = [u_{01} - m_0 u_1 - m_1 u_0 + m_0 m_1 - z] + m_0(u_1 - m_1) + m_1(u_0 - m_0)$$

proves

$$|u_{01} - (m_0 m_1 + z)| \le [1 + |m_0| + |m_1|]c. \tag{IG3}$$

Thus the Gaussian comparison follows from three specified interacting rooted estimates, not from an assertion that the interacting measure converges to a Gaussian measure.

At comparison index i, the station construction supplies these rooted estimates for every pair of basis labels through degree i, with $c = \epsilon_i \to 0$, uniformly over its fixed position window. The Gaussian readings m_j, z are the singleton and pair coefficients of the Gaussian moment profile. In particular $m_0 m_1 + z$ is its full two-point coefficient, by the defining cumulant identity. Write $Q_i(x)$ for the rounded position tuple and $r_i Q_i(x)$ for the actual physical tuple. If x is generic and has a neighborhood in the comparison window, then $Q_i(x) \to x$, so it is eventually generic and stays in that window. Smoothness of the free hierarchy implies convergence, and hence boundedness, of its two singleton readings at $Q_i(x)$. Equation (IG3) yields convergence of the interacting two-point reading to the free two-point reading at x.

The finite-cutoff representatives of a free label converge, along the retained ultrafilter, to a linear lift $Y_{d,i}$. Projecting that lift gives the continuum map $U_{d,i}$. The finite station bounds pass to these limits because the lattice moments with converging finite-dimensional coefficient labels have the calibrated ordered limits. The station construction also preserves the unit, conjugation and signed motions. The free detector proves eventual injectivity of $U_{d,i}$; the field-dimension bound gives surjectivity, so these maps are eventually isomorphisms at each fixed degree. Replacing finitely many exceptional initial maps by any fixed isomorphism leaves all limits unchanged.

The conclusion also holds for convergent moving free labels $v_{i,0}, v_{i,1}$. To see this without assuming a uniform norm bound on all interacting bilinear forms, fix a basis e_a of the degree-d free space and expand

$$B_i(v_{i,0}, v_{i,1}) = \sum_{a,b} v_{i,0}^a v_{i,1}^b B_i(e_a, e_b).$$

Each basis-pair value converges by (IG3), and every displayed coefficient converges. There are finitely many summands. Their limits are therefore exactly the free bilinear form evaluated at the limiting labels. This is the moving-label comparison used for the normalized centered trace.

We finally justify the universal radius selection used in (TLim1). Given an arbitrary positive sequence $s_n \to 0$, a fixed finite degree horizon, any lower bound on the sequence index and station depth, and any positive radius ceiling, the station construction supplies a later sequence member below that ceiling whose value is **exactly** the station's scale radius. Its purchased comparison bounds include the three modes in (IG1), the bounded coefficient lifts and the smeared comparison estimates through that horizon. All their error thresholds can be met together because only finitely many basis pairs and modes occur at this horizon. Recursively increase the horizon to i, require a larger sequence index and station depth, and set the next radius ceiling equal to half the radius just chosen. This produces a strictly increasing subsequence map, strictly decreasing chosen radii, and station depths tending to infinity. At each fixed degree all its comparisons hold eventually.

The exact radius matching has an elementary interval argument. Write $M_k \geq 1$ for the calendar's magnification monitor and $\rho_k(b) = a_k M_k^b$ for the station gap. Eventually $M_k \leq M_{k+1}$, and the purchased block bound gives $L_k \leq M_{k+1}$. Thus, for $b \geq 0$,

$$\rho_k(b) = a_{k+1} L_k M_k^b \leq a_{k+1} M_{k+1}^{b+1} = \rho_{k+1}(b+1).$$

Also $\rho_k(b) \to 0$ for each fixed b. Given a lower station index J, choose $i \geq J$ beyond these eventual bounds. For $0 < s \leq \rho_i(b)$, take the first $k > i$ with $\rho_k(b) \leq s$. Minimality gives $s \leq \rho_{k-1}(b) \leq \rho_k(b+1)$, also when $k = i+1$. The continuous function $u \mapsto \rho_k(b+u)$, $0 \leq u \leq 1$, therefore attains s. The station family was purchased for this entire interval of exponents; its other charges and interpolation parameters stay fixed over the interval. Consequently this adjustment of u preserves its estimates. Given a sequence $s_n \to 0$, choose a later sequence index below the radius ceiling supplied by this argument, and repeat with a larger J. This proves the asserted simultaneous escape of sequence and station indices, with exact equality of radii.

The bounded lifts are then extracted along the already retained ultrafilter and projected as above. Their finite-dimensional limits assemble into one scaling family with exactly those chosen radii. The trace and local-coordinate estimates apply to every such compatible scaling family; they retain the previously fixed physical table and phase. Hence this construction works for the arbitrary sequence supplied to it, including the violating sequence in the proof of the full trace limit.

16.3 One calibration matrix for the three panels

The three comparison panels used in (WS1)–(WS17) differ in their point and smear geometry, but have the same finite Taylor order and the same calibration test system. We justify the common matrix, rather than choosing three unrelated solutions. At a fixed degree and sufficiently late stop, the detector matrix of that calibration system is invertible by (PC1)–(PC13). Select the solved chart at one anchor panel. Its defining equations prescribe its pairings with that test system. For any other panel, transport the same local representative and the same Taylor card to its allowed geometry. The support, threshold, and budget equalities make its equations identical to the anchor equations. Uniqueness of their solution then gives the same field matrix T at that panel. Signed covariance follows by transporting the solved equations under a signed permutation and again invoking uniqueness. This proves simultaneously the tested-chart and signed-intertwining properties needed in (WS1).

More concretely, the inverse chart is T^{-1} after the field dilation and its engineering-weight factor. Applying the defining calibration equations to a basis proves that it solves those equations; the unit equation gives the unit, and conjugating or turning the equations gives conjugation and signed

covariance. Finite dimensionality bounds the finitely many inverse coefficients after the detector lower bound has been paid. These are precisely the coefficients used in the rooted comparison estimate, including on the smear panel. There is no separate choice of an inverse for each spectator word.

For a fixed finite horizon each panel has finitely many local demands, all eventually paid. Take their maximum together with the common detector and monitoring thresholds. At this stage the anchor solution transfers to all three panels at every finer cutoff and every admitted volume. Choose the stops by the cofinal advance already used in (WS16), satisfying these finite maxima and each earlier horizon. Passage to the common coefficient subsequence preserves the tested equations because their finite pairings converge. Therefore the point, displaced-point and smeared comparisons have the same limiting field identification and may be combined in the trace and Ward arguments.

16.4 Common scaling maps and convergence of every derivative

The comparison of a physical field with a free field must use one map for all its readings. Here is the construction, including the compactness and interpolation steps. Fix a finite maximum number of slots, a finite field horizon and a bounded separated configuration window. During this construction these bounds may be enlarged to include detectors and derivatives. The physical table, its null quotient and its limiting correlations remain those already constructed.

Choose radii $r_k \downarrow 0$ and lattice levels n_k with

$$\delta_k = \frac{a_{n_k}}{r_k} \longrightarrow 0. \tag{WS1}$$

For each free coefficient basis vector at each requested degree, the insertion construction gives a column of physical reference coefficients. Its bound is finite for each fixed k, d, uniformly in the late mesh and large period. These bounds are not asserted uniform in k. Extract the finitely many coordinates of each column on their compact balls, first in the mesh limit at a retained period and then in the period limit. Diagonalize over the countable pairs (k, d). The resulting maps $U_{k,d} : F_d^{\mathrm{free}} \to F_d$ initially need not be invertible or commute exactly with the inclusions. Reality, the signed permutation action and the unit pass to the limit since they are finite linear equalities.

The estimates that permit this extraction also supply numbers $A_k \geq 0$, for every fixed word length n, with

$$\delta_k A_k^n \longrightarrow 0. \tag{WS2}$$

The center and comparison estimates are imposed at levels $m \leq j$ satisfying

$$a_m < r_k/Q, \qquad 6na_j < a_m, \qquad a_j < r_k/(4Q).$$

The insertion supports are therefore separated before the comparison is taken. On the rescaled lattice, the aligned comparison tends to the free correlation. On a configuration window in which all the required coordinate projections have a fixed strict gap, the long finite-difference estimate gives

$$|S(U_k v; r_k x)| \leq A_k^n, \qquad |\partial_{x_{\ell,\alpha}} S(U_k v; r_k x)| \leq 4QA_k^n. \tag{WS3}$$

Rounding to the rescaled lattice changes a coordinate by $O(\delta_k)$; the resulting error vanishes by (WS2). Thus the correlations converge to the free ones on these coordinate-generic windows. All bounds here are for a fixed finite word and fixed window. No uniform bound on all fields or all degrees is being taken.

To fit the coefficient maps, use compactly supported smooth tests in the coordinate-generic part of the separated configuration space. Their readings separate the free quotient. Indeed, a field whose every such reading vanishes has zero correlations there by the fundamental lemma for test functions, and then on the entire separated space by continuity and density. Finite dimension now gives finitely many separating readings at each receiving degree. Spectators in these readings can have larger degree than the field being detected; enlarge the common horizon accordingly. If E_d denotes this finite injective reading map, the physical readings of $U_{k,d}$ converge, in matrix norm, to E_d. Injectivity is open for maps on a finite-dimensional domain: choose a left inverse L_d of E_d; when $\|L_d(E_{k,d} - E_d)\| < 1$, the matrix $L_d E_{k,d}$ is invertible. Thus $U_{k,d}$ is eventually injective. The previously constructed spanning family gives $\dim F_d \le \dim F_d^{\mathrm{free}}$, so it is eventually an isomorphism.

We make explicit the finite-horizon estimates preceding these readings. A centered budget is the reflection-positive seminorm for every signed or conjugated copy of the one insertion in a reporting chamber. For a word of n insertions, take n chambers whose centers may be shifted within a finer common chamber. If distinct points have at least one coordinate difference larger than r_k/Q, the inequalities preceding (WS3) permit these shifts so the inspected insertion cells are disjoint. Apply Cauchy–Schwarz in the positive reflection form successively in the separating coordinate directions. Each insertion contributes its centered seminorm at most A_k; no factor depends on the volume or the microscopic cutoff. Thus the finite Wilson word has magnitude at most A_k^n.

Take its ordered mesh and period limits at a coordinate-generic tuple. The fitted coefficients converge in their finite-dimensional balls, and multilinearity permits them to vary along the limit. The same bound passes to the limiting correlation. On the open set of tuples with some separating coordinate for each pair, the correlation is continuous and the generic tuples have full measure and are dense. Therefore the bound extends to every tuple in that open set. This proves the first part of (WS3) on the entire separated window, including coordinate ties.

For a coordinate-generic tuple with all coordinate gaps larger than $1/Q$, move one coordinate a positive rescaled distance $t < (4Q)^{-1}$. On the microscopic lattice perform the integer displacement $\lfloor r_k t/a_N \rfloor$ in that coordinate. Its physical length tends to $r_k t$. The long finite-difference estimate for the same centered budgets, with available gap $r_k/(4Q)$, bounds the difference of the two Wilson words by

$$A_k^n \frac{r_k t}{r_k/(4Q)} = 4Q A_k^n t. \tag{WS8}$$

Both endpoint readings converge with their original placements and coefficients. Pass through the ordered limits, divide by t, and let $t \downarrow 0$ through endpoints that remain generic. Smoothness identifies the limit with the coordinate derivative and proves the second bound in (WS3). This estimate is uniform on a fixed strict coordinate-gap window. It is not used with Q tending to infinity near a projection tie.

The small quantity in (WS2) has the following precise use. The finite suppliers provide it separately for every word in a fixed finite-degree basis and all its signed and conjugated variants. Taking a maximum over finitely many words preserves convergence to zero. Indeed, if $a_{k,w} \ge 0$ and $\delta_k a_{k,w}^n \to 0$ for each of finitely many w, then $\delta_k \max_w a_{k,w}^n = \max_w(\delta_k a_{k,w}^n) \to 0$. The polynomial prices of the solved columns are (PC12)–(PC13); the rooted insertion and aligned Gaussian estimates provide these centered budgets and the aligned word comparison on the same station. The finite estimates must supply these products at their actual station radius and grid; the grid cannot be altered independently after a comparison has been purchased. Thus varying labels in a bounded finite-dimensional set costs a fixed finite sum of these basis bounds, rather than requiring uniformity over an infinite set of labels.

Here is the scale and price bookkeeping in that last statement. For each requested station k,

all degrees through the horizon use one earlier index $a(k) \le k$, with

$$a(k) \longrightarrow \infty, \qquad s_k - s_{a(k)} \le \log \widehat{\beta}_k, \qquad \frac{a_{a(k)}}{a_k} \le \widehat{\beta}_k. \tag{WS11}$$

The last inequality follows by exponentiating the preceding one, since $a_i = e^{-s_i}$. At the comparison radius r_k the actual rescaled mesh is $a_{a(k)}/r_k = \mathrm{mag}_{a(k)}^{-1}$. The common finite matrices are taken to sublimits along the same comparison filter below the retained ordered filter. Each pair of degrees, each spectator and each placement therefore uses these same matrices and this same radius.

The reverse coefficient transport from a level m to $k \ge m$ has price

$$\mathrm{quota}_k(v) \le (a_m/a_k)\, \mathrm{quota}_m(v). \tag{WS12}$$

Indeed one step costs the column allowance $R_{i+1} \le L_i$; multiplying those inequalities gives $\prod_{i=m}^{k-1} L_i = a_m/a_k$. The solved readout at $a(k)$, followed by its exact field-number and engineering dilations, consequently supplies a preimage y for every actual coefficient v, with

$$\|y\| \le 2a_k^{-d}\, \mathrm{quota}_k(v).$$

To check the factor two, its approximate matrix differs from the exact awake coefficient by at most half the awake coefficient size, as in (PC7); the triangle inequality bounds the approximate size by twice the awake size. The field shrink is at most one and the comparison magnification at least one, so undoing their positive dilations cannot increase the weighted coefficient size. Formula (WS12) then gives the displayed bound. This argument applies to the full coefficient space, including equation directions, before restricting to the physical quotient.

The magnification and each fixed finite sum of its engineering and field-number powers have rate zero relative to $\log \widehat{\beta}_k$. For example the label price is bounded by

$$2 \sum_{\alpha} (2\, \mathrm{monitor}_{a(k)})^{f_\alpha} \mathrm{mag}_{a(k)}^{r_\alpha}.$$

For every $\varepsilon > 0$, each fixed power is eventually at most $\widehat{\beta}_{a(k)}^{\varepsilon} \le \widehat{\beta}_k^{\varepsilon}$; the finite sum has the same rate-zero property. These are the actual finite-horizon prices used with the rooted insertion estimates. The centered estimates assert jointly, for a word of length n, the existence of reporting levels $m \le j$ such that

$$a_m < r_k/Q, \quad n(6a_j) < a_m, \quad a_j < r_k/(4Q), \quad A_k^n(a_{a(k)}/r_k) \to 0. \tag{WS13}$$

The first three conditions give the movable disjoint chambers used in (WS8); the last is exactly the price needed in (WS9). They are imposed together with the aligned moment comparison, not inferred merely from the two-field detector limit. In passing from a basis to bounded arbitrary labels, expand every slot in the fixed finite basis and sum the products of the coordinate norms. The finite sum preserves both the vanishing comparison error and the last limit in (WS13).

The smeared station version says the same thing without committing to a grid interpolation: for every smooth test t in the stated separated window, with $\|t\|_\infty \le M$, the ordered upper error is at most $\varepsilon M \prod_\ell \|v_\ell\|$. Its finite-cutoff lift is bounded, real, signed-equivariant, and sends the unit to the actual reference unit. Its scale and normalization are fixed for all those tests. Enlarging the arity, window and degree and decreasing the error preserves the previous requests. The common-order selection first retains these finite data, then uses (WS4) and finite detecting matrices to obtain all derivative estimates. Thus neither the conditional formulation of a station estimate nor the existence of a two-point comparison is being used as a substitute for the joint estimates (WS13).

The simultaneous centered supplier can be described explicitly; it is stronger than a two-point Gram comparison. Fix a degree cap D, an arity cap $n \geq 2$, and a compact separated window. The actual construction provides three calibrated panels at exponents a, b, e, on one family and with the same coefficient matrices. The bounds used for a panel are its solved-column and marked insertion bounds (PC1)–(PC13), together with the rooted comparison of the entire requested finite word. Denote the common monitor by M_k, and write $H_k(t) = M_k^t$. In the source this is $\mathrm{bill}(t, \operatorname{clim} k)$; thus $H_k(t)H_k(u) = H_k(t+u)$. A primary comparison radius is $r_k = a_k H_k(a)$, its mesh spacing is $\delta_k = H_k(-a)$, and a small motion allowance is $p_k = H_k(c)$. Choose the exponents so that

$$c < 0, \quad b + 6 < a, \quad e + 5 < a + c, \quad (a-b)Dn + c < 0, \quad (a-e)Dn - a - c < 0. \tag{WS14}$$

There is no obstruction to these finite numerical choices. Put $x = 7Dn + 1$, $y = x + 6$, $a = 17 + D + x + yDn + y$, $b = a - 7$, $e = a - y$, and $c = -x$. Then $7Dn - x = -1$, $e + 5 = a - x - 1 < a + c$, and $yDn - a + x = -17 - D - y < 0$; also all three panel exponents exceed $D + 6$. These are the source's finite choices. The calibrated panel supplier is chosen with these exponents before extraction.

Here is how that supplier yields centered bounds, rather than merely ordinary moment bounds. For a single basis label, the rooted comparison applied to the singleton at the origin controls the observable minus its free scalar value by the word error $\varepsilon_k \to 0$. The centered root seminorm is subadditive, is homogeneous in scalar coefficients, and assigns norm one to the constant observable. Thus its full observable has root seminorm at most $1 + |r_{\mathrm{free}}|$, eventually. Interior support in the centered chamber holds at every translated origin by locality of the observable; translation of the chamber preserves this root seminorm. This gives a uniform centered bound at the panel's own width, including its constant part.

Changing the panel magnification from b to a multiplies a coefficient of engineering degree $q \leq D$ by $H_k(a-b)^q$. Expand in the fixed finite label basis. The triangle inequality therefore gives the centered price

$$A_k = C\, H_k((a-b)D). \tag{WS15}$$

The same argument at panel e gives $A'_k = C' H_k((a-e)D)$. To include every signed and conjugate variant, first apply these finite transformations to the free label, expand each in the same basis, and take the sum of their finitely many coordinate prices. The coefficient lift intertwines these transformations and sends the unit to the actual reference unit, so the resulting bound is for the required physical transformed token. All slots of the word are included before taking the maximum constant. Thus one reporting width works simultaneously for the word; it is not selected after seeing its offsets.

The available plateau geometry gives, for j_b the reporting width at panel b,

$$n(6a_{j_b}) < a_{j_b - 1}, \qquad \frac{a_{j_b - 1}}{a_k} \leq H_k(b+6), \qquad \frac{a_{j_e}}{a_k} \leq H_k(e+5). \tag{WS16}$$

The first follows because one blocking factor tends to infinity. For the other two, the principal plateau length costs $H_k(b)$ or $H_k(e)$; the rounded length and two buffer lengths cost at most five further monitor powers, and passing to the preceding block costs one more. The block factor is bounded by the monitor and its tenth power dominates the fixed covering-size buffers. Multiplying these inequalities gives exactly the exponents $b+6$ and $e+5$, rather than changing the comparison radius.

Equations (WS14)–(WS16) now give the simultaneous vanishing products needed in the moving-chamber proof:

$$A_k^n p_k \longrightarrow 0, \qquad (A'_k)^n \delta_k / p_k \longrightarrow 0, \qquad 4n\delta_k \leq p_k \quad \text{eventually.} \tag{WS17}$$

Indeed their first two exponents are the last two negative numbers in (WS14); the third uses $-a < c$, which follows from the last inequality there. Also $a_{j_b-1} < r_k/(2R)$ and $a_{j_e} < r_k p_k$ eventually, by the two strict panel gaps in (WS14). These two widths are the two centered budgets needed for the shift across a grid tie and for the comparison away from it. The rooted whole-word error is $C_R\varepsilon_k \to 0$. The moving-chamber decomposition (WS8)–(WS9), with its strip of width p_k, is therefore valid jointly in all slots. Integrating it against a test of supremum M gives the station error $\varepsilon M \prod_i \|v_i\|$, including projection ties. This is the actual three-panel route to the joint station statement. The simpler conditions (WS13) record the part of that geometry used by its one-width formulation; no two-field estimate is being used to supply a whole-word centered estimate.

For clarity about existence on the chosen trajectory, the dense calibrated supplies retain fixed three-panel exponents at each horizon, the foregoing inequalities, and the actual solved charts at a predecessor $g(k) \to \infty$ satisfying (WS11). Their availability is paid by the same delayed local insertion and marked comparison schedule used for (PC1)–(PC13). The finite horizon does not depend on the requested extraction order. Passing their finite matrices to limits on a fixed common subfilter preserves those charts and gives the smeared comparisons for their continuum fields. At each horizon, the finite-dimensional derivative and inclusion adjustments in (WS4)–(WS7) act on these fields. Finally choose an increasing integer horizon whose every finite requirement has already become valid at index k: if K_D is a threshold for horizon D, use the largest $D \le k$ with $\max_{j\le D} K_j \le k$. It tends to infinity. This is the common diagonal; it preserves the physical index k, the fixed family, and the inherited subfilter. It does not retune the grid for individual spectators.

These same readings control the inclusions and derivatives without requiring a bound on the inverse in the physical coefficient norm. In a rescaled position variable, integration by parts is exactly

$$\int t(x) S(r_k\partial_\alpha B, \dots; r_k x)\, dx = -\int (\partial_{\ell,\alpha} t)(x) S(B, \dots; r_k x)\, dx. \tag{WS4}$$

The derivative of a compactly supported test has the same support constraint. Apply the converging injective reading matrix at the receiving degree to the two sides. A left inverse close to the fixed free reading matrix then proves

$$\begin{aligned} U_{k,d+1}^{-1} I_{d,d+1}^{\mathrm{phys}} U_{k,d} &\longrightarrow I_{d,d+1}^{\mathrm{free}}, \\ U_{k,d+1}^{-1} r_k \partial_\alpha^{\mathrm{phys}} U_{k,d} &\longrightarrow \partial_\alpha^{\mathrm{free}}. \end{aligned} \tag{WS5}$$

Every inverse used in this deduction acts on a finite detector matrix with a fixed injective limit. It does not presume that the physical charts are uniformly conditioned.

One can make the first line of (WS5) exact. Choose left inverses P_d of the free adjacent inclusions. Average them over the finite signed permutation group and then average with their complex conjugates. The averaged maps are still left inverses and now respect both symmetries. Define recursively

$$X_{k,0} = U_{k,0}, \qquad X_{k,d+1} = I_{d,d+1}^{\mathrm{phys}} X_{k,d} P_d + U_{k,d+1}(1 - I_{d,d+1}^{\mathrm{free}} P_d). \tag{WS6}$$

Substitution of $I_{d,d+1}^{\mathrm{free}} v$ proves exact compatibility with the adjacent inclusion, hence with every inclusion by composition. Induction in d, using (WS5), gives $U_{k,d}^{-1} X_{k,d} \to 1$. Consequently these maps are eventually isomorphisms, preserve the unit and symmetries, and still satisfy the second line of (WS5). On the finitely many exceptional indices choose any isomorphisms; these choices do not affect a limit. Conjugation by matrices tending to the identity preserves all the finite-dimensional limits just obtained.

It remains to remove coordinate-generic restrictions and to pass to all position derivatives. For a fixed word let

$$R_k(x) = S(X_k v; r_k x). \qquad \partial_{\ell,\alpha} R_k(x) = S(X_k v_1, \dots, X_k D_{k,\alpha} v_\ell, \dots; r_k x), \quad D_{k,\alpha} = X_{k,d+1}^{-1} r_k \partial_\alpha X_{k,d}. \tag{WS7}$$

The matrices $D_{k,\alpha}$ converge to the free derivative matrix. By multilinearity, readings with any converging finite list of coefficient vectors have the same limiting readings as their limits. Also (WS2) applies to the finitely many successor-degree basis words in (WS7). It bounds δ_k times every needed gradient by a quantity tending to zero on a slightly larger separated window. Interpolate from nearby sampled points to arbitrary separated points along short coordinate segments in that larger window. Their total length is $O(\delta_k)$. The fundamental theorem of calculus and (WS7) make the error tend to zero. The argument does not require a positive gap between distinct coordinate projections: the points themselves stay separated.

Repeat (WS7) for each differentiation. For m derivatives use the m successor degrees for the differentiated labels and one additional degree for the interpolation gradient. Thus a horizon D proves the assertion when $d + m < D$. The finite sum over basis components proves convergence of the full derivative tensor in its norm. This convergence is uniform on every compact separated window. Otherwise a sequence of points violating uniform convergence has a convergent subsequence in that compact set; the moving-point interpolation just proved gives a contradiction. Enlarging the initial finite request supplies every prescribed derivative order, word and compact window, on the same retained physical hierarchy.

The interpolation near projection ties uses a different gradient bound from (WS8). On a slightly larger separated window each gradient in (WS7) is a finite sum of readings of successor-degree basis words, with coefficients from $D_{k,\alpha}$. These coefficient matrices converge and hence are bounded. The first, all-separated part of (WS3) applied to those basis words and the finite-maximum argument above therefore give

$$\delta_k \sup_{\text{larger separated window}} \|\nabla R_k\| \longrightarrow 0. \tag{WS9}$$

Round every coordinate to the rescaled mesh, then perturb the rounded coordinates by distinct bounded integer offsets to make every coordinate projection different. The number of forbidden equalities is finite for a fixed word, so one fixed finite list of offsets suffices. The displacement remains $O(\delta_k)$. For large k, all short joining coordinate segments stay in the larger separated window. The integral of the gradient along them is $o(1)$ by (WS9). This proves interpolation also at coordinate ties without losing a power through a shrinking coordinate gap.

There is one final sign choice when these maps are required to identify the normalized physical stress with the free stress. Each enlarged finite request allows one of the two signs of the physical conserved line. If one sign works for all requests, retain it. If a particular request excludes it, enlarge any subsequent request to include that one; its available map must use the opposite sign and restricts to a map for the subsequent request. Hence one fixed sign works for every finite request. Subsequent normalization fixes its amplitude as explained below.

For completeness, a single sequence realizing all the finite requests is obtained as follows. At arity n, exhaust the separated space by

$$K_{n,j} = \{x : \max_\ell |x_\ell| \le j,\ \min_{\ell \ne m} |x_\ell - x_m| \ge j^{-1}\}, \tag{WS10}$$

with the minimum condition omitted for zero or one slot. Enumerate the countable requests consisting of an arity, a finite tuple of degree bounds, a derivative order and an integer j. At stage s, enlarge one finite request to contain the first s requests, their finitely many basis words,

the detecting spectators, and one extra derivative degree. Use the finite construction on the same physical hierarchy and the same chosen stress sign. Along its available comparison sequence choose a station late enough that the radius is below $1/s$, the rescaled mesh is below $1/s$, all the first s derivative errors on their compact sets are below $1/s$, and every required instance of (WS9) is below $1/s$. There are finitely many conditions, each eventually true, so such a station exists. Pad maps above this stage's horizon by arbitrary fixed isomorphisms; every fixed degree eventually lies below the horizon.

The selected stage-s maps and radii form one common sequence. For a fixed enumerated request, all later stages include it and give errors at most $1/s$; hence it has the claimed limit. Every compact separated set lies in some $K_{n,j}$, so this proves all the asserted compact limits. If ordered mesh-period coefficient extractions are refined during these constructions, retain their countable columns and inverse charts together by the ordered compact extraction already used in (WS1): at each period take mesh limits first, then refine the period sequence and diagonalize the countable coordinate list. This preserves the same table and stress. The station diagonal just described is taken after those ordered limits and does not identify them with an arbitrary joint cutoff-volume diagonal. At any one station only its finite horizon is asserted; every fixed request is eventually included.

16.5 Constructing the physical stress and its shell charges

Centering removes the unit direction before stress rank is compared. Shell-flux representability is then proved for the conserved stress line, with the finite Ward residual displayed until its limiting estimate is consumed. Exact conservation of the eventual physical field does not imply vanishing of every finite localized divergence.

We first detail the rank transfer on centered tensors. A scaling isomorphism need not take the free expectation functional to the physical one. Let $c_0(v) = v - \langle v\rangle_0 1$ and $c(v) = v - \langle v\rangle 1$ be the two centering maps. Apply them componentwise to tensors. For a late scaling index, $U_{k,4}$ intertwines reality and signed symmetries and fixes the unit. It therefore induces the map

$$\mathcal{R}_k : \mathcal{K}_0 \longrightarrow \mathcal{K}, \qquad \mathcal{R}_k B = c(U_{k,4} B), \tag{SR1}$$

between the real centered symmetric signed tensor spaces. Its inverse is $A \mapsto c_0(U_{k,4}^{-1} A)$. Indeed the first centering subtracts only componentwise unit multiples. The inverse isomorphism carries those units to units, which the second centering kills; the remaining term is the original centered tensor. Thus these are exact inverse maps even though the two expectations differ. The corresponding degree-five isomorphism $\mathcal{Q}_k$ on real signed polar vectors needs no centering correction.

Write D_0, D for the free and physical divergence maps on these centered tensor spaces. On the fixed free spaces consider

$$Z_k = \mathcal{Q}_k^{-1}\, r_k D \mathcal{R}_k. \tag{SR2}$$

Derivatives of the unit are zero, so the centering correction in (SR1) contributes nothing to (SR2). The derivative comparison (WS7) gives $Z_k B \to D_0 B$ for every B. The free calculation makes D_0 onto and its kernel one dimensional. Choose a right inverse G of D_0. Then $Z_k G \to I$ in matrix norm, by convergence on a finite basis. Eventually $\|Z_k G - I\| < 1$; the geometric series makes $Z_k G$ invertible, hence Z_k onto. Since $r_k > 0$ and both surrounding maps in (SR2) are isomorphisms, D is onto. The exact dimension equalities from (SR1) and $\mathcal{Q}_k$, followed by rank–nullity, now give

$$\dim_{\mathbb{R}} \ker D = \dim_{\mathbb{R}} \ker D_0 = 1. \tag{SR3}$$

This establishes the conserved centered line without identifying the expectation functionals through U_k.

The physical stress is selected on the common field quotient, before its final amplitude is fixed. At weight four the real signed symmetric tensor space includes the vacuum tensor as well as (FC3), so it has dimension four. The rank comparisons transfer this dimension from the free table to the physical table and identify it with the image of the four real reference coefficients

$$J_{ab}(w) = \begin{cases} w_0\,1 + w_1 \sum_i Q_{ai,ai} + w_2 \sum_{i,k} Q_{ik,ik}, & a = b, \\ w_3 \sum_i Q_{ai,bi}, & a \neq b. \end{cases} \tag{SC1}$$

In (SC1) take the actual physical classes of these reference labels, not their free polynomial values. The physical weight-five polar vector space has dimension at most two, by the same rank comparison. Differentiation on the quotient therefore gives a linear divergence map from a four-dimensional real space into a space of dimension at most two:

$$(\nabla J)_b = \sum_a \partial_a J_{ab}. \tag{SC2}$$

Its range is polar because the signed derivative transformation has one sign for each derivative index; summing over that index cancels its squared sign.

Adjoin the real one-point reading of J_{00} to (SC2). The resulting map has target dimension at most three and thus has a nonzero kernel. Choose w in that kernel and normalize its Euclidean coefficient norm to one. The tensor map (SC1) is injective by the four-dimensional rank comparison, so $J(w) \neq 0$. It is conserved and its 00 mean vanishes. Reality makes every diagonal mean real; translation makes each mean constant in position; unsigned permutations make all diagonal means equal; reflection in one of two different indices negates an off-diagonal mean. Thus every mean vanishes. This constructs a centered nonzero real symmetric conserved tensor. It does not yet select its sign or claim its Ward charge is the unit translation generator.

Here is the normalization of this line along the comparison curve. For any nonzero tensor J in it and the fixed nonzero free tensor T^0, put

$$B_k = U_{k,4}^{-1}(r_k^4 J), \qquad \xi_k = \frac{\|B_k\|}{\|T^0\|} > 0, \qquad \widetilde{B}_k = \xi_k^{-1} B_k. \tag{SR4}$$

The inverse and norm here are on the full tensor coordinate space, including its scalar directions. The norm of $\widetilde{B}_k$ is exactly $\|T^0\|$, so finite-dimensional compactness gives a convergent subsequence with limit B of that norm. No positive lower bound on ξ_k is needed. Closedness of the real signed symmetric space puts B in that space. Its free mean is zero: the one-point comparison for bounded converging labels applies to $\widetilde{B}_k$, while its physical image is the positive scalar multiple $\xi_k^{-1} r_k^4 J$, whose expectation is exactly zero. Applying the converging derivative matrices (WS7) to $\widetilde{B}_k$ similarly passes the exact physical divergence zero to the free divergence zero. Thus $B \in \ker D_0$. By its one-dimensionality, $B = cT^0$ for real c; equality of norms gives $|c| = 1$. Choose the single sign of J that makes $c = 1$. Restricting the comparison parameter to this subsequence preserves all earlier limits and exact compatibility laws. This proves the normalized stress comparison used above; its eventual Ward amplitude is fixed separately below.

The finite lattice Ward identity associated with this construction retains its residual explicitly. At spacing a_N, write $T_{ab}^{N,p}(x)$ for the calibrated cylinder (SC1). Define its symmetric divergence, a clover-based directed Wilson variation, and its residual by

$$\begin{aligned} \operatorname{div}_a T_b(x) &= (2a_N)^{-1} \sum_a [T_{ab}(x + e_a) - T_{ab}(x - e_a)], \\ \Delta_b F(x) &= (2a_N^5)^{-1} \sum_i [\operatorname{contact}_{x,+i}(X_{bi}, F) - \operatorname{contact}_{x,-i}(X_{bi}, F)], \\ \mathcal{A}_b(x) &= \operatorname{div}_a T_b(x) - Z_\delta \mathcal{B}_b(x). \end{aligned} \tag{SC3}$$

Here X_{bi} is the signed clover: one quarter of the four oriented plaquette curvature readings, with signs $+,-,-,+$ for $(+b,+i),(+b,-i),(-b,+i),(-b,-i)$. All are transported to the common base. Its adjoint covariance and antisymmetry follow by reversing the corresponding plaquette words. In a color orthonormal basis e_c, the contact is $\sum_c X^c D_{x,\pm i,e_c} F$. The corresponding equation insertion is $\beta\langle X, \mathrm{force}_{x,\pm i}\rangle - \sum_c D_{x,\pm i,e_c} X^c$; $\mathcal{B}_b$ is its same signed sum divided by $2a_N^5$.

Right or left Haar invariance in that one directed link gives

$$\left\langle \left(\beta\langle X, \mathrm{force}\rangle - \sum_c D_c X^c \right) F \right\rangle = \left\langle \sum_c X^c D_c F \right\rangle. \tag{SC4}$$

Differentiate the change of variables in the finite Haar integral; the derivative of the Wilson density gives the force, and the product rule gives the derivative of X and of F. Compactness and continuity of the finite word derivatives justify moving that derivative through the integral. Linearity over the color basis proves (SC4) for the group-valued clover vector field.

For a lattice cutoff χ, a vector field V, an enclosed insertion F and an exterior word E, let

$$W_{ab}(x) = \frac{(\chi(x+e_a)-\chi(x))V_b(x+e_a) + (\chi(x)-\chi(x-e_a))V_b(x-e_a)}{2a_N}. \tag{SC5}$$

The finite flux is $-a_N^4 \sum_{x,a,b} W_{ab}(x)\langle T_{ab}(x) F E\rangle$. Its exact decomposition is

$$\mathrm{Flux} = a_N^4 \sum_{x,b} \chi(x) V_b(x) Z_\delta \langle \Delta_b F(x) E\rangle + a_N^4 \sum_{x,b} \chi(x) V_b(x) \langle \mathcal{A}_b(x) F E\rangle. \tag{SC6}$$

Use (SC4), the product rule, and finite symmetric summation by parts. The term with the discrete derivative of V vanishes for the constant and skew-linear Killing fields in question, by symmetry of T. The shifted V's in (SC5) are essential for this algebra. The derivative of E vanishes when each active Ward center avoids its full local insertion footprint. This is a support condition on the whole cylinder, stronger than just excluding its marked point. Taking absolute values in the second sum gives its exact error bound. No smallness of $\mathcal{A}$, including at the enclosed insertion, follows from this identity alone.

The continuum charges are instead supplied by the paid short-distance product estimates. Their input is the actual isolated product

$$r^{\,h} I_{J_{ab}}(\lceil y/a_N \rceil) I_B(0) - I_u(0), \qquad D + h = d + 4, \tag{SC7}$$

with B of weight d and u in the one fixed horizon D. At a fixed shell its coefficient bound may depend on that shell. It is common to every exterior test at that shell. The estimate is a rooted cylinder estimate on the reporting cell and on all its specified aligned seams and offsets, not only a two-point expectation. The paid product construction and weighted mesh interpolation give, for each finite family of exterior tests annihilating every D-field,

$$\sup_{y \in \{r_m/4 < |y| < 2r_m\}} \left| \sum_\nu c_\nu r_m^h \langle J_{ab}(y) B(0)\, \mathcal{O}_\nu \rangle \right| \longrightarrow 0. \tag{SC8}$$

The schedule uses $r_m > 0$ tending to zero and a finer mesh level $j(m)$ with $s_{j(m)}/s_m \to t > 1$. The product price chooses $t = 1 + 1/[64(d+12)]$, so $(d+12)(1-t) = -1/64 > -1/32$; its exponential gain dominates the growth of the two insertion columns. The rooted remainder after the exterior stencil estimate is at most $2\,\mathrm{stakes}(k)/\,\mathrm{pitts}(m)$, where $\mathrm{pitts}(m) \to \infty$. The common all-order mesh bounds transfer this from the fine reporting grid to the whole annulus, with the finer mesh

small compared with $r_m/128$. These are the annular product estimates used here; the finite Ward residual in (SC6) is not substituted for them.

Here are the rate choices that supply (SC7), rather than leaving its annular approximation as an input. Write $q_i = \log(a_i^{-1})$. The calendar has $q_{i+1}/q_i \to 1$. For $t > 1$, let $j(i)$ be the first index with $q_{j(i)} \geq tq_i$. Minimality gives $q_{j(i)-1} < tq_i$, so squeezing between these inequalities and $q_{j(i)}/q_{j(i)-1} \to 1$ gives $q_{j(i)}/q_i \to t$; also $j(i) > i+1$ eventually. A positive sequence of logarithmic rate b here means $\log f_i/q_i \to b$. Products add rates, quotients subtract them, and composition with j multiplies them by t. The estimates remain valid with arbitrarily small rate losses.

For $D + h = d + 4$, the two-field product and shell rescaling bounds give the available local approximation radius a logarithmic rate

$$-D + M, \qquad M = (d+12)(1-t). \tag{SC8a}$$

The d and 4 are the field weights, the h is the shell rescaling, and the additional eight is the pair of four-dimensional placement bounds. The adjacent insertion remainder starting at level k obeys the explicit induction bound

$$m_{k,i} \leq b \left(\frac{a_i}{a_k}\right)^{D+1} \frac{f_i}{f_k}, \tag{SC8b}$$

where $f_i > 0$ is its retained error allowance. The induction step is exactly the allowance inequality $\eta_i L_i^{D+1} f_i \leq f_{i+1}$, combined with $a_i = L_i a_{i+1}$. Choose $t = 1 + [64(d+12)]^{-1}$. Then $M = -1/64 > -1/32$. The test normalization $c_i = e^{q_i/32}$ has rate $1/32$, while the accumulated coefficient-normalization gain from level k to i is $e^{(q_i - q_k)/8}$, of rate $1/8$. Substitution of (SC8a)–(SC8b) into the finite product approximation leaves the strict inequalities

$$\frac{1}{32} - (D+1) < -D + M, \qquad \frac{1}{32} - M - \frac{1}{8} < 0. \tag{SC8c}$$

All column counts, fixed-frame norms and logarithmic factors have rate zero. The two strict margins in (SC8c) therefore pay them eventually, give a positive admissible approximation radius, and make its rooted error at most 2 $\text{stakes}_{p,k}/c_i$, the bound in (SC7). The support requirements hold as well: the shell radius tends to zero, and the finest mesh times every finite stencil count tends to zero relative to the fixed physical chamber size. Thus the finite two-field approximation applies at every point in the required annulus, with bounded coefficient representatives. This is the asserted source of (SC7).

Why do the approximants in (SC7) disappear from (SC8)? For a finite exterior family, let $L_N : E_D \to \mathbb{C}$ be its combined finite-cutoff reading. Its limiting reading is zero by hypothesis, and convergence on a basis of E_D gives $\|L_N\| \to 0$. At each fixed shell $\|u_N\| \leq A$, so $|L_N(u_N)| \leq A\|L_N\| \to 0$, although the approximant itself need not converge. The rooted estimate bounds the remaining product discrepancy against each exterior word by reflection positivity and its spectator cell bounds. The mesh interpolation then gives (SC8). This argument integrates annihilators, rather than selecting a new field at every integration point.

Let χ be smooth, compactly supported, one near the enclosed point and zero near the other insertions. For a Killing field V, define

$$Q_V(B;\chi) = -\sum_{a,b} \int (\partial_a \chi)(y) V_b(y - x_0) \langle J_{ab}(y) B(x_0)\, \mathcal{O} \rangle \, d^4 y. \tag{SC9}$$

Its integrand is supported on a compact shell disjoint from all insertions, so it is an ordinary integrable smooth function. Conservation and symmetry imply independence of the cutoff: for their difference integrate the divergence of $(\chi_1 - \chi_2) V_b \langle J_{ab} B \mathcal{O} \rangle$. The divergence of J is zero, and

the symmetric tensor contracts to zero with the skew derivative of V. The difference vanishes near every insertion and has compact support, so this integration by parts has no contact or infinity boundary term.

Take V homogeneous of degree $i = 0$ or 1 and shrink χ to radius r_m. Changing variables $y = x_0 + r_m z$ in (SC9) gives the factor r_m^{3+i}: four powers from volume, minus one from differentiating the cutoff, and i from V. Choose $h = 3 + i$ in (SC8). Its bound times the finite integral of $|\partial_a \chi(z) V_b(z)|$ tends to zero. Therefore every finite linear relation among the ordinary D-field exterior readings is satisfied by the flux readings, with $D = d + 1$ for translations and $D = d$ for rotations.

Finite-dimensional linear algebra now represents all these flux readings by one field. Indeed, if functionals $A_\nu \in E_D^*$ and values b_ν satisfy every finite relation $\sum c_\nu A_\nu = 0 \Rightarrow \sum c_\nu b_\nu = 0$, define a linear functional on their span by $A_\nu \mapsto b_\nu$, extend it to E_D^*, and use $E_D^{**} = E_D$. Faithfulness of the quotient makes the representing class unique. Initially use generic grid configurations; translation and continuity of the shell integral and of the kernels extend the same identity to every separated configuration, including coordinate ties.

Linearity of (SC9) in B and faithfulness prove that the representing field depends complex linearly on B. Its rotation charge depends real linearly on the skew matrix, by the same argument. These are the actual unnormalized translation and angular charge maps on the physical quotient. Their normalization and full rotation laws are established by the subsequent angular-flux and comparison arguments; centering and the finite-dimensional existence argument alone do not fix them.

The shell functional also has an ordinary face-flux interpretation. Let a closed coordinate box contain x_0 strictly in its interior and every exterior insertion strictly outside the closed box. Integrate each a-component of the current on its upper face minus its lower face, with three-dimensional Lebesgue measure in the remaining coordinates. This sum equals (SC9). To prove it choose a smooth window supported strictly inside the box and equal to one near x_0. Multiply the current by $1 - \chi$. It is zero near the only enclosed insertion and smooth at all points of the box. Its divergence equals $-\sum_a \partial_a \chi \, F_a$, by conservation and the Killing identity. The ordinary coordinate divergence theorem applies to this smooth truncated current. On every face $\chi = 0$, so its face flux is the original one. The volume integral is (SC9), and cutoff independence gives the equality for any allowed window. This argument treats faces with their outward orientations and does not integrate with four-dimensional ambient measure on the surface.

The bounded coefficient in (SC7) is required at each fixed shell before its lattice limit. It need not be uniform as the shell shrinks. The order is: fix the finite exterior annihilator and a reporting level; buy a sufficiently fine seam level; fix a shell index; then take the ordered mesh and period limits. At that shell the finite-dimensional annihilator norm tends to zero and the approximant coefficient is bounded. Only after this limit is taken does the shell radius go to zero. The rooted error tends to zero by (SC8a)–(SC8c). The same aligned-cell estimate applies on every seam and offset, so it permits moving a shell node freely in the annulus. This is the joint two-field approximation required for the shell construction.

Equation terms may be restored before this estimate is used. For one chamber their difference has zero expectation against every function supported outside that chamber, by the link Haar identity and avoidance of all the active links. Replace one chamber factor at a time in the full replicated array. All other factors are supported in complementary chambers, so each replacement leaves the integral unchanged. Finite induction replaces every chamber, including the conjugated reflected copies. Taking the reciprocal-cardinality power then gives equality of the two rooted seminorms. An equality only for a single unreplicated expectation would not suffice. No contact error at overlapping stencils is suppressed in this argument: overlap is excluded by the full footprint

condition, and actual contacts are kept in (SC4)–(SC6).

16.6 The centered stress trace

The positive two-point representation gives monotonicity in positive Euclidean time before the short-distance limit is known. The comparison can be selected inside any proposed sequence of radii on the same table and phase. Monotonicity sandwiches nearby generic distances around each selected radius; applying this to a hypothetical violating sequence proves the full right-hand limit. A single favorable scaling subsequence would not suffice for this step.

Use the selected character vacuum of the same hierarchy and its conserved physical stress tensor. Let $T = \sum_{\mu=0}^{3} T_{\mu\mu}$ be its Euclidean trace, a degree-four physical field, and let

$$\Theta = T - \langle T \rangle\, 1.$$

Its pinned reconstructed field is Hermitian and Lorentz scalar, and

$$\langle \Omega, \Theta(f)\Omega \rangle = 0, \qquad \Theta(f) = T(f) - \langle T \rangle \int f \; I \tag{TR1}$$

on the common field domain. To check the identity operator in this formula, first insert a unit label in any Euclidean correlation and use the hierarchy's unit law to remove it. The two corresponding strict complex-time functions agree on imaginary times and hence everywhere by holomorphic uniqueness. Smearing in that slot gives $\int f$ times the remaining function, by Fubini, and passing to the tempered boundary gives the same Wightman identity. Pair the resulting word vectors with arbitrary word vectors. Their dense span implies that unit insertion is $(\int f)I$ on the common domain. Linearity then proves (TR1). Similarly the Euclidean one-point function of Θ is zero; holomorphic uniqueness and the boundary limit prove its vanishing vacuum matrix element. Reality and scalar covariance follow from the stress tensor and the real invariant unit.

Let

$$E(s) = \langle \Theta(-\tfrac{s}{2}e_0)\Theta(\tfrac{s}{2}e_0) \rangle_E, \qquad s > 0.$$

The normalization used in the source gives

$$\lim_{s \downarrow 0} s^8 \log^2(1/s)\, E(s) = A_G > 0, \qquad A_G = 12 \frac{d_2(G)^2}{b_0(G)^2} \frac{\dim \mathfrak{g}}{\pi^4}, \quad d_2(G) = -\frac{b_0(G)}{2}. \tag{TR2}$$

Thus $A_G = 3 \dim \mathfrak{g}/\pi^4$ in these conventions. Positivity follows from $b_0(G) > 0$, $\dim \mathfrak{g} > 0$ for the simple nonabelian Lie algebra, and the nonzero coefficient d_2. These constants refer to the physical stress normalization and the Wilson-calibrated running charge.

We record the mechanism establishing the full positive-distance limit from the scaling construction. Orthogonal invariance of Θ and translation invariance identify its separated two-point function at any two points with E of their Euclidean distance. An orthogonal reflection takes their difference to a time-axis vector of the same norm; a common translation then puts them at the opposed points above. The degree-four scaling comparison identifies the rescaled trace with the scalar free-field label multiplied by $4d_2$, after the logarithmic running normalization. The free scalar two-point contraction is

$$\frac{3 \dim \mathfrak{g}}{\pi^4 \, |x_0 - x_1|^8}.$$

The running normalization contributes $b_0^{-2}/4$, while the trace coefficient contributes $(4d_2)^2$, giving exactly $A_G |x_0 - x_1|^{-8}$. This proves the normalized limit along the extracted scaling radii and their nearby generic configurations.

To remove the nearby-distance error, use the positive spectral representation

$$E(s) = \int_{\overline{V}_+} e^{-sp_0}\, d\mu_\Theta(p). \tag{TR3}$$

It is real and nonnegative. Since $p_0 \geq 0$ almost everywhere for this measure, $0 < s \leq t$ implies $e^{-tp_0} \leq e^{-sp_0}$, and integration proves $E(t) \leq E(s)$. Finiteness at every positive s is part of the positive-time spectral representation. Thus this monotonicity is available before proving the short-distance limit; it is not inferred from that limit.

We give the precise selection statement needed here. Fix the physical table, its phase decomposition and its stress tensor. The trace construction selects a comparison ultrafilter below the table's mesh-then-volume filter, its retained local data, and one free hierarchy. For **every** sequence $s_n > 0$ tending to zero, it then provides a strictly increasing map $j : \mathbb{N} \to \mathbb{N}$, a scaling family on this same physical table, and its comparison data, whose radii satisfy the exact identity

$$r_i = s_{j(i)}. \tag{TLim1}$$

The scaling comparison applies to these radii. In particular the statement is not merely that one favorable sequence $r_i \to 0$ exists. Nor does it select a different physical hierarchy or phase when the sequence s_n changes. The construction and this universal sequence property were established in the preceding interacting-to-Gaussian comparison.

First prove convergence along the radii of any such retained comparison. For $1/2 < u < 2$, put

$$x_0^{(u)} = 0, \qquad x_1^{(u)} = (u/2, u/2, u/2, u/2).$$

The four coordinates of the two points are distinct, so this is a generic two-point configuration, and its Euclidean separation is u. Its maximum-coordinate separation is $u/2 > 1/4$, while its configuration norm is less than 2. Consequently it has a neighborhood in the fixed comparison window with norm bound 100 and maximum-coordinate separation at least $1/100$.

Let $\widehat{x}_i^{(u)}$ be the rounded configurations in that comparison, and set $d_i(u) = |\widehat{x}_{i,0}^{(u)} - \widehat{x}_{i,1}^{(u)}|$. Rounding convergence gives $d_i(u) \to u$. The actual physical configuration is exactly $r_i \widehat{x}_i^{(u)}$, not just an asymptotic replacement of it. The interacting comparison, followed by the running-clock identity and the free scalar Wick contraction, gives

$$c(r_i)E(r_i d_i(u)) \longrightarrow A_G u^{-8}, \qquad c(r) = r^8 \log^2(1/r). \tag{TLim2}$$

The centering and scalar homogeneity used in this identity are retained in the trace-label calculation below. The same normalization $c(r_i)$, independent of u, is used on both sides of the coming inequalities.

Fix $a < 1 < b$ in $(1/2, 2)$. Eventually $0 < d_i(a) < 1 < d_i(b)$. Monotonicity and $c(r_i) \geq 0$ give

$$c(r_i)E(r_i d_i(b)) \leq c(r_i)E(r_i) \leq c(r_i)E(r_i d_i(a)).$$

By (TLim2), their outer limits are $A_G b^{-8}$ and $A_G a^{-8}$. Given $\epsilon > 0$, first choose a, b close enough to 1 that both outer limits differ from A_G by less than $\epsilon/2$. Then choose i large enough for the two convergences and the bracketing inequalities. This proves $c(r_i)E(r_i) \to A_G$. The complex pairing has zero imaginary part by (TR3), so this also proves complex convergence.

Finally, suppose the full right-hand limit at zero failed. There would be $\epsilon > 0$ and, for every n, a radius $0 < s_n < 2^{-n}$ such that

$$|c(s_n)E(s_n) - A_G| \geq \epsilon.$$

Apply the universal selection statement to this particular sequence. Equation (TLim1) retains this inequality at every r_i, whereas the preceding sandwich proves that its left side tends to zero. This contradiction proves (TR2) for the full small-scale limit. No density or ratio condition on a single favorable sequence of radii is being assumed.

The positive measure in (TR3) is unique and is the trace spectral measure. With the Fourier sign and normalization used in the reconstruction, it satisfies

$$\|\Theta(f)\Omega\|^2 = \int |\widehat{f}_+(p)|^2 \, d\mu_\Theta(p), \qquad \langle \Theta(f)\Omega, \Theta(g)\Omega\rangle = \int \overline{\widehat{f}_+(p)}\widehat{g}_+(p) \, d\mu_\Theta(p). \tag{TR4}$$

Its support lies in the forward cone, it is Lorentz invariant and of temperate growth, and every positive time damping has finite total mass. The full two-point Wightman distribution is its relative Fourier transform, also on joint Schwartz tests. Existence follows from the reconstructed joint momentum measure on the dense family of smeared trace vectors; these quadratic forms patch to a single positive measure. Uniqueness follows by testing all $|\widehat{f}_+|^2$, polarizing, and using compactly supported test functions to identify the measure. The regularized point-vector construction (TR5) and the measure argument (TR10)–(TR17) below give the patching and ultraviolet limit explicitly.

The spectral construction can be made directly from the regularized point vectors. Let v_s, $s > 0$, be the single trace shot with positive imaginary-time height s, centered at the origin. The shot identities already proved give

$$v_{s+t} = e^{-tH} v_s, \qquad W_2 \text{ at heights } (-s, s) = \langle v_s, U(x_1 - x_0) v_s\rangle.$$

Let μ_s be the diagonal joint-momentum measure of v_s. Functional calculus gives $\mu_{s+t} = e^{-2tp_0}\mu_s$. Define

$$d\mu_\Theta(p) = e^{2p_0} \, d\mu_1(p). \tag{TR5}$$

The semigroup identity shows equivalently $d\mu_\Theta = e^{2sp_0} d\mu_s$ for every $s > 0$: compare both heights through their sum. Hence $e^{-2sp_0}\mu_\Theta = \mu_s$ is finite, and its characteristic function is the regularized two-point function. This proves (TR3), with $s/2$ in the shot identity. Smearing the same identity and taking the previously established boundary limit yields (TR4) and the full joint-test Fourier representation. The momentum cone support is inherited from the shot measures, and scalar Lorentz covariance identifies each Lorentz pushforward with the same unique measure. The polynomial growth estimate for these measures follows from the common polynomial shot-height bound; this estimate will be detailed below.

The free contraction constant in (TR2) has a direct finite calculation. For $x \neq 0$, put

$$R_{ai}(x) = \delta_{ai} - 2x_a x_i/|x|^2.$$

This is an orthogonal reflection, since multiplying its two rows and summing the common index gives δ_{ab}. The free curvature covariance has the form

$$\mathrm{Curl}_{abij}(x) = D(x)(R_{ai}R_{bj} - R_{aj}R_{bi}), \quad D(x) = 4 \text{ normal } |x|^{-4}.$$

Orthogonality yields

$$\sum_{a,b,i,j} \mathrm{Curl}_{abij}(x)^2 = 24D(x)^2.$$

Indeed summing first over j, then i, gives $D^2(2 - 2\delta_{ab}^2)$; summing the a, b indices gives $24D^2$. The scalar curvature-square field has coefficient $1/4$ in each sum, two Wick contractions between its two quadratic factors, and $\dim \mathfrak{g}$ color choices. Thus its two-point function is $(1/16)\, 2 \dim \mathfrak{g}\, 24D(x)^2 =$

$3\dim\mathfrak{g}/(\pi^4|x|^8)$ with the source's Green-function normalization. This supplies the free constant used above.

We next derive the ultraviolet measure limit from (TR2). Write

$$Z(\rho)=\rho^{-8}\log^2\rho,\qquad \nu_\rho=Z(\rho)\,(p\mapsto p/\rho)_*\mu_\Theta,\quad \rho>1.$$

These are positive Lorentz-invariant cone-supported measures. For $s>0$,

$$\int e^{-sp_0}\,d\nu_\rho=Z(\rho)E(s/\rho)\longrightarrow A_G s^{-8},\tag{TR6}$$

because $Z(\rho)/Z(\rho/s)\to s^{-8}$. Lorentz invariance extends (TR6) to every strictly future timelike damping vector. The comparison measure is

$$d\nu_\infty(p)=\frac{A_G}{1536\pi}\,(p_0^2-|\mathbf{p}|^2)^2 1_{\{p_0\ge|\mathbf{p}|\}}\,d^4p.\tag{TR7}$$

Its time-axis Laplace integral is $A_G s^{-8}$, as can also be checked directly: integrate spatial radius $r\in[0,p_0]$, obtaining

$$4\pi\int_0^{p_0}(p_0^2-r^2)^2r^2\,dr=\frac{32\pi}{105}p_0^7,$$

and then $\int_0^\infty e^{-sp_0}p_0^7\,dp_0=7!s^{-8}$. The product is $1536\pi s^{-8}$. Positivity, cone support and finite timelike Laplace transforms allow the Laplace convergence theorem to identify the compact-test limit with (TR7); its proof is unfolded next.

There is a uniform polynomial bound strong enough for Schwartz tests. Since $e^{-p_0/R}\ge e^{-1}$ on $p_0\le R$,

$$\nu_\rho\{p_0\le R\}\le e\int e^{-p_0/R}\,d\nu_\rho\le CR^8,\qquad R\ge1,\tag{TR8}$$

uniformly for sufficiently large ρ. For the last inequality use (TR2) at $(\rho R)^{-1}$, and $Z(\rho)\le R^8Z(\rho R)$. The same bound holds for ν_∞. On the forward cone, Euclidean norm is comparable to p_0; dyadic annuli and Schwartz decay therefore make the tails of every Schwartz test uniformly small. Combining this with compact-test convergence gives

$$Z(\rho)\int\varphi(p/\rho)\,d\mu_\Theta(p)\longrightarrow\int\varphi\,d\nu_\infty\quad(\varphi\in\mathcal{S}(\mathbb{R}^4)).\tag{TR9}$$

For joint two-point tests use their relative wave transform. The normalized position dilation $f_\rho(x)=\rho^d f(\rho x)$ satisfies $\widehat{f}_{\rho,+}(p)=\widehat{f}_+(p/\rho)$ by change of variables. The same identity for the joint wave transform carries (TR9) to the source's relative and joint Wightman distribution limits.

We now prove the polynomial growth and Laplace continuity assertions used in (TR6)–(TR9), then the finite mass-shell obstruction. The subsequent stress correction and Ward calculation derive the trace coefficient and its Wilson normalization.

16.6.1 Polynomial bounds and the Laplace continuity argument

Which topology the spectral passage controls The spectral reader is first defined at positive time, where exponential damping makes its integral finite. Passing from these Laplace readings to a tempered distribution requires uniform polynomial growth bounds on the rescaled positive measures and a uniqueness argument on the allowed tests. The proof below supplies those bounds before taking the Schwartz limit. Convergence at a few positive-time values would not by itself justify arbitrary Schwartz smearing or the later spectral comparison tests.

Here are the measure arguments used above. They apply to positive measures; no assertion about convergence of arbitrary signed measures is needed. Suppose μ is supported in the closed future cone and

$$\mu\{p : p^0 \le R\} \le AR^k \qquad (R \ge 1). \tag{TR10}$$

The damped point-vector estimate supplies this inequality directly. Indeed, on the energy slab the weight $e^{-2p^0/R}$ is at least e^{-2}. If the squared norm of the vector at height R^{-1} is at most A_0R^k, its spectral representation gives (TR10) with $A = e^2A_0$. A Euclidean ball of radius R is contained in that energy slab. On the dyadic sets $2^j \le 1+\|p\| \le 2^{j+1}$, the mass is at most $A2^{k(j+1)}$. Consequently

$$\int (1 + \|p\|)^{-k-1}\, d\mu(p) \le 2^{k+1}A. \tag{TR11}$$

The sets may overlap on their boundaries, which only enlarges this upper bound. This proves the polynomial temperateness needed for Schwartz testing.

For completeness, let μ_i, ν be supported in the future cone C, with finite Laplace transforms at every u in its open dual cone C°, and suppose

$$\int e^{-u\cdot p}\, d\mu_i(p) \longrightarrow \int e^{-u\cdot p}\, d\nu(p) \quad (u \in C^\circ). \tag{TR12}$$

The pairing here is the real Euclidean pairing used for the damping parameter. On the one-point compactification $C \cup \{\infty\}$, the functions $s_u(p) = e^{-u\cdot p}$, extended by zero at infinity, are continuous: for $p \in C$,

$$u \cdot p \ge \tfrac{1}{2}(u^0 - \|\mathbf{u}\|)\|p\|.$$

Their products satisfy $s_u s_v = s_{u+v}$. Together with the constant function they separate points. They distinguish infinity from finite points by positivity; if two finite points have equal pairings with every $u \in C^\circ$, perturbing the time unit vector by a sufficiently small multiple of any real vector shows that the two points coincide. The real Stone–Weierstrass theorem therefore makes their linear span uniformly dense in the continuous functions on this compactification.

Fix $u \in C^\circ$ and damp all the measures by $e^{-u\cdot p}$. Their total masses converge by (TR12). Integrals of the constant and of every s_v converge, since the latter are precisely the Laplace transforms at $u + v$. Uniform approximation and the bounded total masses give convergence against every continuous function on the compactification. If f is continuous and compactly supported in the ambient space, $e^{u\cdot p}f(p)$, restricted to C and extended by zero at infinity, is such a function. Its damped integral equals $\int f\, d\mu_i$. This proves convergence against real compactly supported continuous tests, and then against complex tests by taking real and imaginary parts.

If (TR10) holds with common A, k for all these measures and the limit, this convergence extends to every Schwartz function. To see this explicitly, put $q = 2^{k+1}A$ and choose B with $|f(p)| \le B(1 + \|p\|)^{-k-2}$. Let χ_R be a smooth cutoff equal to one on the radius-R ball, supported in the radius-$R+1$ ball, with $0 \le \chi_R \le 1$. For $R \ge 1$,

$$\int |f - \chi_R f|\, d\mu_i \le Bq/R, \qquad \int |f - \chi_R f|\, d\nu \le Bq/R.$$

Choose R to make both tails smaller than a prescribed third of the error, then apply compact-test convergence to $\chi_R f$. This proves (TR9) from (TR6)–(TR8), with all tail estimates uniform. The same density argument also proves uniqueness of the positive cone-supported measure from its all-future Laplace transform.

16.6.2 The finite polynomial comparison class and its obstruction

The mass-shell exclusion is a finite comparison statement. The allowed functions are the finite complex linear combinations of finite products of the specified mass-shell integrals, with the stated integer and homogeneity restrictions. Their radial expansions have the prescribed power and logarithmic form, incompatible with the inverse square logarithm of the trace. The separate Maxwell argument retains its own locality, sector and comparison hypotheses.

The polynomial-field exclusion has no bound on the number of colors or on the degree of the two chosen polynomials. Each elementary contraction, however, has a finite sum of fixed mass-shell kernels with finite polynomial momentum numerators. More generally, the proof permits a real continuous homogeneous spatial numerator P of nonnegative integer degree j, an energy power $k \geq 0$, and a mass $m \geq 0$. Write

$$\omega_m(p) = \sqrt{m^2 + |p|^2}, \qquad a(p) = \frac{\omega_m(p)^k P(p)}{2\omega_m(p)}, \qquad p_+ = (\omega_m(p), p). \tag{TR13}$$

Values at the isolated massless zero momentum do not affect the integrals. A contraction track with n elementary factors acts on a total-momentum test by

$$L(\varphi) = \int_{(\mathbb{R}^3)^n} \prod_{\alpha=1}^{n} a_\alpha(p_\alpha)\, \varphi\left(\sum_{\alpha=1}^{n} (p_\alpha)_+\right) \prod_{\alpha=1}^{n} dp_\alpha. \tag{TR14}$$

The empty track is evaluation at zero. Allowed comparison kernels are finite complex linear combinations of these tracks. This includes units, arbitrary constant contraction coefficients and statistics signs; positivity of the comparison kernel is unnecessary.

Indeed, expand each polynomial into finitely many monomials. A cross-vertex pairing is a bijection between the occurrences at the two vertices. For each arc, expand its finite sum of elementary shell kernels, and then expand each polynomial numerator into monomials in energy and spatial coordinates. Each resulting term is (TR14); all arc coefficients, monomial amplitudes and pairing coefficients multiply into its constant coefficient. Summing the finite choices proves the asserted inclusion for any two finite polynomial fields. An infinite set of available generator labels causes no difficulty, since only finitely many occur in the two fields being compared.

At positive imaginary time s, a track is absolutely damped and its value factors:

$$L(e^{-sp^0}) = \prod_{\alpha=1}^{n} \int_{\mathbb{R}^3} a_\alpha(p) e^{-s\omega_{m_\alpha}(p)}\, dp. \tag{TR15}$$

Polar coordinates and homogeneity reduce an elementary factor to

$$\frac{1}{2}\left(\int_{\mathbb{S}^2} P(\theta)\, d\theta\right) \int_0^\infty r^{j+2} (m^2 + r^2)^{(k-1)/2} e^{-s\sqrt{m^2+r^2}}\, dr. \tag{TR16}$$

These formulas also give absolute integrability for compact total-momentum tests: such a test is bounded by a constant times e^{-sp^0} on its support.

The small-s radial expansion used here has the following exact meaning. For every integer N, the radial factor differs by $O(s^N)$ from a finite sum of terms $s^a(\log s)^b$, where $a \in \mathbb{Z}$ and $b \in \mathbb{N}$. Finite sums and products preserve this property. For the product statement, first bound each factor by some integer power of s, then take the other factor's expansion far enough to pay that bound. Thus the real part of every finite combination of (TR15) admits, in particular, such a finite expansion with bounded remainder. The radial expansion itself will be derived below from the massive integral; it is not an assumption on the comparison fields.

No real function with such a bounded-remainder expansion can have a nonzero finite limit after multiplication by $s^{8}\log^{2}(1/s)$. To prove this, the bounded remainder contributes zero. Combine equal terms in the finite expansion and set $s=2^{-n}$. A nonzero term becomes a constant times

$$\left(2^{-(8+a)}\right)^{n} n^{b+2}. \tag{TR17}$$

Order the finitely many terms lexicographically by their exponential base and then their polynomial degree. If the largest base is below one, every term tends to zero. If it is at least one, its term diverges in magnitude and every lower term divided by it tends to zero. Dividing a supposed finite limit by this largest term would force its nonzero coefficient to vanish. These two cases prove the assertion.

It remains to justify that equality only on the product tests stated in the theorem would imply equality of these imaginary-time values. Every smooth compactly supported momentum test v can be represented as the relative wave transform of a tensor product of two Schwartz position tests. Choose the first Fourier test equal to one on a ball containing $\operatorname{supp} v$; choose the second by Fourier surjectivity so that the relative product is v. Equivalently one may multiply v by the first Fourier test and its conjugate, whose squared modulus is one on this support. Hence equality on the indicated tensor tests gives equality of $\int v\,d\mu$ and the finite track combination for every such v.

Choose smooth compact cutoffs χ_n tending pointwise to one, with $0\le\chi_n\le 1$, and apply this equality to $v_n(p)=\chi_n(p)e^{-sp^{0}}$. On the trace side, finite damped mass permits dominated convergence. On each track side, the absolute damped integrability in (TR15) permits the same passage. Therefore equality of the tested two-point functions would imply

$$E(s)=\sum_{\ell=1}^{M} c_\ell L_\ell(e^{-sp^{0}}) \qquad (s>0).$$

Taking real parts and applying (TR17) contradicts the positive limit (TR2). This proves the stated exclusion for every finite polynomial comparison of this class.

The promised radial expansion follows from a convergent-integral argument, so it does not require summing a formal asymptotic series. Put

$$M_{j,k}(s)=\int_0^\infty \frac{r^{j+2}}{\sqrt{m^2+r^2}}(m^2+r^2)^{k/2}e^{-s\sqrt{m^2+r^2}}\,dr.$$

For $s>0$ these integrals are finite. The cases $k=0,1$ are dominated respectively by $r^{j+1}e^{-sr}$ and $r^{j+2}e^{-sr}$; higher cases follow from the exact recurrence

$$M_{j,k+2}=m^2M_{j,k}+M_{j+2,k}. \tag{TR18}$$

The gamma integral gives $M_{j,0}(s)\le\Gamma(j+2)s^{-j-2}$. Differentiation under the integral is justified locally at $s>0$ by the integrable bound with damping $s/2$, and gives $M'_{j,k}=-M_{j,k+1}$. Integration of the derivative of $r^{j+3}e^{-s\sqrt{m^2+r^2}}$ gives

$$sM_{j+2,0}=(j+3)M_{j,1}, \qquad \frac{d}{ds}(s^{j+3}M_{j,1})=-m^2s^{j+3}M_{j,0}. \tag{TR19}$$

The radial boundary terms vanish at zero and infinity. The derivative is integrable by the preceding bounds.

Here is the finite-order integration rule used with (TR19). Say that f has expansion to order $a\in\mathbb{Z}$ if f is a finite linear combination of integer powers of s times nonnegative powers of $\log s$,

up to $O(s^a)$. If $f' = g$, g has expansion to order a, and $a \neq -1$, then f has expansion to order $a+1$. Every finite expansion has an antiderivative of the same kind: integrate $s^{-1}(\log s)^b$ directly, and use integration by parts for all other powers. After subtracting this antiderivative it remains to treat $|f'| \leq Cs^a$. If $a < -1$, integrate from s to a fixed positive point to obtain $f = O(s^{a+1})$. If $a > -1$, the derivative is integrable at zero, so f has a finite limit and differs from that constant by $O(s^{a+1})$. The arbitrary integration constant is itself an allowed expansion term. This proves the rule.

Start with the bound at order $b = -j-2$ for $M_{j,0}$. If it has expansion to some integer order $a \geq b$, multiply by s^{j+3}, apply the integration rule to the second identity in (TR19), and divide back by that power. Since $j+3+a \geq 1$, the exceptional integration exponent is avoided. This yields an expansion of $M_{j,1}$ to order $a+1$. Its negative is the derivative of $M_{j,0}$; integrating improves the latter to at least order $a+1$. If $a+1 = -1$, first weaken the derivative remainder to order -2 and integrate that remainder, obtaining the required order -1. Otherwise the integration rule gives order $a+2$, which can be weakened to $a+1$. Induction supplies expansions of $M_{j,0}$ to every integer order. The preceding deduction then does the same for $M_{j,1}$, and (TR18) proves it for every k. This proves the expansion asserted after (TR16), and closes the polynomial-field exclusion.

16.6.3 The Maxwell comparison and the reduction to local sector estimates

The Maxwell conclusion concerns a specified comparison of representations. For m copies use the photon measure $d\lambda(k) = d^3k/(2|k|)$, with its irrelevant zero-momentum value assigned arbitrarily, and the Hilbert direct sum

$$\mathcal{F}_{\mathrm{amb}} = \bigoplus_{n \geq 0} \bigoplus_{s:\{1,\dots,n\} \to \{1,\dots,m\} \times \{1,2,3\}} L^2((\mathbb{R}^3)^n, \lambda^{\otimes n}). \tag{TR20}$$

The physical subspace consists of waves symmetric under simultaneous permutation of momenta and labels and transverse in each slot. Transversality means invariance under applying $\Pi_{ab}(k) = \delta_{ab} - k_a k_b/|k|^2$ to that slot's spatial component. Translation is multiplication by the plane wave at total momentum $P(k) = \sum_i (|k_i|, k_i)$. The vacuum is the unit wave in sector zero.

The comparison hypothesis is a complex linear isometry $U : \mathcal{H} \to \mathcal{F}_{\mathrm{amb}}$ whose image is physical, which takes the selected vacuum to this vacuum and intertwines translations. It requires neither surjectivity nor an identification of all fields. For a list h of real Maxwell electric test functions, its one-particle wave is

$$z_h(c, a; k) = \sum_b \Pi_{ab}(k)|k|\, \widehat{h}_{c,b,+}(|k|, k).$$

Its coherent wave has sector values

$$\mathcal{E}(h)_{n,s}(k) = e^{-\|z_h\|^2/2} \frac{i^n}{\sqrt{n!}} \prod_{\alpha=1}^{n} z_h(s_\alpha; k_\alpha). \tag{TR21}$$

Schwartz decay and $\lambda(\{|k| \leq R\}) = \pi R^2$ make z_h square integrable. The sector norm sum is the exponential series, so (TR21) defines a physical Hilbert vector.

Put $\sigma(f) = U\Theta(f)\Omega$. The excluded condition is the vacuum matrix identity

$$\langle \mathcal{E}(-h), \sigma(f) \rangle = \langle \sigma(\bar{f}), \mathcal{E}(h) \rangle \tag{TR22}$$

for every compact smooth f and compact smooth real list h whose supports are spacelike separated from that of f. The proof actually needs this only for f supported in one fixed ball of any positive radius. Thus it yields the precise small-region variant as well.

Here is the initial spectral reduction, which uses only the comparison hypothesis. Let $|Uv|^2\,d\Lambda$ denote the squared-wave measure on the direct sum (TR20). Its pushforward under P is the translation spectral measure of v: their characteristic functions are both $\langle v, T(b)v\rangle$. For $v = \Theta(f)\Omega$, this measure is $|\widehat{f}_+|^2\,d\mu$. Polarization, also after bounded momentum multipliers, therefore gives

$$\widehat{g}_+(P)\,\sigma(f) = \widehat{f}_+(P)\,\sigma(g) \quad \text{almost everywhere on the wave space.} \tag{TR23}$$

For clarity, the squared norm of the difference equals the corresponding norm in $L^2(\mu)$, where the two products are identical. This proves (TR23) without pointwise evaluation of a Hilbert vector at exceptional momenta.

The remaining sector argument constructs common amplitudes $q_{n,s}$ from (TR23), proves from (TR22) that they satisfy the local position identities, and uses their finite dilation spectrum. Its numerical end is particularly useful to state now. Write $R_j = 2^j$, $B_{n,j} = \{k : |k_i| \le R_j\ \forall i\}$, and let $\nu_q = |q|^2\lambda^{\otimes n}$. Positivity gives

$$\nu_{q_{n,s}}(B_{n,j}) \le \mu\{p : p^0 \le (n+1)R_j\}. \tag{TR24}$$

The trace asymptotic implies that the right side, multiplied by R_j^{-8}, tends to zero. The local sector estimate proves that every sector with $n \ge 4$ vanishes and that each of the remaining finitely many labeled sectors has

$$\nu_{q_{n,s}}(B_{n,j}) \le C_{n,s}R_j^6. \tag{TR25}$$

Summing gives $\mu\{p : p^0 \le R_j\} \le CR_j^6$, since an energy slab is contained in each sector's corresponding barrel. This contradicts the $R^8/\log^2 R$ high-energy law already proved. The common-amplitude construction and the local sector estimate are unfolded next.

Choose a countable sequence of Schwartz functions v_j whose positive Fourier transforms have modulus one on the radius-j ball. At each wave-space point choose the least j with $\widehat{v}_{j,+}(P) \ne 0$, and define

$$q = \sigma(v_j)/\widehat{v}_{j,+}(P). \tag{TR26}$$

This is measurable, since the selection is made from a countable measurable partition. For any fixed f, intersect the countably many full-measure sets in (TR23) for $g = v_j$. On that intersection, $\sigma(f) = \widehat{f}_+(P)q$. Thus the construction gives one amplitude for all f, with the usual separate almost-everywhere interpretation for each test. Weighted spectral equality then implies

$$\mu = P_*(|q|^2 d\Lambda), \qquad \mu(B) = \sum_{n,s} \int_{\{P(k)\in B\}} |q_{n,s}(k)|^2\,d\lambda^{\otimes n}(k). \tag{TR27}$$

To remove the Fourier weights, use the same v_j: their squared weights equal one on an exhausting sequence of momentum balls. Equality on each ball and countable exhaustion prove equality of the measures. This proves (TR24), since $P^0 = \sum_i |k_i| \le nR_j$ on $B_{n,j}$. The harmless $n+1$ in that bound also covers the vacuum sector.

We record explicitly the algebraic and measure-theoretic part of the local sector estimate. Treat amplitudes as measurable functions modulo almost-everywhere equality. For a slot i, let

$$(D_i q)(k_1, \dots, k_n) = q(k_1, \dots, 2k_i, \dots, k_n).$$

The operators D_i commute. The photon measure scales with degree two: the pushforward under $k \mapsto rk$ is $r^{-2}\lambda$. For a parcel $B(d) = \{|k_i| \le d_i\}$ this gives

$$\nu_{D_i q}(B(d)) = \tfrac{1}{4}\nu_q(B(d_1, \dots, 2d_i, \dots, d_n)). \tag{TR28}$$

The space of amplitudes with finite mass on every dyadic barrel and $R_j^{-8}\nu_q(B_{n,j}) \to 0$ is a real vector space, stable under every D_i. Addition uses $|q+r|^2 \le 2(|q|^2+|r|^2)$; stability follows from (TR28) and containment in the next barrel, with the fixed loss 2^6. Thus all finite polynomials in the D_i preserve this growth restriction.

The local distribution argument gives an integer K, common to the finitely many slots, such that

$$\prod_{a=0}^{K-1}(D_i - 2^{a-1})q = 0 \qquad (1 \le i \le n). \tag{TR29}$$

Here the exponent $a-1$ includes the electric-test factor $|k_i|$. To split these identities while preserving the growth restriction, work entirely in the preceding invariant real vector space. A linear operator annihilated on a vector by a polynomial with distinct real roots decomposes that vector into its finitely many eigencomponents. One direct induction takes $v = \prod_{a<t}(D_i - 2^{a-1})q$, which is an eigenvector of eigenvalue 2^{t-1}; subtract

$$\left(\prod_{a<t}(2^{t-1} - 2^{a-1})\right)^{-1} v$$

from q and reduce the number of factors by one. All denominators are positive. Since the slot operators commute, this splitting can be repeated in every slot inside the eigenspaces already selected and inside the kernels of the other annihilating polynomials. Hence q is a finite real sum of simultaneous eigenwaves, each satisfying the growth restriction.

Let $D_i w = 2^{a_i-1}w$ for such an eigenwave. Equation (TR28) yields

$$\nu_w(B(d_1,\dots,2d_i,\dots,d_n)) = 2^{2a_i}\nu_w(B(d)), \qquad \nu_w(B_{n,j}) = R_j^{2\sum_i a_i}\nu_w(B_{n,0}). \tag{TR30}$$

If $a_i = 0$, shrink only its radius successively by halves. The parcel masses remain equal; continuity of a finite measure from above identifies their limit with the mass of the set $k_i = 0$, which is zero. Starting with any barrel proves $w = 0$. Thus every nonzero eigenwave has $a_i \ge 1$. If $\sum_i a_i \ge 4$, the growth restriction and (TR30) again force its base barrel mass to vanish, so $w = 0$. The surviving waves therefore have $n \le \sum_i a_i \le 3$, and (TR30) bounds their barrel masses by a constant times R_j^6. Finite summation using the squared triangle inequality proves (TR25), including sector zero.

The weaker projected and averaged versions require no new sector estimate. A self-adjoint bounded map P which is the identity on $U\mathcal{H}$ can be moved across each inner product in (TR22), removing it from the coherent vector. For the open-region version, compact spacelike-separated supports have open neighborhoods with spacelike-separated closures, by compactness and openness of the spacelike relation. The neighborhood of $\operatorname{supp} f$ can also have closure inside the prescribed ball. Thus its asserted region identity implies the ball-local identity used above.

More generally, suppose a probability space indexes maps P_z preserving inner products between physical waves and fixing $U\mathcal{H}$ pointwise. Since coherent vectors are physical, $\langle P_z\mathcal{E}(h), Uv\rangle = \langle \mathcal{E}(h), Uv\rangle$, and likewise with the arguments reversed. The scalar integrands in the proposed averaged identity are therefore constant in z. Integration removes the averaging and gives the excluded local identity. This proves the averaged exclusion with exactly these hypotheses, without requiring boundedness, linearity, or a projection property of the individual P_z.

We now prove the distribution step leading to (TR29). The polynomial slab bound for μ and (TR27) imply, for each fixed sector amplitude, a weighted square-integrability bound

$$\int |q(k)|^2 \prod_i (1+|k_i|)^{-2d}\, d\lambda^{\otimes n}(k) < \infty \tag{TR31}$$

for some integer d. Indeed a sector ball of radius R lies in the total-energy slab $(n+1)R$, so its squared-wave measure has polynomial ball growth; apply (TR11) and then dominate its inverse radial weight by the product weight with d sufficiently large. Cauchy–Schwarz and $\int(1+|k|)^{-4}d\lambda(k)<\infty$ show that $|q(k)|\prod_i(1+|k_i|)^{-d-2}$ is integrable. Every integral below is controlled by this one after taking sufficiently high Schwartz seminorms. Slot dilation preserves (TR31), since $1+2|k|\le 2(1+|k|)$.

Use $\widehat{g}_-(p)=\widehat{g}_+(-p)$. After selecting a label sector, the consequence of (TR22) is

$$\int q(k)\prod_j |k_j|\widehat{g}_{j,-}((k_j)_+)\,d\lambda^n=\int \overline{q(k)}\prod_j |k_j|\widehat{g}_{j,+}((k_j)_+)\,d\lambda^n \tag{TR32}$$

whenever every g_j is compactly supported in the spacelike region relative to zero. Notice that no separation between two different g_j's is required. To pass from the ball-local identity to zero, smear the trace with a real normalized bump $a_\ell(x)=\ell^4 a(\ell x)$. For large ℓ its support is inside the chosen ball and spacelike to all the finitely many supports. Its Fourier transform tends pointwise to one and has modulus at most one. The sector coefficient identity extracted from (TR21) therefore tends to (TR32) by the integrable bound just established. Extraction of those coefficients, including the use of symmetry and transversality, is given below.

Keep all slots but i tested, and define the two equal-time data functionals

$$\begin{aligned}
A_b(f)&=\int q(k)|k_i|^{b+1}\widehat{f}_-(k_i)\prod_{j\neq i}|k_j|\widehat{g}_{j,-}((k_j)_+)\,d\lambda^n,\\
B_b(f)&=\int \overline{q(k)}|k_i|^{b+1}\widehat{f}_+(k_i)\prod_{j\neq i}|k_j|\widehat{g}_{j,+}((k_j)_+)\,d\lambda^n,\\
T_b(f)&=A_b(f)-(-1)^b B_b(f).
\end{aligned} \tag{TR33}$$

Here f is a spatial Schwartz test. These are continuous complex linear distributions. More precisely their norms are bounded by a constant, depending on q and the other tests, times the fixed seminorm

$$2^J(\|\widehat{f}_-\|_{0,0}+\|\widehat{f}_-\|_{J,0}+\|\widehat{f}_+\|_{0,0}+\|\widehat{f}_+\|_{J,0}),\qquad J=d+2+b+1. \tag{TR34}$$

The order of this seminorm does not depend on the other tests.

If f has compact support avoiding zero, insert into slot i the spacetime test $f(\mathbf{x})\partial_t^b a_\ell(t)$. Its support is eventually spacelike to zero. Its two Fourier transforms in the time variable carry the factors $(i|k_i|)^b$ and $(-i|k_i|)^b$, respectively. Applying (TR32), then taking $\ell\to\infty$, proves $T_b(f)=0$. Dominated convergence follows from (TR31) and the modulus-one bound for the Fourier transform of the normalized bump. Thus T_b is supported at the spatial origin.

A distribution controlled by a fixed continuous Schwartz seminorm and supported at zero depends on a uniformly finite jet at zero. Here is the proof with the uniformity needed in this application. Bound that seminorm by a finite maximum of Schwartz seminorms and let L be the largest derivative order occurring. If all derivatives of f through order L vanish at zero, multiply f by a fixed smooth cutoff $\chi(sx)$ equal to one near zero. Point support gives $T_b(f)=T_b(\chi(s\,\cdot)f)$. On the cutoff support $|x|\le C/s$, Taylor's theorem bounds the derivative of f of order $L-j$ by a constant times s^{-j-1}. Each order-j cutoff derivative costs s^j; the Leibniz formula therefore bounds all the finitely many required seminorms by C_f/s. The weights are bounded on this shrinking support. Letting $s\to\infty$ proves $T_b(f)=0$. Hence some L_b, depending only on the seminorm in (TR34), suffices for every choice of the other remote tests and every slot.

Let $C_2 f(x)=2^3 f(2x)$. Change variables in the photon measure to obtain

$$T_b[D_i q](f)=2^{-3-b}T_b[q](C_2 f). \tag{TR35}$$

The order-j jet of $C_2 f$ is 2^{3+j} times that of f. For $b = 0$, the induced factor on this jet is 2^j, the root $2^{(j+1)-1}$ in (TR29); for $b = 1$, it is 2^{j-1}, the root indexed by j. Choose $K = L_0 + L_1 + 2$. Applying the polynomial in (TR29) to q transfers by (TR35) to a polynomial on tests that kills every jet through L_b, for both $b = 0, 1$. Therefore the resulting amplitude r has $T_0[r] = T_1[r] = 0$ on all spatial tests, for every choice of the other remote tests.

These two zero data imply $r = 0$, as follows. Take a compact smooth momentum test w avoiding zero. In the $b = 0$ identity use the inverse negative Fourier test of $|k|w(k)$; in the $b = 1$ identity use that of $w(k)$. Multiplication by $|k|$ preserves smoothness on this support. The negative-frequency parts of the two identities agree, as do their positive-frequency parts, but the latter have opposite signs. Adding the identities isolates the negative-frequency integral and proves

$$\int r(k)|k_i|^2 w(k_i) \prod_{j \neq i} |k_j| \widehat{g}_{j,-}((k_j)_+)\, d\lambda^n = 0. \tag{TR36}$$

Translate each g_j a sufficiently small amount; it remains remote by openness and compactness. Multiply w by the corresponding shell plane wave in its slot, still a smooth test avoiding zero. Equation (TR36) then says that the Fourier transform, on a neighborhood of zero in all spacetime translation variables, of its integrable integrand vanishes.

An integrable function on a product of forward massless shells with this property vanishes almost everywhere. To prove it, restrict its Fourier transform to a line in a strictly future direction. Nonnegative shell energies give a bounded continuous function on the closed upper half-plane, holomorphic in its interior; the derivative is justified by the bound $ue^{-yu} \leq 1/y$. A zero boundary interval extends by Schwarz reflection and the holomorphic identity theorem to the whole line. First continue along strictly future lines, then express an arbitrary displacement as the difference of two sufficiently large strictly future displacements. The Fourier transform vanishes everywhere. Restricting to spatial translations and using ordinary Fourier uniqueness of the resulting finite complex measure proves the assertion.

Finally choose a sequence of remote bumps whose shell Fourier transforms are eventually nonzero at every fixed momentum, obtained by shrinking a normalized bump around a fixed spacelike point. Choose also spatial cutoffs avoiding zero that are eventually one at every nonzero momentum. Intersect the countably many almost-everywhere conclusions for their integrands in (TR36). Almost every shell momentum is nonzero, so all the multiplicative factors are eventually nonzero. It follows that $r = 0$ almost everywhere. Apply this to each slot's polynomial amplitude to obtain (TR29). This completes the locality-to-dilation argument.

To finish the promised extraction in (TR32), define for a physical vector x

$$L_x(g_1, \dots, g_n) = \sum_s \int \overline{x_{n,s}(k)} \prod_i z_{g_i}(s_i; k_i)\, d\lambda^n.$$

Cauchy–Schwarz makes every term integrable. Decomposition of the Hilbert inner product by orthogonal sectors gives the convergent coherent expansion

$$\langle x, \mathcal{E}(h) \rangle = e^{-\|z_h\|^2/2} \sum_{n \geq 0} \frac{i^n}{\sqrt{n!}} L_x(h, \dots, h). \tag{TR37}$$

For $\langle \mathcal{E}(-h), x \rangle$, the corresponding coefficients are $i^n \overline{L_x(h, \dots, h)}/\sqrt{n!}$. Apply (TR22) to th for every positive real t. The identical positive Gaussian factors cancel. The resulting power series is zero for every $t > 0$, and each coefficient is zero. One can justify coefficient uniqueness without an analytic-series theorem: after the first N coefficients vanish, divide by t^N; the tail is dominated for $0 < t \leq 1$ by the absolutely summable series at $t = 1$, so $t \downarrow 0$ recovers the next coefficient.

Each L_x is multilinear and symmetric. Symmetry follows by permuting the momentum variables and the labels together, using physical symmetry of x and invariance of the product photon measure. The diagonal identity for real lists therefore polarizes to arbitrary real lists. For example, expand the diagonal polynomial at $h+tu$; its derivative replaces one slot by u, and symmetry identifies the sum with n times that term. Induction on n recovers all mixed terms. Compact tests supported spacelike to the trace support form a complex vector space, and every such list is a complex linear combination of real lists by its real and imaginary parts, with the same support restrictions. Complex multilinearity then gives the polarized identity for all these lists, with complex conjugation of the lists on the conjugated side.

Transversality removes the polarization matrices in these integrals. For a fixed slot expand $z_g = \Pi\, z_g^{\mathrm{raw}}$; move the real symmetric matrix Π onto the wave component in the finite sum over spatial labels. That component is unchanged because x is transverse. Repeat in every slot. More explicitly, exchanging the selected component label and the summed matrix index is a bijection of the finite label-index pairs; all other slot factors remain fixed. No regularity of wave representatives is used, since there are only finitely many almost-everywhere transverse identities at each particle number.

Finally take in slot i a list having just the component s_i nonzero, with scalar test g_i. After the projectors are removed, all label sectors except s vanish in the sum. The surviving polarized identity is

$$\int x_{n,s}(k) \prod_i |k_i| \widehat{g}_{i,-}((k_i)_+)\, d\lambda^n = \int \overline{y_{n,s}(k)} \prod_i |k_i| \widehat{g}_{i,+}((k_i)_+)\, d\lambda^n,$$

where $x = \sigma(f)$ and $y = \sigma(\bar{f})$. The Fourier conjugation identity $\overline{\widehat{\bar{g}}_+(p)} = \widehat{g}_-(p)$ fixes the signs. Taking the real shrinking trace bump as described above yields (TR32). Together with (TR26)–(TR36) and the high-energy contradiction, this completes all the Maxwell exclusions asserted here.

16.6.4 Passing a running trace label through the scaling limit

There are two separate ingredients in the trace limit: convergence of the running label in its finite-dimensional field space, and convergence of the moments evaluated on that label. The latter passage does not require a uniform convergence assertion on an infinite-dimensional space. Let E_d be the fixed free field space of degree d, choose a basis e_a, and let $U_i : E_d \to E_d^{(i)}$ be the scaling identifications. The source proves, for every pair of basis vectors,

$$\langle U_i e_a(x_i) U_i e_b(y_i) \rangle_\chi \longrightarrow \langle e_a(x) e_b(y) \rangle_{\mathrm{free}} \tag{TR38}$$

along the same chosen refinement and in every selected phase under consideration. The rounded scaled configuration tends to the target configuration; its rounding error is at most $4n/\mathfrak{m}_i$, where the mesh magnification $\mathfrak{m}_i \to \infty$. Generic separated configurations stay in their open admissible region. The comparison error is bounded by $o_i(1 + M(x_i, y_i))$, where $o_i \to 0$ and the finite profile norm M is continuous at the separated target. This proves (TR38).

If $v_i, w_i \in E_d$ tend to v, w, expand them in the chosen finite basis. Every coordinate converges, and bilinearity expresses the moment as the finite sum

$$\sum_{a,b} (v_i)_a (w_i)_b \langle U_i e_a(x_i) U_i e_b(y_i) \rangle_\chi.$$

Taking the limit term by term gives the free moment of v, w. This is the passage used for the running centered trace, rather than an interchange of an uncontrolled varying field with a distributional limit.

The centering convention is also harmless at this step for a specific reason. If $T_{\mu\nu}$ is the constructed stress tensor, replace it by $T_{\mu\nu} - \langle T_{\mu\nu} \rangle\, 1$. This preserves symmetry, reality, conservation and its charges: derivatives of the unit vanish, and a constant symmetric unit tensor has zero charge contribution in the stated Ward construction. Its trace differs from the phase-centered trace by a constant multiple of the unit. Under the inverse scaling map the latter contributes

$$B_i r_i^4\, c\, 1,$$

which tends to zero by the running-band estimate. The scaling maps identify the unit exactly. Thus the tensor normalization can first be proved with the convenient centered tensor and then transferred to the exact selected-vacuum trace.

In the source's coupling convention

$$b_0 = \frac{11\, C_{\mathrm{adj}}}{48\pi^2}, \qquad \beta_{\mathrm{slope}} = 2b_0, \qquad d_2 = -\frac{b_0}{2}. \tag{TR39}$$

The last identity follows from the stress coefficient $d_2 =$ forceR after substituting the displayed slope; its derivation from the stress comparison is treated with that comparison below. Once the running centered trace converges to $4d_2\, \mathcal{O}_{\mathrm{scale}}$, the preceding finite-basis argument gives its two-point limit. The free Wick calculation yields $3 \dim \mathfrak{g}/(\pi^4 |x-y|^8)$. The running charge contributes $(2b_0)^{-2}$ after multiplication by $\log^2(1/r_i)$. Consequently the constant is

$$(2b_0)^{-2} (4d_2)^2 \frac{3 \dim \mathfrak{g}}{\pi^4} = \frac{12 d_2^2}{b_0^2} \frac{\dim \mathfrak{g}}{\pi^4} = \frac{3 \dim \mathfrak{g}}{\pi^4},$$

as claimed in (TR2). The powers and logarithms cancel before taking the limit: the two running labels each carry $B_i r_i^4$, and the prefactor is $(\log(1/r_i)/B_i)^2$.

The running-clock cancellation used here follows directly from the calibrated clock, including the extra magnification in the scaling sequence. Write

$$\widehat{\beta}_i = \beta_* + g\, t_i, \qquad r_i = e^{-t_i} a_i, \qquad g = 2b_0 > 0.$$

The chosen magnifications satisfy $a_i \geq 1$ eventually and $0 \leq \log a_i \leq \mathrm{Gate}(\widehat{\beta}_i)$, where $\mathrm{Gate}(x)/x \to 0$. The calibrated coupling sequence and the band charge satisfy $\beta_i/\widehat{\beta}_i \to 1$ and $\beta_i/B_i \to 1$. Therefore

$$\frac{\log(1/r_i)}{B_i} = \left(\frac{t_i - \log a_i}{\widehat{\beta}_i} \right) \frac{\widehat{\beta}_i}{\beta_i} \frac{\beta_i}{B_i} \longrightarrow \frac{1}{g}. \tag{TR40}$$

Here $t_i/\widehat{\beta}_i = (1 - \beta_*/\widehat{\beta}_i)/g$, and the magnification term tends to zero by its bound. Also $\log(1/r_i) \to \infty$, because $r_i \downarrow 0$ along the scaling sequence. The elementary exponential limit $\ell e^{-4\ell} \to 0$, with $\ell = \log(1/r_i)$, gives

$$B_i r_i^4 = \frac{\log(1/r_i) r_i^4}{\log(1/r_i)/B_i} \longrightarrow 0.$$

This proves both the logarithmic factor and the disappearing unit contribution used in the trace argument.

16.6.5 The three finite candidates and their common extraction

All three finite tensor candidates, their individual stencil derivatives and their divergences are retained in one extraction. Thus a divergence is the sum of derivatives of those same finite functions.

The correction does not combine unrelated limiting candidates whose coefficient coordinates happen to agree.

Here are the actual finite functions behind the three columns used below. Write $C_{ab}(x) = \mathrm{curv}(\mathrm{Hol}_{\square(a,b)}(x))$, with the curvature coordinate and normalization of (SQ1)–(SQ12), and put $P_{ija}(x) = \langle C_{ia}(x), C_{ja}(x)\rangle$. At a fixed issued ticket, the three tensors are

$$\begin{aligned} J^{\varnothing}_{ij}(x) &= \mathrm{Meld}(J)_{ij}(x), \\ J^{\mathrm{sh}}_{ij}(x) &= \begin{cases} 0 & i = j, \\ \sum_a P_{ija}(x) & i \neq j, \end{cases} \\ J^{\mathrm{sc}}_{ij}(x) &= \frac{\delta_{ij}}{4} \sum_{a,b} P_{aab}(x). \end{aligned} \tag{TF11}$$

The first is exactly the restored, localized principal tensor (TF0d), including its six antisymmetric differences and its scalar forward difference. The other two use elementary plaquette holonomies; there is no new stationary current in their definition. Their free coefficient vectors in the three-coordinate tensor basis are respectively the principal vector, $(0, 0, 1)$, and $(0, 1/4, 0)$. In the source's three-coordinate convention the principal free vector is $-\,\mathrm{weyl}$. Its physical sign is fixed together with the Ward charge in (TR41)–(TR44); it cannot be changed independently of that normalization.

All sixteen entries of each tensor, all sixteen individual stencil derivatives, and all four divergences are extracted on the same filter. Explicitly these are the origin functions $J^b_{ij}(0)$, $\mathrm{elm}_i\, J^b_{ij}(0)$, and $\sum_i \mathrm{elm}_i\, J^b_{ij}(0)$, for $b = \varnothing, \mathrm{sh}, \mathrm{sc}$. The tensor entries have degree-four coefficient records and the derivatives have degree-five records. These records include both the finite-cutoff label and its retained limit; extraction of the joint correlation book gives the corresponding fields L_i, Ω_i, Δ_i. Thus extraction of a divergence is extraction of the sum of its actual finite derivatives, not an independently selected vector limit.

The local bounds needed here have a direct source. A plaquette has four links. If each matrix link has derivatives through order k bounded by the chart price d, its holonomy has price $\mathrm{travelCost}(G, k, 4, d)$. Applying the curvature map gives price $z(d) = \mathrm{curvatureCost}(G)\,\mathrm{travelCost}(G, k, 4, d)$. The bilinear product has price $2^k \|\langle\cdot,\cdot\rangle\| z(d)^2$. Sum the finitely many axes and multiply the scalar branch by $1/4$. This gives a bound independent of the torus and the ticket, supported in radius four. One stencil derivative enlarges the radius to eight. The scheduled small-field extension makes these bounds into the coefficient and correlation bounds used in extraction. The issued principal tensor has these same bounds by (TF0d) and its localization estimates. The condition $20(\mathrm{high}(x) + \mathrm{rim}(x) + 10) < n$ makes all integer offsets used in the coefficient tests nonwrapping.

Taylor projection commutes with the stencil divergence on these records: expand the finite translates and use the stencil moments from (TF0b). The degree-four stress coordinate therefore maps to its degree-five divergence coordinate. For the principal branch the antisymmetric difference corrections have zero relevant tensor coordinate; the scalar difference has zero trace coordinate. Consequently its free divergence scalar coordinate vanishes. This is the specific low-degree statement used later; it does not say that its entire finite divergence vanishes. The finite action, localized current, surgery polynomial and coefficient equations are issued by one ticket, so the six terms of (LCE6) concern this same first branch. The two elementary branches supply the correction columns in (TR41).

With $\mathcal{Y}_i T = r_i^4 U_i^{-1} T$ and $\mathcal{P}_i = B_i \mathcal{Y}_i$, their normalized limits and correction are precisely (TR41)–(TR44): solve the two divergence coordinates for $c_{i,0}, c_{i,1}$, then set $\widetilde{T}_i = B_i L_i + c_{i,0}\Omega_i + c_{i,1}\Delta_i$. This specifies the finite origin, normalization, and limiting meaning of all three candidates and both

correction objects.

16.6.6 Correcting the stress tensor and fixing its amplitude

The two correction coordinates are solved by an eventually invertible finite-dimensional detector matrix. Its kernel agrees with the true divergence kernel on the stated stress space. This converts small detector errors into an exactly conserved tensor, while the weighted scalar sensor fixes the trace amplitude. The local estimates following the inverse argument supply the quantitative bounds needed for this stronger conclusion.

The coefficient in (TR39) is earned by an exact finite-dimensional correction, not by inserting an assumed trace anomaly. The three lifted stress candidates are denoted L_i, Ω_i, Δ_i. Their free limits are respectively the Maxwell stress tensor, a traceless correction, and a scalar tensor with trace $4\mathcal{O}_{\mathrm{scale}}$. All lie in the centered symmetric, real, rotation-covariant stress space. Let K_i be the two-component detector map for their divergences, and define

$$S_i(c_0, c_1) = c_0\Omega_i + c_1\Delta_i, \qquad X_i = B_i K_i L_i.$$

The maps $K_i S_i$ converge to an invertible free map. In finite dimensions this gives eventual invertibility and a uniform inverse bound: identify matrices in fixed bases, write the nearby matrix as $A(I + A^{-1}E_i)$, and use the convergent geometric inverse when $\|A^{-1}E_i\| < 1/2$. Solve

$$(K_i S_i)c_i = -X_i. \tag{TR41}$$

The available detector estimate bounds $\|X_i\| \le CR_i$, hence $\|c_i\| \le C'R_i$, with $R_i/B_i \to 0$. The detector map has exactly the same kernel as the true divergence on this stress space. Consequently $\widetilde{T}_i = B_i L_i + S_i c_i$ is exactly conserved for every sufficiently large i.

The distinguished scalar detector ℓ_i satisfies, for a nonnegative control t_i with $R_i \le t_i$,

$$\ell_i(X_i) \to -g, \qquad t_i\ell_i(K_i\Omega_i) \to 0, \qquad t_i\big(\ell_i(K_i\Delta_i) + 4\big) \to 0. \tag{TR42}$$

The coefficient bound therefore kills the two error products. Applying ℓ_i to (TR41) gives

$$0 = \ell_i(X_i) + c_{i,0}\ell_i(K_i\Omega_i) + c_{i,1}\ell_i(K_i\Delta_i).$$

Equivalently,

$$4c_{i,1} = \ell_i(X_i) + c_{i,0}\ell_i(K_i\Omega_i) + c_{i,1}\big(\ell_i(K_i\Delta_i) + 4\big).$$

Thus $c_{i,1} \to -g/4 = \mathrm{forceR}$. The factor four is the exact scalar detector value, and the sign follows from the displayed equation. The primitive comparisons supplying (TR42) are unfolded with the detector construction.

Write $\mathcal{Y}_i T = r_i^4(U_i^{-1}T)$, component by component, and $\mathcal{P}_i = B_i\mathcal{Y}_i$. The normalized lifted candidates satisfy $\mathcal{P}_i L_i \to T_{\mathrm{Maxwell}}$, $\mathcal{P}_i\Omega_i \to T_\Omega$, and $\mathcal{P}_i\Delta_i \to T_\Delta$. Therefore

$$\mathcal{Y}_i\widetilde{T}_i = \mathcal{P}_i L_i + \frac{c_{i,0}}{B_i}\mathcal{P}_i\Omega_i + \frac{c_{i,1}}{B_i}\mathcal{P}_i\Delta_i \longrightarrow T_{\mathrm{Maxwell}}. \tag{TR43}$$

For the scaled trace, the stronger traceless estimates give $B_i\,\mathrm{tr}(\mathcal{P}_i L_i) \to 0$ and $c_{i,0}\,\mathrm{tr}(\mathcal{P}_i\Omega_i) \to 0$. The remaining trace tends to

$$B_i\,\mathrm{tr}(\mathcal{Y}_i\widetilde{T}_i) \longrightarrow 4\,\mathrm{forceR}\,\mathcal{O}_{\mathrm{scale}}. \tag{TR44}$$

Only the first correction coefficient requires a growth bound here; the scalar coefficient has the actual finite limit just proved.

The conserved centered stress space is one-dimensional. To see why this rank statement follows from the same construction, K_iS_i being onto makes K_i onto. Its kernel equals the divergence kernel. The real dimension of the centered stress space agrees with the corresponding free dimension by the scaling and unit identifications. The detector target has the free divergence-target dimension. Rank–nullity and the free calculation, whose divergence kernel has dimension one, therefore show that the physical divergence kernel has dimension one as well.

Let R be the actual centered stress tensor with its translation Ward charges. It is nonzero. Indeed the free derivative map is nonzero: its divergence target contains a nonzero coordinate and is attained; if every observer vanished on every first derivative, the separating observers would make that target zero, a contradiction. Choose an observer with a nonzero derivative value y. Its free translation-flux Ward identity has value y. The corresponding physical shell-flux test of R, with positions, cutoff and fields scaled together and with the single factor r_i prescribed by the derivative, tends to y. If $R = 0$, every such test would be zero, contradicting $y \neq 0$.

Since the kernel is one-dimensional, $\widetilde{T}_i = v_i R$ for real v_i eventually. The shell-flux test of R tends to y, while its test on $\widetilde{T}_i$ also tends to y, by (TR43) and the tested derivative comparison. Linearity makes the latter v_i times the former. Divide by the nonzero limiting value to obtain $v_i \to 1$. Then $v_i \neq 0$ eventually, and multiplying (TR43) and (TR44) by v_i^{-1} transfers both limits to R. Centering back to the selected-vacuum trace uses the vanishing unit term proved in (TR40). This proves the running trace-label limit required for (TR2).

There is also an exact scalar-coordinate formulation. The invariant scalar degree-four space has the real basis $1, 4\mathcal{O}_{\mathrm{scale}}$, with injective coordinate map $(e, d) \mapsto e\,1 + 4d\,\mathcal{O}_{\mathrm{scale}}$. A linear left inverse exists and is continuous in this finite-dimensional space. Applying it to the running trace limit gives coefficients e_i, d_i with

$$U_i^{-1}(r_i^4\Theta) = e_i\,1 + 4d_i\,\mathcal{O}_{\mathrm{scale}}, \qquad B_i d_i \to d_2, \qquad B_i e_i \to 0.$$

The exact scalar membership, followed by this left inverse, justifies these simultaneous coefficients without selecting unrelated subsequences.

The two correction columns have a concrete free-coordinate description. The free divergence detector is a surjective map $V : \mathbb{R}^3 \to \mathcal{D}$ with

$$\ker V = \mathbb{R}(1, -\tfrac{1}{4}, 1).$$

The two columns are $\Omega = (0, 0, 1)$ and $\Delta = (0, \frac{1}{4}, 0)$. Their images are independent: if $x\Omega + y\Delta = c(1, -\frac{1}{4}, 1)$, coordinate zero gives $c = 0$, and the other coordinates give $x = y = 0$. They span the quotient: for $w \in \mathbb{R}^3$,

$$(w_2 - w_0)\Omega + (4w_1 + w_0)\Delta = w - w_0(1, -\tfrac{1}{4}, 1). \tag{TR45}$$

Thus their detector images form a basis of $\mathcal{D}$. This is the free invertibility used in (TR41).

Here are the precise error transfers behind (TR42). Let s_i be the moving map from the finite bare degree-five coordinate space to the free divergence target, with $s_i \to s_\infty$ in operator norm. For a lifted candidate A_i, its detector value $p_i = K_iA_i$ and its bare divergence coordinates u_i satisfy

$$H_i^N(p_i - s_iu_i) \to 0 \qquad \text{for every fixed } N \geq 0. \tag{TR46}$$

The increasing clock H_i dominates the positive control t_i. For the two correction candidates, $u_i \to u_*$, with $s_\infty u_* = V\Omega$ or $V\Delta$, respectively. In particular (TR46) at $N = 0$ gives the two column limits.

The scalar functionals ℓ_i are uniformly bounded and approximate a fixed bare functional $\mathcal{L}$ with the stronger rate

$$\sup_{\|u\|\le 1} |t_i(\ell_i(s_i u) - \mathcal{L}(u))| \longrightarrow 0. \tag{TR47}$$

For each correction candidate the bare comparison also gives $t_i(\mathcal{L}(u_i) - \mathcal{L}(u_*)) \to 0$. Decompose the full scalar error into the three terms

$$\begin{aligned} t_i(\ell_i(p_i) - \mathcal{L}(u_*)) = {} & t_i\ell_i(p_i - s_i u_i) \\ & + t_i(\ell_i(s_i u_i) - \mathcal{L}(u_i)) \\ & + t_i(\mathcal{L}(u_i) - \mathcal{L}(u_*)). \end{aligned}$$

The first tends to zero by boundedness of ℓ_i, $t_i \le H_i$, and (TR46) at $N = 1$. The second tends to zero by (TR47) and boundedness of u_i. The third is the stated bare comparison. The exact free scalar readings are $\mathcal{L}(u_\Omega) = 0$ and $\mathcal{L}(u_\Delta) = -4$, giving the last two limits in (TR42).

For the principal candidate, a renormalized bare divergence v_i satisfies

$$X_i - s_i v_i \to 0, \qquad \|v_i\| \le C R_i, \quad R_i \le t_i, \qquad \mathcal{L}(v_i) \to -g. \tag{TR48}$$

Although v_i need not stay bounded, (TR47) still pays its growth:

$$|\ell_i(s_i v_i) - \mathcal{L}(v_i)| \le \frac{\varepsilon_i}{t_i}\|v_i\| \le C\varepsilon_i \to 0.$$

Uniform boundedness of ℓ_i handles the first difference in (TR48), proving $\ell_i(X_i) \to -g$. Operator-norm boundedness of s_i, the bound on v_i, and $R_i \ge 1$ eventually give $\|X_i\| \le C' R_i$. Thus the same estimates supply both the scalar limit and the norm bound needed to solve the correction equation.

The elementary absorption principle used throughout this step is: if $g_i \ge 0$, $|a_i| \le C g_i$, and $g_i x_i \to 0$ in a normed space, then $a_i x_i \to 0$. It follows immediately by taking norms. In (TR48), the first difference is obtained by applying this principle to the two preexisting comparisons: the physical detector differs from the candidate's bare coordinate reading by an error smaller than every inverse power of H_i, and the candidate's renormalized divergence differs from v_i by an error of the same strength. Since $B_i \le 2H_i$, both errors remain negligible after multiplication by B_i. This explains why growing running factors can be inserted in (TR42)–(TR44).

The scalar membership used after (TR44) follows from the same free degree-four classification. A real scalar u, fixed by all signed orthogonal symmetries, gives a diagonal tensor $J_{ab} = \delta_{ab} u$. It belongs to the classified free symmetric stress space. Writing it in that space's real coordinates w, its trace has the form

$$\operatorname{tr} J = 4w_0\, 1 + 4(w_1 + 4w_2)\mathcal{O}_{\mathrm{scale}} = 4u.$$

Thus $u = w_0\, 1 + (w_1 + 4w_2)\mathcal{O}_{\mathrm{scale}}$, as required. The scalar-coordinate map is injective: expectation kills $\mathcal{O}_{\mathrm{scale}}$ and sends the unit to one, forcing the unit coefficient to vanish; the scalar field is nonzero by its positive free two-point norm, forcing the remaining coefficient to vanish. The running inverse scaling preserves conjugation and signed symmetries, so it preserves this exact scalar membership.

The physical map K_i has the exact formula

$$(K_i T)_a = B_i r_i^5\, U_{i,5}^{-1}((\operatorname{div} T)_a). \tag{TR49}$$

The real symmetry restrictions make this a map between the stated real tensor spaces. Its prefactor is positive and $U_{i,5}$ is invertible, so $K_i T = 0$ if and only if $\operatorname{div} T = 0$. This proves the kernel identification used above. The bare-to-physical degree-five map s_i is likewise the componentwise

renormalized lift followed by $U_{i,5}^{-1}$. Equivariance under conjugation and signed symmetries makes its image lie in the free divergence target. Applying the componentwise divergence comparison to these exact formulas yields (TR46).

We give the finite-dimensional repair which supplies (TR47). This is needed because the limiting degree-five map has a one-dimensional kernel. Let E, H be finite-dimensional real spaces, $f : E \to H$ a surjection with $\ker f = \mathbb{R}x$, $x \neq 0$, and choose a right inverse U. Suppose $s_i : E \to H$ have exact right inverses $u_i : H \to E$ eventually, with $u_i \to U$. Suppose also that $w_i \to x$ and $s_i w_i \to 0$. There is a linear functional $A : E \to \mathbb{R}$ such that

$$y = Ufy + A(y)x.$$

Indeed $y - Ufy$ lies in the one-dimensional kernel and its scalar coordinate is linear. Set $V_i = u_i f + A \otimes w_i$. Then $V_i \to I_E$, hence V_i^{-1} exists eventually and tends to I_E. Also

$$s_i V_i y = fy + A(y)s_i w_i \longrightarrow fy.$$

Multiplication by V_i^{-1} gives $s_i \to f$ in operator norm. Pointwise convergence is sufficient in this argument because all operator spaces are finite-dimensional: evaluate on a fixed finite basis and reconstruct the operator continuously.

Let $\mathcal{L} : E \to \mathbb{R}$ be the bare scalar functional and e_i the required precision scale. Assume the stronger two estimates

$$e_i s_i w_i \to 0, \qquad e_i \mathcal{L}(w_i) \to 0. \tag{TR50}$$

Define $\ell_i = \mathcal{L}u_i$. These functionals converge to $\mathcal{L}U$, so they are uniformly bounded. Put $Q_i = e_i(\ell_i s_i - \mathcal{L})$. On the vector $V_i y = u_i fy + A(y)w_i$, the right-inverse identity cancels the first part exactly and gives

$$Q_i V_i y = A(y)\left[\mathcal{L}u_i(e_i s_i w_i) - e_i \mathcal{L}(w_i)\right] \longrightarrow 0.$$

Since $V_i^{-1} \to I_E$, it follows that $Q_i \to 0$ in operator norm. This proves (TR47) with $e_i = t_i$. It also proves the convergence of s_i needed for (TR46)–(TR48); no prior uniform convergence on its near-kernel direction was assumed.

In the present application x is the nonzero free equation-of-motion vector, the kernel of the free degree-five map. The exact right inverses u_i are obtained by lifting free target vectors at the same retained scale and undoing the renormalized coordinate map. The projection-lift identity makes $s_i u_i = I$ exactly; the common scaling comparison gives $u_i \to U$. The near-kernel vectors w_i are the issued equation-of-motion coordinates. Their two paid estimates are precisely (TR50). The bare scalar functional annihilates the limiting equation-of-motion vector. Thus this argument constructs the detector used in (TR42), with the precision scale $t_i = \mathrm{Age}_i^{1/5}$.

For the correction candidates the sharper scalar estimate in (TR47) is compatible with their qualitative coordinate limits for an additional reason. Before taking the continuum limit their issued coordinates differ from the ideal orbit coordinates by at most $\mathrm{loss}_i = \mathrm{Age}_i^{-1/4}$. The ideal scalar reading equals the target scalar reading exactly. Bounded linearity therefore bounds the scalar discrepancy by a constant times loss_i. Multiplication by $t_i = \mathrm{Age}_i^{1/5}$ gives the decaying power $\mathrm{Age}_i^{-1/20}$. This proves the third error limit used after (TR47). The qualitative joint degree-four and degree-five coordinate limits follow by taking the same finite-level orbit comparison, whose coefficients approach the prescribed vector, and using the vanishing unweighted loss before passage along the original extraction filter.

For the principal divergence coordinates, the explicit prelimit estimate which supplies the final limit in (TR48) is

$$\begin{aligned} \|v_i\| &\le 2\|\mathcal{O}\|\,\mathrm{Age}_i^{1/6}, \\ |\mathcal{L}(v_i)+g| &\le \|\mathcal{M}\|\mathrm{Age}_i^{-1/4} + 4\|\mathcal{M}\|\mathrm{Age}_i^{-1} + \tfrac{1}{12}H_i^{-30} + \tfrac{1}{12}H_i^{-100}. \end{aligned} \tag{TR51}$$

Here $\mathcal{O}$ is the fixed orbit map on degree-five coordinates and $\mathcal{M} = \mathcal{L}\mathcal{O}$. To pass this bound to the continuum coordinates, first compare the issued representative with its projection to the retained coordinate core, with error $\mathrm{Age}_i^{-1/4}$. The retained representative has norm at most $\mathrm{Age}_i^{1/6}$. Apply the bounded orbit map for the first bound and the bounded functional for its error contribution. The finite-level scalar calculation gives the remaining three terms and the value $-g$. Continuous linear maps and closed norm inequalities preserve these bounds along the extraction limit. Every term on the right of the second inequality tends to zero. The scalar finite-level calculation is developed below, and (TS1)–(TS5) supply the weighted near-kernel estimates required in (TR50).

The scalar functional has an explicit invariant-coordinate formula. For the bare divergence vector parametrized by $w \in \mathbb{R}^3$,

$$\rho(\mathrm{bareVector}\, w) = -4(w_0 + 4w_1). \tag{TR52}$$

The coordinate parametrization is an isomorphism for simple $\mathfrak{g}$. The equation-of-motion vector is the image of $(1, -1/4, 1)$, which (TR52) annihilates. More generally, if the vector belongs to the equation ideal, its coordinates satisfy $2(w_0 - w_2) = 0$ and $2w_1 + \frac{1}{2}w_2 = 0$; these force $\rho = 0$. The field-space functional $\mathcal{L}$ is ρ after the degree-five coordinate isomorphism. This proves both its equation-of-motion annihilation and the readings $0, -4$ on the two columns of (TR45).

At finite level the same reading can be computed before quotienting. For a real regular local vector $V = (V_a)$, the coefficient functional $\mathcal{M}$ on its degree-five cards satisfies

$$\mathcal{M}\Big(\theta^{-2}(\mathrm{altoD}\, V_a)_a\Big) = \frac{1}{12}\,\mathrm{Re}\,\mathrm{Holt}(V). \tag{TR53}$$

Here θ is the actual cutoff threshold. The local-card identity gives $\mathrm{altoD}\, V = \theta^2\,\mathrm{DCard}\, V$; the factors cancel exactly. The card reading identifies the coefficient contraction with the local polynomial contraction Holt. Reality and nonwrapping of the support justify that identification without periodic alias terms. Thus (TR53) is a literal coefficient identity with the stated cutoff, rather than a remainder estimate.

Here are the local estimates underlying the coefficient calculation. Write x for the large coupling parameter and put

$$\begin{aligned} \Omega_a(x) &= \exp\{a(\log\log x)^2\}, \qquad \theta = x^{-1/2}\Omega_1(x), \\ c_x &= \theta^2\Omega_{-2}(x), \quad t_x = \theta^2\Omega_{1/8}(x), \quad \sigma_x = \theta^2 x^{-1/100}. \end{aligned} \tag{LCE1}$$

All assertions in this paragraph concern sufficiently large $x > 1$. Products of the Ω_a add their subscripts; for every fixed a and $\varepsilon > 0$, $\Omega_a(x) \le x^{\varepsilon}$ eventually. This follows by dividing $a(\log\log x)^2$ by $\log x$, whose ratio tends to zero.

The regular local fields entering the exact split have the following derivative envelopes on the common healthy chart: the nonlinear principal discrepancy and the finite Taylor surgery error have envelope c_x; the principal comparison field and its loop field have envelope t_x; the finite-loop tail has envelope c_x. These are envelopes through the prescribed chart order, on the stated local support, including measurability, gauge covariance and chart regularity. The large-x threshold is common to the stated local estimates. At each fixed admitted x, exhaust the finite loop and tail

shelves after specifying the required positive error allowance. Absolute convergence makes their errors arbitrarily small. These shelves may depend on x; a single finite truncation uniform for all large couplings is not needed.

For clarity, the nonlinear estimate has a quantitative Taylor proof. Let $s = \theta(1 + \mathrm{high}(x))^P$, where P is the fixed stage power. In exponential link coordinates the discrepancy K between the nonlinear curvature expression and its linear comparison has derivatives through order four at zero bounded by x^{-30}. Its derivatives through the required higher order on $\|y\| \le Ds$ are bounded by $2\Omega_{1/200}(x)$; $D \ge 1$ is a fixed chart radius. The change from color coordinates to link directions has norm at most one, so these low-jet bounds persist under that change. Taylor's formula with integral remainder, after rescaling $y = sp$, gives, uniformly for $\|p\| \le D$,

$$\|\operatorname{minor}_s(\beta K)\|_{C^k} \le 4x\, 2\Omega_{1/200}(x)\, C_{5,D} s^5 \le \sigma_x, \qquad k \ge 5, \quad |\beta| \le 4x. \tag{LCE2}$$

The low-order terms fit this same estimate because $x^{-30} \le 2\Omega_{1/200}(x) s^{5-j}$ for $0 \le j \le 4$. The final inequality follows from $s^5 = x^{-5/2}$ times a fixed subpower factor, whereas $\sigma_x = x^{-1-1/100}\Omega_2(x)$. Thus the fifth-order Taylor remainder has a strictly decaying power to spare. The finite-difference surgery uses $h = \theta x^{-1/24}$. Since $h/s \le x^{-1/24}$, choosing its order $m \ge 48$ makes every fixed subpower multiple of $s(h/s)^{m+1}$ at most $\sigma_x h$ eventually. This controls the surgery error after division by its difference step. The support margins and nonwrapping conditions used here are the finite chart conditions already imposed when choosing the station; none is a new continuum hypothesis.

We spell out how these local bounds survive projection and summation. Fix the degree-five frame and its local section, and denote their fixed reading prices by $r \ge 1$ and $v \ge 1$. If the section has growth exponent $p \ge 0$, put $K_0 = (1+p)r + v$. A field with derivative envelope c_x has weighted coefficient size at most rc_x. This is the coefficient projection estimate: each coordinate is a bounded derivative reading, its fixed reading weight is included in r, and summing the coordinate absolute values bounds the coefficient norm.

For the secondary surgery terms indexed by a finite set I, let Q_i be their supports and g_i their derivative envelopes before multiplying by the difference coefficient. The actual summable tail bound is

$$\sum_{i \in I} e^{s_0 \operatorname{cover}(Q_i)} \left| \frac{\mathrm{teas}_i}{h} \right| g_i \le \sigma_x. \tag{LCE3}$$

Here s_0 is the fixed support weight of the local budget, distinct from the chart radius s. Multiplying each summand by its projection price at most $r+v$ and using the triangle inequality gives coefficient size at most $(r + v)\sigma_x$ for their sum. The same summation constructs its local representative and bounds its approximation remainder; this is why a bound for each individual term is not used in place of the required bound for their sum.

The principal local support has bounded weighted cover, so its support weight is at most a fixed $A = \exp(P_0\ \mathrm{TileNo})$. Its derivative envelope is $2c_x + 2t_x \le 4t_x$. Combining it with (LCE3) bounds both the total projected coefficient and its local representative error by

$$K_0(4At_x + \sigma_x) \le (4A + 1)K_0\theta^2\Omega_{1/8}(x) \le \theta^2\Omega_{1/6}(x). \tag{LCE4}$$

The last inequality follows from $\Omega_{1/8}/\Omega_{1/6} \to 0$. Similarly,

$$K_0\sigma_x \le c_x, \qquad rc_x \le \theta^2\Omega_{-1}(x). \tag{LCE5}$$

Indeed their respective ratios are $K_0 x^{-1/100}\Omega_2(x)$ and $r\Omega_{-1}(x)$, both tending to zero. These inequalities prove the total envelope and all four small coefficient envelopes used next. They also show explicitly why the local errors still vanish after division by θ^2.

We identify those six contributions in terms of the finite functions, to avoid ambiguity between this split and the six antisymmetric tensor differences in (TF0d). At the current scale let Y_b be the localized current, $\widetilde{J}$ the restored tensor, $\mathcal{P}_A$ the finite principal action, and L_Q the finite ordinary loop sum. Write $\mathcal{Q}_Q$ for its augmented quiet sum, with the zero and first path moments removed as in (TF6a)–(TF6c). Define $D_Y f$ by differentiating f along the left logarithmic link current, and write $\operatorname{div}_\nu Y$ for its Haar divergence. The six objects are

$$\begin{aligned} W_b &= \operatorname{altoD}\{\beta(\textstyle\sum_a \mathscr{D}_a \widetilde{J}_{ab} - D_{Y_b}\mathcal{P}_A)\}, & B_b &= \operatorname{div}_\nu Y_b - (\operatorname{div}_\nu Y_b)(1), \\ K_b &= D_{Y_b}\mathcal{Q}_Q, & D_b &= D_{Y_b}(\mathcal{Q}_Q - L_Q), \\ E_b &= \operatorname{altoD}\{h^{-1}\mathcal{T}_{m,h}(Y_b, -\beta\mathcal{P}_A - L_Q)\}, & F_b &= \sum_{\gamma\in\mathcal{F}} \operatorname{weighted}_{P_\gamma} H_\gamma. \end{aligned} \tag{LCE6}$$

In this display W, E, F already denote coefficient vectors; B, K, D are scalar configuration functions until altoD is applied. The sign identity is $D_Y L_Q = K - D$, so the loop contribution is $-K + D$. The operator altoD means the degree-five Taylor reader with its actual threshold weights, on the common local support; the secondary readers retain their own supports P_γ. The current here is the localized version of (TF0a), transported through the stationary descent, not an independently chosen continuum vector.

The term $\mathcal{T}_{m,h}$ is a concrete finite-difference remainder. If $t_i > 0$ are the distinct nodes and ℓ_i their Lagrange polynomials on m nodes, use

$$w_i = \ell_i(0)/t_i, \qquad \mathcal{D}_{m,h} f = \sum_{i<m} w_i\{f(t_i h) - f(0)\}.$$

Their exact moment identities are $\sum_{i<m} w_i t_i^{j+1} = 1_{\{j=0\}}$ for $0 \le j < m$, by interpolation of t^j at zero. Therefore $\mathcal{D}_{m,h} f - h f'(0)$ starts in order h^{m+1}, with its integral Taylor remainder. Apply this to the changed current pairing times the exponential action ratio on each active link, and average the color directions over the Haar adjoint orbit of a unit color basis vector. Sum over the finitely many active links. This is precisely $\mathcal{T}_{m,h}$, including its current derivative and density derivative. Its bounds after division by h are (LCE2).

The low-jet estimate for the first term in (LCE6) uses the quantitative stationary error as well as the identity (TF0c). Here is that step in detail. Write $F = \operatorname{boom} J_b(0)$, $R' = D_Y \mathcal{P}_{\mathrm{main}}$, and E for the stationary error. Equality of the analytic jets of $F - E$ and R' through order four gives

$$D^j(F - R')(0) = D^j E(0), \qquad 0 \le j \le 4. \tag{TF12}$$

The issued stationary chart bounds the analytic coefficient norm of E by $\operatorname{shock}(x) = x^{-100}$ on radius $\sigma = \operatorname{east}(x)$. Converting the coefficient bound to ordinary derivatives through order four costs the previously defined analytic derivative price $g_G(\sigma/5, 4)$, called gable in the source. Localizing the tensor to $\operatorname{Meld} J$ costs at most $x^{-70} g_G(v/5, 4)$, where $v = \sigma/100$.

There are two further replacements: the original current Y by its localized current Y_c, and the full principal action by its finite retained action $\mathcal{P}_A$. Use the exact identity

$$D_Y \mathcal{P}_{\mathrm{main}} - D_{Y_c}\mathcal{P}_A = D_{Y - Y_c}\mathcal{P}_{\mathrm{main}} + D_{Y_c}(\mathcal{P}_{\mathrm{main}} - \mathcal{P}_A). \tag{TF13}$$

The current difference has weighted price $p = x^{-70}$; the localized current has price at most $q = \operatorname{trophy}(x) + p$. The full action derivative has bound

$$\begin{aligned} U(x) &= \operatorname{bankCost}(B(x)) K \operatorname{momentPrice}_P(D/2)[\operatorname{GeoSize} + 2 \text{ paths}], \\ T(x) &= \operatorname{bankCost}(B(x)) K x^{-100} + e^{-D \operatorname{rim}(x)/4} U(x). \end{aligned} \tag{TF14}$$

Here $D > 0$ is the issued path width; K, P depend on G and the fixed derivative order four, and the geometric and path prices are fixed by the stage. Cofinal finite path cuts give the tail derivative bound T. The weighted product rule for contracting a current with an action derivative bounds (TF13) by

$$V(x) = \mathrm{vocation}_G(v, p, U) + \mathrm{vocation}_G(v, q, T). \tag{TF15}$$

The prices in (TF15) have the following explicit meanings. Set

$$\begin{gathered} c_G(a) = \#\,\mathrm{Loc}(G,1)\,\frac{e^{2a}6!\,2^4}{a^6}, \qquad g_G(a,k) = (1 + c_G(a))^k, \\ \mathrm{vocation}_G(a, C, D) = 16\, c_G(a/5) C\, g_G(a/5, 3) D. \end{gathered} \tag{TF20}$$

Here $\mathrm{Loc}(G,1)$ is the fixed finite set of first-order link direction types (axis and color); its cardinality has no volume dependence. The lattice shell count bounds the sum of $e^{-a\,\mathrm{distance}}$ over these directions by $c_G(a)$. Expanding a multilinear derivative in its coordinate directions and summing these decays gives $g_G(a,k)$. For a current-action contraction, write $D_Y F = \sum_\ell Y_\ell\, \partial_\ell F$. The localized current has its first three derivatives bounded by $Ce^{-(a/5)\,\mathrm{distance}(\ell,0)} g_G(a/5, 3)$; the action derivative has four derivatives bounded by D and vanishes at the identity. In its fourth product derivative the term requiring a fourth derivative of the current is multiplied by that zero. The remaining binomial coefficients sum to $2^4 = 16$. Summing in ℓ gives exactly (TF20). The vanishing action derivative follows because its one-color linear coefficient is invariant, and simplicity excludes an invariant color vector; it holds for both the full action and its invariant finite tail. This proves the contraction bound in (TF15) using only three current derivatives and four action derivatives. The scheduled radii and cuts give $T = O(x^{-98})$, $V = O(x^{-65})$, and

$$x^{-100} g_G(\sigma/5, 4) + x^{-70} g_G(v/5, 4) + V(x) = O(x^{-64}) \le x^{-30} \tag{TF16}$$

eventually. These bounds are uniform in the admitted finite torus and hold for every sufficiently large finite cut at that fixed stage. Equations (TF12)–(TF16) supply the second, third, and fourth derivatives as well as the zeroth and first ones. They are the quantitative hypothesis used in the rescaled Taylor remainder (LCE2). In that application the higher derivative bound and $x^{-30} \le 2\ \mathrm{thinM}(x) s^{5-j}$, $0 \le j \le 4$, give the order-five rescaled remainder, with $s = \theta(1 + \mathrm{high})^{\text{stage power}}$. Summing the Taylor terms gives the stated $\theta^2 x^{-1/100}$ envelope.

The quiet derivative K has zero scalar contraction by the two path-moment cancellations (TF6a)–(TF6c), whose full-family justification is given immediately below. The derivative of the appended quiet correction is D. The term B keeps the Haar Jacobian, and its centered constant is explicitly removed. The last term is obtained by expanding each finite surgery exponential into support factors, with coefficient w_i/h; its entire weighted sum is bounded in (LCE3). Thus the four small terms are exactly $W, \mathrm{altoD}\, D, E, F$, and the two non-small terms have their actual scalar calculations.

For completeness, the retained quiet moments are small because the *full* issued loop family has zero moments. For each ordered pair a of plaquette types and pair b of color indices, define

$$c_\gamma^{a,b} = \mathbf{1}_{\{\text{types of } \gamma = a\}}\, \mathrm{coeff}_\gamma(b), \qquad z_\gamma = \mathrm{offset}_{\gamma,1} - \mathrm{offset}_{\gamma,0}\,. \tag{TF17}$$

The path bounds include absolute summability of these coefficients and of each $c_\gamma^{a,b}(z_\gamma)_l$, together with their normalization

$$\sum_\gamma c_\gamma^{a,b} = 0, \qquad \sum_\gamma c_\gamma^{a,b}(z_\gamma)_l = 0 \quad (l = 1, \dots, 4). \tag{TF18}$$

These are curvature loop value and first-jet normalizations, distinct from the potential jet normalization. They are carried in the issued path data as the summable loop terms and their zero sums; joining path families preserves them by absolute summation. The path-data construction retains these identities along with the quadratic metric normalization, so (TF18) applies to the family actually used in (LCE6).

Let L be the retained finite set. Subtracting its sum from (TF18) gives

$$\sum_{\gamma\in L} c_\gamma^{a,b} = -\sum_{\gamma\notin L} c_\gamma^{a,b}, \qquad \sum_{\gamma\in L} c_\gamma^{a,b}(z_\gamma)_l = -\sum_{\gamma\notin L} c_\gamma^{a,b}(z_\gamma)_l. \tag{TF19}$$

Since $|z_{\gamma,l}| \le \operatorname{length}\gamma$ and the coefficient is bounded by its atom budget, an omitted bound $\sum_{\gamma\notin L}\operatorname{budget}_a(\gamma)\,(1+\operatorname{length}\gamma) \le A$ bounds every retained moment in (TF19) by A. Thus the coefficients appended at displacement zero and the four unit displacements in (TF6a) have total coefficient magnitude at most $5\,\#\{\text{color pairs}\}\#\{\text{type pairs}\}(A+4A)$. Indeed there are five displacement choices for each type pair, and each coefficient has magnitude at most $A+4A$; the unit-displacement coefficients individually have the sharper bound A. Their paths have length at most ten, so a weighted atom budget at width a costs at most an additional e^{10a}. Fixed-order local function and derivative bounds follow by the finite plaquette product estimate used for (TF11). The exponential path estimate gives $A = q + e^{-\sigma r/4}\,\mathrm{momentPrice}_1(\sigma/2)B$ after a sufficiently large cut and radius r, for width $\sigma > 0$ and full loop price B. This proves the small appended-tail bound used for D. The augmented finite family, meanwhile, has *exactly* zero constant and first moments by its explicit coefficients. Its Hessian against a supported direction and a direction with affine curl therefore vanishes by (TF6b), proving the scalar cancellation for K. Smallness of the correction and exact cancellation of the augmented family are separate consequences of (TF18) and (TF6a), respectively.

The issuance is constructive: for each of these regular scalar functions subtract the local section of its weighted Taylor card and put the masked difference in one regular activity slot, as in (TF5). Add the four principal formatted insertions and the finitely many secondary ones. Their coefficient is exactly the sum of their Taylor cards by linearity, and their support-weighted prices add. Equations (LCE3)–(LCE5) then prove both the bound on the issued coefficient and the paid remainder of its representative. This proves existence of the candidate with its coefficient identity and its budgets simultaneously.

The principal issued divergence decomposes into six vector contributions:

$$V = W + \operatorname{altoD} B - \operatorname{altoD} K + \operatorname{altoD} D + E + F. \tag{TR54}$$

The decomposition is exact because the local split satisfies $y + D = K$, and the regular-card map is additive on this regular split. The central term K has zero Holt contraction by the finite-level cancellation identity, so its scalar reading is zero by (TR53). The term B gives $\operatorname{Re}\operatorname{Holt}(B)/12$. Its difference from the finite-level scalar payment $\operatorname{midPay}/12$ is at most calorie/12. Each of the four remaining terms $W, \operatorname{altoD} D, E, F$, after multiplication by θ^{-2}, has coefficient norm at most $\operatorname{Wave}(-1)$. Applying the operator norm of $\mathcal{M}$ and summing gives

$$\left|\mathcal{M}(\theta^{-2}V) - \frac{\operatorname{midPay}}{12}\right| \le 4\|\mathcal{M}\|\operatorname{Wave}(-1) + \frac{\operatorname{calorie}}{12}. \tag{TR55}$$

The total vector's coefficient-size envelope is $\theta^2\operatorname{Wave}(1/6)$, so $\|\theta^{-2}V\| \le \operatorname{Wave}(1/6)$. These estimates use the coefficient-size bound, which dominates its norm, and the actual positive threshold to cancel θ^2. The scalar payment comparison $|\operatorname{midPay} - \operatorname{landPay}| \le \operatorname{shock}$, with $\operatorname{landPay}/12 = -g$, gives the last two terms and limiting value in (TR51). This reduces that limit to the local scalar

cancellation, the four explicit local envelopes and the payment comparison; it introduces no independent continuum assumption.

The finite scalar payment itself is obtained by differentiating a massive comparison, with the mass parameter kept strictly positive until after the large-volume estimates. Let $\mathcal{R}_i(\mu)$ be its scalar resolvent primitive and $\mathcal{S}_i(\mu)$ its differentiated shell quantity. The source proves

$$\frac{d}{dx}\operatorname{Re}\mathcal{R}_i(e^x) = -\operatorname{Re}\mathcal{S}_i(e^x), \qquad |\operatorname{Re}\mathcal{R}_i(e^x) + Lx| \le C \quad (x \le 0), \tag{TR56}$$

eventually for every fixed x, where $L = -12g$ and the same C works for every such x. The derivative is a derivative of a convergent finite-volume resolvent expression; strict positivity of its massive matrix gives its inverse throughout this calculation.

For any fixed $a < b \le 0$, the real mean value theorem gives a point $c_i \in (a, b)$, hence $\mu_i = e^{c_i} \in [e^a, e^b]$, with

$$|\operatorname{Re}\mathcal{S}_i(\mu_i) - L| \le \frac{2C}{b-a}. \tag{TR57}$$

Only the two endpoint estimates in (TR56) are used, so their separate eventual assertions can be intersected. No uniform convergence in the mass parameter is needed. The mean-value choice may depend on the finite-volume index.

Let $\mathcal{X}_i$ be the local payment before its final second-derivative correction. The comparison supplies

$$\mathcal{X}_i - \mathcal{S}_i(\mu_i) - \mu_i^2\mathcal{K}_i(\mu_i) \to 0, \qquad |\mu_i^2\mathcal{K}_i(\mu_i)| \le C_1\mu_i, \tag{TR58}$$

for every moving mass in any fixed interval $[\delta, 1]$, $\delta > 0$. To prove $\operatorname{Re}\mathcal{X}_i \to L$, first choose $b \le 0$ so negative that $C_1 e^b < \varepsilon/4$. Then choose $b - a$ large enough that the right side of (TR57) is at most $\varepsilon/4$. The interval lower endpoint e^a is now fixed and positive, so (TR58) applies to the mean-value masses. Its vanishing error and the two preceding bounds prove the desired convergence. The choice order is essential: the volume limit is taken with this fixed positive lower mass, after choosing the error tolerance.

We derive the local resolvent bound in (TR58). Use a periodic lattice of side n, angular momentum spacing $h = 2\pi/n$, and normalized momentum weight n^{-4}. Write $k_x = h\,\mathrm{lift}(x)$ for its centered momenta. If $R \ge h$, any subset of momenta with $|k_x| \le R$ has at most

$$(3R/h)^4 \tag{LR1}$$

members. Indeed their centered integer coordinates have absolute value at most $\lfloor R/h \rfloor$. The lift is injective, and the integer coordinate box has at most $(2\lfloor R/h \rfloor + 1)^4 \le (3R/h)^4$ points. This upper bound remains true when the coordinate box wraps on the torus.

For $h \le \mu \le 1$, group momenta according to

$$2^j\mu \le |k_x| + \mu \le 2^{j+1}\mu, \qquad j \ge 0.$$

Assign boundary points to one of their adjacent groups. A group has at most $(6 \cdot 2^j\mu/h)^4$ points by (LR1); its summand $(|k_x| + \mu)^{-5}$ is at most $(2^j\mu)^{-5}$. Therefore

$$n^{-4}\sum_x (|k_x| + \mu)^{-5} \le \frac{6^4}{(2\pi)^4\mu}\sum_{j\ge 0} 2^{-j} \le \frac{2 \cdot 6^4}{(2\pi)^4\mu}. \tag{LR2}$$

Only finitely many groups are nonempty at a given n; extending their sum by zeros justifies the infinite geometric majorant.

Here is the integrand estimate to which (LR2) applies. At $z = |k| + \mu > 0$, scale an external momentum direction p by z. The zero-field matrix, one-field vertex and two-field vertex have sizes $z^2, z, 1$ with bounded scaled derivatives through the finite orders in use. These follow from their vanishing origin jets and the analytic symbol bounds proved in the massive-symbol section. Ellipticity gives the inverse size z^{-2}. Differentiating $A^{-1}A = I$ gives

$$D(A^{-1}) = -A^{-1}(DA)A^{-1},$$

$$D^2(A^{-1})[u,v] = A^{-1}(DA[v])A^{-1}(DA[u])A^{-1} + A^{-1}(DA[u])A^{-1}(DA[v])A^{-1} - A^{-1}(D^2A[u,v])A^{-1}.$$

These formulas, and one further differentiation when needed, retain the same inverse size under scaled derivatives. Hence the contracted resolvent-insertion integrands have scaled analytic-jet bounds of size z^{-2}, with constants depending on the fixed directions and finite matrix dimensions but not on k, μ, n.

For such an analytic integrand V, let $[V]_j(p) = D^jV(0)[p,\dots,p]/j!$ be its homogeneous Taylor coefficient. The scaled bound gives

$$z^3|[V]_3(p)| \le Cz^{-2}/3!, \quad \text{hence} \quad |[V]_3(p)| \le C'z^{-5}.$$

For the second coefficient it gives $|[V]_2(p)| \le C'z^{-4}$. The angular cube and $\mu \le 1$ give a fixed upper bound $z \le R_*$, so this too is bounded by $C'R_*z^{-5}$. Summing the finite matrix entries and applying (LR2) bounds both contracted coefficients by a constant times μ^{-1}.

The local readings are explicit finite coefficient selectors. For a scalar function on integer arguments set

$$\operatorname{Pan} f = \tfrac{1}{2}(f(1) - 2f(0) + f(-1)), \qquad \operatorname{Flop} f = \tfrac{1}{12}[8(f(1) - f(-1)) - (f(2) - f(-2))].$$

On powers of degree at most two, Pan selects the quadratic coefficient; on powers of degree at most three, Flop selects the linear coefficient. The mixed selector is $i \operatorname{Pan}_a \operatorname{Flop}_b$, evaluated on the specified two-direction momentum contour, with the corresponding powers of h^{-1}. Its fixed finite coefficients bound it by a finite sum of the directional third coefficients just estimated. The local second selector is Pan on the opposed-momentum contour, with h^{-2}. Analytic Taylor remainder bounds show that each selector differs from its indicated derivative coefficient by a term tending to zero as $h \to 0$, with mass restricted to a fixed $[\delta, 1]$. For $\delta > 0$, the massive matrices have a common positive lower bound, so their analytic inverse neighborhoods and the higher derivatives needed for these remainders are bounded throughout this interval. This also allows a mass μ_n depending on n. Absorb a final selector error at most 1 into μ_n^{-1}, since $\mu_n \le 1$.

The physical local lens is the prescribed half-weighted physical resolvent contraction minus its gauge comparison contraction. Both have the preceding bounds; their matrix tickets identify the actual finite matrices with the Fourier inverse matrices. Linearity and the triangle inequality therefore prove, with constants fixed before the positive lower mass cutoff is requested,

$$\exists C_{d,a,b} \ge 0 \;\; \forall \delta > 0 \; \forall (\mu_n)_n, \quad [\delta \le \mu_n \le 1 \; \forall n] \implies |\mathcal{K}_{d,a,b,n}(\mu_n)| \le C_{d,a,b}\mu_n^{-1} \quad \text{eventually.} \tag{LR3}$$

The eventual threshold may depend on δ and the moving mass family. Once δ is fixed, $h \le \delta \le \mu_n$ eventually, as required in (LR2). Sum (LR3) over the finitely many tensor indices. Multiplication by μ_n^2 yields $C_1\mu_n$, the remainder bound in (TR58).

For the local second-derivative correction, its exact massive comparison has error tending to zero at each fixed $0 < \mu \le 1$, and its massive term is at most $C\mu^2\mu^{-1} = C\mu$ by the second-selector version of the same argument. Given $\epsilon > 0$, choose $0 < \mu \le 1$ with $C\mu < \epsilon/2$, and only then take

the volume large enough that the comparison and selector errors total less than $\epsilon/2$. This proves vanishing of that correction, including the actual second Taylor coefficient rather than merely its finite difference. Consequently

$$\mathrm{midPay}_i = \mathrm{Re}\left(\mathrm{crest}_i - \tfrac{1}{2}\sum_{d,a}\mathrm{floft}_{d,a,i}\right) \longrightarrow -12g. \tag{TR59}$$

The crest comparison identifies its limit with that of $\mathcal{X}_i$; each displayed correction tends to zero. This proves the scalar payment limit used to choose the finite-level accuracy in (TR55).

The exact quiet-loop cancellation in (TR54) is local. For a finite roster of gauge-invariant loop functions p_α, their periodized sum is

$$\mathcal{P}(U) = c_{\mathrm{reset}} \sum_{\alpha} \sum_{\text{seats } s} p_\alpha(s \cdot U). \tag{TR60}$$

Each summand is analytic on the permitted group charts, so the finite sum is analytic and gauge invariant. If Y_a is an analytic gauge-covariant vector field supported in a nonwrapping finite neighborhood, its directional derivative $r_a = Y_a\mathcal{P}$ is again analytic, local and gauge invariant. At the identity it vanishes, since the vector field does. Gauge covariance follows by differentiating the identity for $\mathcal{P}$ after conjugating the infinitesimal direction; locality follows by differentiating only links in the common support.

The quiet roster is chosen so that the Hessian of (TR60) pairs to zero between a direction supported in that neighborhood and the prescribed distant polynomial directions of degree at most two. The loop lengths are bounded by M; the separation condition $2(t+1+1+M)+6 < n$ ensures this statement uses actual separated links rather than wrapped copies. The derivative of Y_a at the identity maps every direction into the same supported neighborhood. Differentiating $r_a = Y_a\mathcal{P}$ twice therefore expresses its Hessian pairing as two pairings of the Hessian of $\mathcal{P}$, each with one supported direction. Both vanish. In the third local difference quotient entering Holt, the remaining terms with constant-direction legs vanish by gauge invariance and the local Ward identity; the other terms are exactly these zero Hessian pairings. The finite difference quotient is consequently zero. The local coefficient identity expressing Holt as the finite sum of these third differences proves $\mathrm{Holt}((Y_a\mathcal{P})_a) = 0$.

To obtain the quiet roster from the original loop sum, the construction adds its finite boundary compensation roster. If S denotes the original truncated loop sum and U this compensation, the quiet sum is exactly $P = S + U$. Differentiation gives $Y_aS = Y_aP - Y_aU$, the split $y = K - D$ in (TR54). The compensation has the small boundary envelope calorie; the original truncated derivative has its larger local envelope. The compensated roster retains gauge invariance and the common loop-length bound, so the preceding cancellation applies. This accounts for the exact zero term and for its separate small boundary correction.

The derivative in (TR56) can also be read directly from its determinant definition. Its local primitive is a finite contraction of

$$\tfrac{1}{2}\log\det(A(U)+\mu^2 I) - \log\det(B(U)+\mu^2 I). \tag{TR61}$$

The first matrix contains the physical quadratic term and gauge-fixing term; the second is its gauge comparison matrix. For a fixed positive μ, both are invertible near the identity. The real matrix derivative formula gives

$$\frac{d}{d\mu}\log\det(F+\mu^2 I) = 2\mu\ \mathrm{tr}(F+\mu^2 I)^{-1}. \tag{TR62}$$

Joint analyticity in the mass increment and the link chart permits the finite local coefficient derivatives to commute with this mass derivative. Apply (TR62) to (TR61), then the normalized local contraction $\mathrm{yard}_{d,a} = |\mathrm{sites}|^{-1}\,\mathrm{cob}_{d,a}$, and sum over the ordered off-diagonal axis pairs $d \neq a$. The resulting derivative is $-\mu$ times the corresponding resolvent contraction. Substitution $\mu = e^{x}$ supplies the second factor of μ, proving the derivative identity in (TR56).

For its bound, the period-limit coefficient comparison identifies the normalized local contraction with the massive momentum integral up to an eventual error at most one. The latter obeys

$$|\alpha(\mu) + g\log(1/\mu)| \leq C_0 \quad (0 < \mu \leq 1),$$

where C_0 is independent of μ and of the allowed regulator segment. The ordered-axis contraction is twelve times this coefficient. Hence

$$|\operatorname{Re}\mathcal{R}_i(\mu) + 12g\log(1/\mu)| \leq 12(C_0 + 1)$$

eventually for each fixed μ, which is (TR56). The momentum-integral coefficient and its uniform bound follow from the logarithmic momentum integral and weighted massive-symbol calculation below.

The roster condition behind that Hessian zero has a simple exact form. For every pair of plaquette axes and color labels, its charges c_α and relative offsets δ_α satisfy

$$\sum_\alpha c_\alpha = 0, \qquad \sum_\alpha c_\alpha(\delta_\alpha)_a = 0 \quad \text{for each axis } a. \tag{TR63}$$

These are zeroth- and first-moment cancellations. For any additive real function F on the offset lattice, $\sum_\alpha c_\alpha F(\delta_\alpha) = 0$, by expressing F in its four coordinate values. The Hessian of each two-leg loop term is the sum of its two cross-leg curvature products. Shift the site summation so its first leg is at the same site for every α. Where the supported direction has nonzero discrete curvature, that site lies in its one-link enlargement. The other direction's curvature, on this neighborhood enlarged by the loop length, is affine:

$$\operatorname{curl} T(x + \delta_\alpha) = C + F(x) + F(\delta_\alpha).$$

Consequently summing the charge-weighted products gives a fixed first-leg factor times $(C + F(x))\sum_\alpha c_\alpha + \sum_\alpha c_\alpha F(\delta_\alpha) = 0$. Both leg orders vanish in this way. The same argument applies to each signed rotation; the supported neighborhood grows by at most one link. Summing over seats and rotations proves the Hessian cancellation used above. Thus the zero in (TR55) ultimately rests on the explicit finite moment identities (TR63), affine curvature of the test directions and the stated nonwrapping margins.

16.6.7 The logarithmic momentum integral

The massive resolvent estimates use a positive mass lower endpoint chosen after the error tolerance and before the volume limit. Shell counting separates the near-zero lattice momenta from the remaining uniformly controlled modes. The later moving-mass statement is obtained by combining those estimates with its scalar payment and selector bounds. It does not interchange the massless and volume limits without control.

We now derive the uniform bound used after (TR62). Use real angular coordinates $k \in (-\pi, \pi]^4 = \mathcal{B}$, Euclidean radius $r = |k|$, and normalized momentum weight $c_4 = (2\pi)^{-4}$. The

actual massive two-point symbol gives a continuous scalar integrand $F_\mu(k)$, its half second derivative in the chosen external momentum direction. The momentum coefficient has the exact integral formula

$$\alpha(\mu) = c_4 \int_{\mathcal{B}} \operatorname{Re} F_\mu(k)\, d^4k. \tag{TR64}$$

To obtain this formula from the torus symbol, parametrize the unit-period torus by its half-open fundamental cube and change variables $k = 2\pi x$; the Jacobian is $(2\pi)^4$. The symbol is real analytic in external momentum and internal angular coordinates for fixed $\mu > 0$. Its first and second external derivatives are bounded on the product of a compact external interval and the closed internal cube. Dominated differentiation therefore moves both derivatives under the integral. The half in the symbol's second coefficient and the factor two in the definition of α give exactly (TR64).

The massive symbol estimates have the following two forms, with constants $A, D \geq 0$ independent of $0 < \mu \leq 1$:

$$|F_\mu(k)| \leq \frac{A}{(r+\mu)^4}, \qquad |F_\mu(k) - P_G(k)| \leq \frac{D}{r^3} + \frac{D\mu^2}{r^6} \quad (\mu \leq r), \tag{TR65}$$

where

$$P_G(k) = 4C_{\mathrm{adj}} \left(\frac{k_1^2}{r^6} - \frac{4k_1^2 k_0^2}{r^8} - \frac{1}{r^4} \right). \tag{TR66}$$

The distinct axes $0, 1$ specify the chosen external direction and polarization. The comparison construction establishes the first estimate at scale $r + \mu$, and the second on the subfamily $\mu \leq r$ at scale r; packaging all admissible symbols and masses into these two families gives common constants. The estimates apply to the same mass μ, without shifting it to a preferred radius.

Let $\mathcal{A}_\mu = \{\mu < r \leq 1\}$. In four dimensions the unit ball has volume $\pi^2/2$, so polar integration gives

$$\begin{aligned} \int_{\mathcal{A}_\mu} r^{-4}\, d^4k &= 2\pi^2 \log(1/\mu) =: J_\mu, \\ \int_{\mathcal{A}_\mu} r^{-3}\, d^4k &= 2\pi^2(1-\mu), \\ \int_{\mathcal{A}_\mu} r^{-6}\, d^4k &= \pi^2(\mu^{-2} - 1). \end{aligned} \tag{TR67}$$

Rotational invariance determines the required angular moments without a sphere normalization convention. Equal diagonal moments and $\sum_a k_a^2 = r^2$ give $\int_{\mathcal{A}_\mu} k_a^2/r^6 = J_\mu/4$. For $a \neq b$, rotate a unit vector to $(e_a + e_b)/\sqrt{2}$ and to $(e_a - e_b)/\sqrt{2}$. Add their fourth-power identities; the odd terms cancel and give $\int k_a^4/r^8 = 3\int k_a^2 k_b^2/r^8$. Summing the mixed moments over b then gives

$$\int_{\mathcal{A}_\mu} \frac{k_a^2 k_b^2}{r^8}\, d^4k = J_\mu/24 \quad (a \neq b).$$

Substitution into (TR66) yields

$$c_4 \int_{\mathcal{A}_\mu} P_G(k)\, d^4k = c_4\, 4C_{\mathrm{adj}} \left(\tfrac{1}{4} - \tfrac{4}{24} - 1 \right) 2\pi^2 \log(1/\mu) = -\frac{11 C_{\mathrm{adj}}}{24\pi^2} \log(1/\mu) = -g \log(1/\mu). \tag{TR68}$$

Split the entire cube into $r \leq \mu$, $\mathcal{A}_\mu$, and $\mathcal{B} \setminus \{r \leq 1\}$. The unit ball is contained in the cube. On the first region (TR65) bounds the integrand by $A\mu^{-4}$, whose integral is at most $A\pi^2/2$. On the

last region it bounds the integrand by A, giving at most A vol($\mathcal{B}$). On the annulus the difference from the pole integrates, by (TR67), to at most

$$2D\pi^2(1-\mu) + D\pi^2(1-\mu^2) \le 3D\pi^2.$$

All inequalities are absolute before summing the regions. Equations (TR64) and (TR68) prove

$$|\alpha(\mu) + g\log(1/\mu)| \le c_4\left[A\pi^2/2 + A\ \mathrm{vol}(\mathcal{B}) + 3D\pi^2\right]. \tag{TR69}$$

This supplies the mass-independent constant in (TR56). It also displays the origin of the factor eleven and the normalization $g = 2b_0$.

The pole (TR66) comes from explicit vector and ghost traces. For a normalized adjoint-color generator T_b, $\mathrm{tr}(T_b^2) = -C_{\mathrm{adj}}$. This is the adjoint Casimir identity in an orthonormal color basis. For the external polarization in axis 1 and momentum te_0, the continuum vector vertex is

$$(p(k)I_4 + Q(t)) \otimes T_b, \qquad p(k) = -2ik_1, \qquad Q(t)_{ab} = 2it(\delta_{b0}\delta_{a1} - \delta_{a0}\delta_{b1}).$$

The ghost vertex is $p(k)T_b$. The free propagators are the scalar inverses $r^{-2}I$, and the shifted denominator is $D(t) = r^2 + 2tk_0 + t^2$. Spatial and color traces factor. Direct multiplication gives $\mathrm{tr}[(pI + Q(-t))(pI + Q(t))] = 4p^2 - 8t^2$. Thus the two bubble traces, near $t = 0$, are

$$H_{\mathrm{gh}}(t) = \frac{p(k)^2(-C_{\mathrm{adj}})}{r^2 D(t)}, \qquad H_{\mathrm{vec}}(t) = \frac{(4p(k)^2 - 8t^2)(-C_{\mathrm{adj}})}{r^2 D(t)}.$$

Their denominators are nonzero for sufficiently small t when $r > 0$. Differentiating twice, using $D'(0) = 2k_0$, $D''(0) = 2$, and $p(k)^2 = -4k_1^2$, gives

$$H''_{\mathrm{gh}}(0) - \tfrac{1}{2}H''_{\mathrm{vec}}(0) = 2P_G(k). \tag{TR70}$$

This proves the pole and its sign, including the half weight of the vector determinant.

For the actual massive symbol the same logarithmic determinant differentiation separates each second coefficient into an upper term and a lower bubble term. Denote their second derivatives by U_v, L_v for the vector matrix and U_g, L_g for the ghost matrix. The exact scalar combination is

$$F_\mu(k) = \tfrac{1}{2}\left[\tfrac{1}{2}(U_v - L_v) - (U_g - L_g)\right]. \tag{TR71}$$

This equality is before any bound. The upper terms and bubble terms each have common bounds of order T^{-4} in the admissible scaled family. Since multiplication by $1/2$ decreases the absolute value, (TR71) yields the first bound in (TR65) upon taking $T = r + \mu$.

On the family $\mu \le r$, take $T = r$. The more precise estimates are

$$\begin{aligned} |U_v| &\le u/T^2 + v\mu^2/T^6, & |U_g| &\le x/T^2 + y\mu^2/T^6, \\ |L_v - H''_{\mathrm{vec}}(0)| &\le A(T + \mu^2/T^2)/T^4, & |L_g - H''_{\mathrm{gh}}(0)| &\le B(T + \mu^2/T^2)/T^4. \end{aligned} \tag{TR72}$$

Combining (TR70)–(TR72) bounds $|F_\mu - P_G|$ by

$$\frac{uT + xT + A + B}{T^3} + \frac{(A + B + v + y)\mu^2}{T^6}.$$

The angular cube has $T \le R_* = 1 + 2\pi$, so

$$D = \max(uR_* + xR_* + A + B,\ A + B + v + y)$$

is a valid common constant in the second bound of (TR65). All constants are chosen on the family of all admissible regulator symbols, masses and internal momenta; restricting afterward to $\mu \le r$ preserves those choices. This proves the scalar assembly of the two bounds. Their upper-term and bubble estimates are the matrix estimates (TR72), whose propagator and derivative bounds are developed next.

16.6.8 Weighted matrix derivatives for the massive symbols

Weighted matrix differentiation retains the order of factors in the finite determinant and symbol calculations. The vector and ghost contributions are combined only with their prescribed multiplicities and Fourier measures; the angular factor eleven belongs to this normalization. The subsequent position Ward tests determine the measured coupling rather than assuming that a bare coupling has already been measured.

Uniformity for moving positive masses The regulator mass can vary along the retained sequence, but it remains strictly positive when the finite inverse and determinant are differentiated. The weighted symbol estimates state how their constants depend on that mass and the annular momentum scale. Only after the corresponding uniform bounds and volume errors have been proved is the mass sent toward zero. This order is what permits assembling the vector and ghost contributions without exchanging a singular limit with an uncontrolled inverse.

We describe the matrix calculus behind (TR72). All matrices here have fixed finite index sets. Use a submultiplicative matrix norm $\|\cdot\|_*$ for which $|\operatorname{tr} M| \le n\|M\|_*$. An analytic one-variable matrix family $A(t)$, at scale $T > 0$ and size $L \ge 0$, has the required weighted bound if

$$T^j \|A^{(j)}(0)\|_* \le CL \qquad (j = 0, 1, 2), \tag{TR73}$$

with one common C for the whole indexed family. Product differentiation gives sizes $L_1 L_2$, with constant at most $4C_1C_2$: the second derivative has the four terms $A''B, A'B', A'B', AB''$. Sums add constants, and scalar multiplication multiplies the size by a bound on the scalar.

If $A(0)$ is invertible, A has size $L > 0$, and $\|A(0)^{-1}\|_* \le B/L$, then the inverse has size L^{-1} at the same scale. Indeed,

$$(A^{-1})' = -A^{-1}A'A^{-1}, \quad (A^{-1})'' = 2A^{-1}A'A^{-1}A'A^{-1} - A^{-1}A''A^{-1}.$$

The three weighted bounds are at most B/L, CB^2/L, and $(2C^2B^3 + CB^2)/L$. Their sum is a valid common constant. Analyticity of the inverse near zero follows either from its determinant formula or from a local geometric inverse. For two invertible families the exact identity

$$A^{-1} - B^{-1} = -A^{-1}(A - B)B^{-1} \tag{TR74}$$

transfers a relative error size EL into inverse error size E/L. Products of several differing factors are compared by replacing one factor at a time; (TR73) controls every term. These statements retain the raw derivatives, including their binomial coefficients.

The underlying symbol has a zero-field matrix $A(k)$, one-field vertices $V(q, k)$ and their incoming counterparts $V(-q, k+q)$, and the two-field vertex $W(-q, q, k)$. Along $q = te$, $\|e\| \le 1$, the symbol estimates give sizes $T^2, T, 1$, respectively, in (TR73), for both shifts $k \pm te$. The shift-independent zero-field matrix has the same T^2 bound. Products of symbol families preserve this collection of bounds: a one-field vertex of a product has its two Leibniz terms, and its two-field vertex has the four terms corresponding to both placements and both cross placements. Inversion preserves it with overall size inverted, by (TR74) and the differentiated inverse formulas. This accounts for all placements of the external legs.

These sizes follow from precise Taylor cancellations at the origin. The zero-field matrix differs from $|k|^2 I$ by a function whose jets of orders below three vanish. The one-field vertex differs from its continuum linear vertex by a function whose jets of orders below two vanish. Their derivatives through orders three and two, respectively, are uniformly bounded on the fixed angular ball. Taylor's estimate therefore makes their remainders of orders T^3 and T^2, including their

weighted derivatives through order two along the admissible rays. The continuum one-field vertex itself is linear, so has size T. The two-field vertex has uniform derivatives through order two at the unscaled size one. The symbol's fixed analytic-strip bounds give all these derivative constants by Cauchy estimates, with chain-map norms included. For instance, bound the highest required derivative on the ball, use the vanishing lower jets, and integrate successively along the segment from zero to k to obtain $\|D^j f(k)\| \le C|k|^{u-j}$ for $j \le u$. This proves the stated scaled Taylor bounds.

The admissible scale obeys

$$|k| \le T \le |k| + \mu, \qquad \mu \le T, \qquad 0 < \mu \le 1.$$

On the angular cube the lattice momentum magnitude $s(k)$ satisfies $|k|/4 \le s(k)$. Hence

$$s(k)^2 + \mu^2 \ge T^2/64. \tag{TR75}$$

The ghost massive matrix is this scalar times the identity. The physical vector quadratic form, including its gauge-fixing term, is bounded below by $\min(\varepsilon, 1)(s(k)^2 + \mu^2)I$, where $\varepsilon > 0$ is the common physical ellipticity constant. The elementary coercive inverse estimate gives both inverses norm at most a fixed constant times T^{-2}. Thus the inverse calculus above applies uniformly, even as both momentum and mass tend to zero.

Let $G(k)$ denote the massive inverse and V_-, V_+ the incoming and outgoing one-field vertices. The lower bubble matrix is

$$G(k)V_-(-te, k + te)G(k + te)V_+(te, k).$$

Its size is $T^{-2}T\,T^{-2}T = T^{-2}$. Taking its second derivative and trace by (TR73) gives CT^{-4}, as used in (TR71). The upper matrix is $G(k)W(-te, te, k)$. Split W into its massless vertex and mass-regulator vertex. In the massless part the first inverse is independent of t; its norm is $O(T^{-2})$, while the raw second derivative of the massless two-field vertex is bounded without a scale loss. This part therefore contributes CT^{-2}. The regulator's two-field vertex has size $\mu^2 T^{-2}$ at derivative scale T, and multiplication by $G(k)$ gives size $\mu^2 T^{-4}$. Its second derivative costs another T^{-2}, yielding $C\mu^2 T^{-6}$. This proves the upper estimates in (TR72), and the rough T^{-4} estimate follows from the full sizes and the bounded range of T.

For the bubble comparison take $T = |k|$, $\mu \le T$, and compare with the continuum matrices. The zero-field matrix discrepancy has size $T^3 + \mu^2$. Write

$$E = T + \mu^2/T^2, \qquad T^3 + \mu^2 = ET^2.$$

Equation (TR74) gives inverse discrepancy size ET^{-2}. The vertex discrepancy has size $T^2 + \mu^2/T \le ET$; the continuum vertex has size T. Telescoping the four bubble factors therefore bounds their matrix difference at size ET^{-2}. Its second derivative and trace are at most CET^{-4}, exactly the last two estimates in (TR72). This completes the matrix power counting underlying (TR65); the origin-jet identities and common strip bounds invoked here are the corresponding symbol inputs.

The origin cancellations used in this calculus can be stated independently of the family bounds. For a real analytic function, the statement $f(tx) = O(t^u)$ on every real ray implies that all its jets of degrees below u vanish. Indeed its first nonzero homogeneous Taylor polynomial would be nonzero on some real diagonal, by polarization, contradicting the ray estimate. Conversely, vanishing of these jets and a bounded order-u derivative give the uniform Taylor estimate used in (TR73). Thus ray constants may depend on the ray; only the final derivative bounds must be common.

For the lattice difference symbol $q_a(k) = e^{ik_a} - 1$, $q_a(k) - ik_a$ vanishes to order two and $q_a(k)$ vanishes to order one. Therefore

$$\sum_a q_a(-k)q_a(k) - |k|^2$$

vanishes to order three: write each product discrepancy as the sum of one second-order remainder times a first-order factor and one first-order linear factor times a second-order remainder. The ghost zero-field matrix is this lattice scalar times the identity, proving its required third-order cancellation.

For the vector zero-field matrix, its exact gauge-fixed form is

$$A_{\mathrm{vec}}(k) = d_1(k)^* M(k) d_1(k) + d_0(k) d_0(k)^*.$$

The metric coefficient satisfies $M(0) = I$. Subtract the same expression with $M = I$. Its discrepancy is the product of two first-order exterior differences and the first-order metric discrepancy, hence vanishes to order three. The free exterior identity says that the expression with $M = I$ is the lattice scalar times the identity. Combining with the preceding ghost calculation proves the vector third-order cancellation. This uses the actual quadratic symbol normalization, including its longitudinal gauge-fixing term.

The one-field ghost vertex has the exact component expression

$$\left[-e^{-i(k_a+p_a)} q_a(k) + q_a(-k-p) e^{ik_a}\right] (T_b)_{xy}.$$

Rescale p, k by a real t, divide by t, and let $t \to 0$. The exponentials tend to one and $q_a(tk)/t \to ik_a$, giving $-i(2k_a + p_a)(T_b)_{xy}$. Analyticity then makes the discrepancy from this linear expression vanish to order two. The gauge-fixing vector contribution has corresponding linear expression

$$[-ik_j \delta_{ia} - i(k_i + p_i)\delta_{ja}](T_b)_{xy}.$$

Adding the Wilson curvature vertex gives

$$\delta_{ij}[-i(2k_a + p_a)] + 2i(p_j \delta_{ia} - p_i \delta_{ja}),$$

times T_b, exactly the continuum vertex used in (TR70). The Wilson curvature contribution is obtained by differentiating its group commutator to second order and then applying the third-order chain rule; the Hessian of a commutator at identity is the sum of the two products of first derivatives with their reversed products subtracted. Terms from a group factor with zero first jet vanish. This fixes the curvature-vertex sign.

The permitted perturbing symbol has its one-field part vanishing to order two by its prescribed soft order three. Adding it therefore preserves the same continuum linear vertex. The common strip estimates on the base symbol and the perturbation give the derivative bounds required after these jet identities. This explains how the actual, possibly perturbed, symbol supplies the cancellations in (TR73), while retaining one continuum pole for the entire permitted family.

16.6.9 The determining position-space Ward tests

The coefficient can be read in the position-space frame by a finite determining system. This step distinguishes the formal first even coefficient from the later real remainder estimate. Let $\mathcal{O} = \mathcal{O}_{\mathrm{scale}}$, let T^0 be the normalized free Maxwell stress tensor, and write δv for the tensor $\delta_{ab} v$. For a real invertible degree-four frame P_s preserving the stated conjugation and symmetries, the physical tensor decomposition is

$$P_s^{-1}(s^4 T_{ab}) = \kappa_s T^0_{ab} + d_s \delta_{ab} \mathcal{O} + e_s \delta_{ab} 1. \tag{TR76}$$

Here all three coefficients are real and chosen simultaneously from the exact tensor identity. The physical coefficient result above gives $d_s/u_s^2 \to d_2$ for any coupling u_s with the same leading normalization, in particular for the Wilson charge. Since $u_s \to 0$, it also gives $d_s \to 0$.

Use the separated two-point configuration $x = (e_0, 0)$, with one exterior field $\mathcal{O}$. At the scaled configuration sx, pull both labels through P_s. Define E_s as the first-point divergence of the moment with the stress label P_sT^0; define B_s as the same divergence with $P_s(\delta\mathcal{O})$. Thus B_s is just the time derivative in the first position of the two-scalar moment. These derivatives are with respect to the unscaled configuration x, so the chain rule supplies a factor s multiplying the physical differentiated field. For a smooth shell cutoff χ around the first point, define R_s as the corresponding translation shell-flux test with stress P_sT^0, inserted field $P_s\mathcal{O}$ and exterior field $P_s\mathcal{O}$; define Q_s with stress $P_s(\delta\mathcal{O})$ in its place.

Conservation and the translation Ward identity give exactly

$$\kappa_s E_s + d_s B_s = 0, \qquad \kappa_s R_s + d_s Q_s = B_s. \tag{TR77}$$

To check the scale factors, the shell-flux integral changes by s^3: its volume contributes s^4 and its cutoff derivative contributes s^{-1}. Multiplying (TR76) by this flux functional and using the chain-rule factor s on the translated field yields the same factor s^4 on both sides, which cancels. The unit term has zero divergence and zero translation shell flux; these are exact unit identities. No coincident insertion or contact identity is used in (TR77).

The position-frame jet comparison on compact separated configurations gives

$$B_s \to B_0, \qquad R_s \to B_0, \qquad Q_s \to Q_0(\chi). \tag{TR78}$$

The second limit uses the free translation Ward identity. The first is explicit. Put $c_G = 3\dim\mathfrak{g}/\pi^4$. The free scalar moment on tx is $c_G t^{-8}$, so

$$B_0 = -8c_G \neq 0. \tag{TR79}$$

Consequently the second equation in (TR77), together with $d_s \to 0$, implies $\kappa_s R_s \to B_0$. Dividing by $R_s \to B_0 \neq 0$ gives $\kappa_s \to 1$. Divide the first equation by u_s^2 and by κ_s; then

$$\frac{E_s}{u_s^2} = -\frac{d_s/u_s^2}{\kappa_s} B_s \longrightarrow -d_2 B_0. \tag{TR80}$$

Thus the tested conservation column is known through its first possible even coefficient; the other three columns are needed only at order zero.

For clarity, the determining formal system uses an unknown order-zero stress amplitude κ and a two-entry invariant jet (d_0, d_1), the latter corresponding to powers $1, u^2$. Its verified columns are

$$B = B_0, \quad E_0 = 0, \quad E_1 = -d_2 B_0, \quad R = B_0, \quad Q = Q_0(\chi).$$

The equations are

$$\kappa E_0 + d_0 B = 0, \quad \kappa R + d_0 Q = B, \quad \kappa E_1 + \kappa' E_0 + d_1 B + d_0 B' = 0 \quad \text{for every } \kappa', B'. \tag{TR81}$$

Since $B_0 \neq 0$, the first forces $d_0 = 0$, the second forces $\kappa = 1$, and the third forces $d_1 = d_2$. Conversely those values satisfy every equation. In the signed parity convention this gives the trace jet $(0, 0, -2b_0, 0)$, because the trace multiplies the scalar coefficient by four and $4d_2 = -2b_0$. This formal parity statement alone asserts no real order-four remainder; that stronger estimate requires the normalized frame construction below.

16.6.10 A normalized position coupling and its real remainder

A finite formal parity jet is distinguished from a real remainder estimate. Here the exact normalized coefficient identities give a bounded quotient at order four, and hence an actual $O(g_X^4)$ error. The weaker frame class retains only the stated leading coefficient after its positive scalar branch and nonzero two-point normalization are imposed; the stronger remainder is not transferred to arbitrary frames.

The degree-four frame can be adjusted by two independent positive real factors while fixing the unit. Let C be the centering projection and Π the rotation-invariant projection onto the scalar direction in the centered field space. They satisfy $C^2 = C$, $C\Pi C = \Pi C$, and $\Pi C \Pi C = \Pi C$. The free scalar obeys $\Pi C \mathcal{O} = \mathcal{O}$, while the free Maxwell stress components obey $C T^0_{ab} = T^0_{ab}$, $\Pi C T^0_{ab} = 0$. Define

$$A_{x,y} = x\Pi C + y(C - \Pi C) + (I - C), \qquad x, y > 0. \tag{TR82}$$

The three summands are complementary projections. Consequently $A_{x,y}^{-1} = A_{1/x,1/y}$. It fixes the unit, preserves expectation, commutes with rotations and conjugation, multiplies $\mathcal{O}$ by x and multiplies T^0 by y. These facts follow by applying the three projections successively; in particular none of the centered components feeds the unit component.

Put $P'_s = P_s A_{x_s,y_s}$. Multilinearity of the moments and their derivatives gives the exact transformations

$$E'_s = y_s x_s E_s, \quad B'_s = x_s^2 B_s, \quad R'_s = y_s x_s^2 R_s, \quad Q'_s = x_s^3 Q_s. \tag{TR83}$$

For example the shell test R contains one stress label and two scalar labels, accounting for its three factors. Let $b = B_0$, which is real and nonzero by (TR79). Define, for sufficiently small positive s,

$$x_s = \sqrt{\frac{b}{\operatorname{Re} B_s}}, \qquad y_s = \frac{b}{x_s^2 \operatorname{Re} R_s}. \tag{TR84}$$

The quotients tend to one by (TR78), so $x_s, y_s > 0$ eventually and both tend to one. Outside that final interval take any total choice of frame; only the eventual identities are used. Equations (TR83)–(TR84) impose

$$\operatorname{Re} B'_s = b, \qquad \operatorname{Re} R'_s = b \tag{TR85}$$

exactly. This establishes existence of the normalized frame.

Conversely any positive adjustment of the form (TR82) satisfying (TR85) tends to the identity. Indeed $x_s^2 \operatorname{Re} B_s = b$ implies $x_s^2 \to 1$; positivity chooses $x_s \to 1$. Then $y_s x_s^2 \operatorname{Re} R_s = b$ gives $y_s \to 1$. Thus all such choices preserve the limiting frame and its reality, symmetry, unit and centering properties.

They also preserve every compact separated jet comparison. For any finite tuple of labels, expand $P_s^{-1} P'_s v_j = A_{x_s,y_s} v_j$ in a fixed finite basis. Its coefficients converge to those of v_j. Expand the moment multilinearly in that basis, commute its finite sum with the fixed position derivatives, and bound each difference by the sum of its coefficient error times the bounded free derivative and its bounded coefficient times the original comparison error. Compactness bounds the free derivatives uniformly on the separated compact set. Finitely many terms then give any prescribed error. This proves all the comparisons for P', without a further extraction.

Set

$$c = (-d_2)b \neq 0, \qquad \lambda_X(s) = \frac{\operatorname{Re} E'_s}{c}, \qquad g_X(s) = \sqrt{\lambda_X(s)}. \tag{TR86}$$

This is a measured position-space charge, defined from a separated derivative test rather than from the physical stress coefficient. First apply (TR80) with the independently calibrated Wilson charge

g_W. Equations (TR83) and $x_s, y_s \to 1$ show $\operatorname{Re} E'_s/g_W(s)^2 \to c$. Hence $\lambda_X/g_W^2 \to 1$, so $\lambda_X > 0$ eventually,

$$g_X/g_W \to 1, \qquad \log(1/s)\, g_X(s)^2 \to (2b_0)^{-1}. \tag{TR87}$$

The positive square root fixes the branch. In particular this coupling is calibrated to the Wilson charge before it is used in the Ward equations.

Transfer (TR76) to P', denoting its coefficients again by κ, d, e. The coefficient comparison survives this asymptotically identity change, so $d/g_X^2 \to d_2$. The exact Ward equations (TR77) still hold. Take their real parts, use (TR85), and write $q_s = \operatorname{Re} Q'_s \to q_0 = \operatorname{Re} Q_0$. They give

$$\kappa_s - 1 = -\frac{d_s q_s}{b}, \qquad d_s = d_2 g_X(s)^2 \kappa_s. \tag{TR88}$$

The second equality uses the definition (TR86), not a truncated expansion. Therefore

$$\frac{\kappa_s - 1}{g_X(s)^2} \to -\frac{d_2 q_0}{b}, \qquad \frac{d_s - d_2 g_X(s)^2}{g_X(s)^4} \to -\frac{d_2^2 q_0}{b}. \tag{TR89}$$

For the first limit use $d/g_X^2 \to d_2$ in the first identity; it also implies $\kappa_s \to 1$. The second follows by substituting the second identity into its numerator. Since a convergent real quotient is bounded, (TR89) proves

$$d_s = d_2 g_X(s)^2 + O(g_X(s)^4), \qquad 4d_s = -2b_0 g_X(s)^2 + O(g_X(s)^4). \tag{TR90}$$

It also implies the signed cubic remainder is little o, by dividing the order-four bound by g_X^3 and using $g_X \to 0$. No full perturbative series is required for these assertions.

Finally take the trace and expectation of (TR76). Since the aligned scalar is centered, T^0 is traceless and the unit has expectation one, one obtains $e_s = s^4 \langle \operatorname{tr} T \rangle / 4$. Subtracting the selected-vacuum expectation instead gives the exact centered-field comparison

$$s^4 \Theta = 4 d_s \operatorname{center}_\chi(P'_s \mathcal{O}).$$

This identifies the scalar in the normalized frame with the actual physical trace used in the spectral and exclusion arguments.

The projection in (TR82) is constructed by compact-group averaging. Average the real orthogonal action on degree-four fields against normalized Haar measure, on the left and on the right as in the aligned frame construction, and compose the two averages. Haar invariance shows that its image is fixed by every orthogonal transformation and that it is unchanged when its argument is rotated. It fixes every invariant vector; hence it is idempotent. Since the action commutes with conjugation, the average does also. Invariance of expectation shows that averaging preserves expectation, so it commutes with centering in the required identities. The scalar $\mathcal{O}$ is fixed. Each off-diagonal component of T^0 changes sign under reflection of just one of its indices; its invariant projection is therefore its own negative and is zero. The projected diagonal components are equal by rotation covariance. Their sum is zero because T^0 is traceless; all diagonal projections are consequently zero. This proves every projection property used in (TR82).

The leading scalar coefficient also transfers to the weaker position normalizations asserted by the theorem. Such a frame need only be real, rotation-aligned, unit-preserving, have its centered scalar in the prescribed positive branch, and have the centered scalar two-point norm at the fixed separated test tend to the free nonzero value. The branch classification places its centered scalar on the same positive ray as that of a reference plotted frame:

$$\operatorname{center}(L_s \mathcal{O}) = a_s P_s \mathcal{O}, \qquad a_s > 0.$$

The ratio of the two centered two-point norms equals a_s^2. Both norms tend to the same nonzero free value, so $a_s^2 \to 1$, and positivity gives $a_s \to 1$. Taking the trace and centering in the tensor decomposition yields $s^4 \operatorname{center}(\operatorname{tr} T) = 4d_s^{(L)} \operatorname{center}(L_s\mathcal{O})$, and the analogous identity for P. The centered scalar is nonzero eventually by its nonzero two-point norm. Cancelling it gives $d_s^{(L)} a_s = d_s^{(P)}$, hence $d_s^{(L)}/g_s^2 \to d_2$ for every equally normalized coupling. Unit mixing in $L_s\mathcal{O}$ disappears under centering. This proves precisely the leading transfer, without transferring the stronger order-four bound to arbitrary frames.

All spectral pairings used in this discussion are ordinary integrals at each fixed positive scale. A dilated compact continuous test remains compact continuous and is integrable against a cone-supported measure with finite future Laplace damping, by the compact-test argument following (TR12). A dilated Schwartz test remains Schwartz and is integrable by (TR10)–(TR11). On the limiting cone measure the same bounds apply to its polynomial density $(p_0^2 - |\mathbf{p}|^2)^2$. Thus the compact and Schwartz versions of the spectral limit, and their relative and joint position versions, carry their integrability assertions as well as their limiting identities.

16.6.11 The same scalar diagnostic in the selected vacuum

The scalar diagnostic is first normalized in the unselected hierarchy and then interpreted in the selected vacuum through an exact centered label identity and positivity. Convexity and secants control the selected scalar function. This argument keeps the phase fixed; it does not obtain a desired result by replacing the selected vacuum as the radius changes.

The frame was positioned using the unselected hierarchy. Its interpretation in the selected vacuum follows from the exact centered label identity and positivity. Let $v_s = P_s\mathcal{O}$, let $m_s = \langle v_s \rangle_\chi$, and write $w_s = v_s - m_s 1$. The coefficient comparison gives

$$s^4\Theta = 4d_s w_s, \qquad \log(1/s)\, 4d_s \longrightarrow -1. \tag{TR91}$$

For any fixed separated two-point shape x, rotation and translation reduce the trace moment to $E(s|x_1 - x_2|)$. The short-distance limit (TR2), with the fixed dilation of the logarithm accounted for, gives

$$s^8 \log^2(1/s) \langle \Theta(sx_1)\Theta(sx_2) \rangle_\chi \longrightarrow c_G |x_1 - x_2|^{-8}.$$

Substitute (TR91) and divide by the square of its nonzero limiting coefficient. It follows that $\langle w_s(sx_1) w_s(sx_2) \rangle_\chi$ tends to the free scalar two-point function. The exact unit remainder in the uncentered label comparison gives $m_s = O(s^4/|d_s|)$, hence $m_s \to 0$, since $s^4 \log(1/s) \to 0$ and the second limit in (TR91) is nonzero. The unit insertion rule gives

$$\langle v_s(sx_1) v_s(sx_2) \rangle_\chi = \langle w_s(sx_1) w_s(sx_2) \rangle_\chi + m_s^2. \tag{TR92}$$

Thus the uncentered selected scalar moment has the same pointwise free limit at every separated shape.

The derivative diagnostic requires a further argument; differentiating this pointwise limit would be unjustified. For fixed $s > 0$, the real function

$$U_s(t) = \operatorname{Re} \langle v_s(ste_0) v_s(0) \rangle_\chi, \qquad t > 0,$$

is convex. In fact (TR91) and (TR92) express it as

$$U_s(t) = \frac{s^8}{(4d_s)^2} E(st) + \operatorname{Re}(m_s^2), \tag{TR93}$$

for all sufficiently small s. The first prefactor is nonnegative, and $E(st) = \int e^{-stp^0}\, d\mu(p)$ is convex in t, since each real exponential is convex and its integrals are finite. The additive constant does not affect convexity.

We use the following elementary consequence. If differentiable convex functions U_s on $(0,\infty)$ converge pointwise to V, and V is differentiable at 1, then $U_s'(1) \to V'(1)$. For fixed $a < 1 < b$, convexity bounds $U_s'(1)$ between the secant slopes from a to 1 and from 1 to b. Pass to the pointwise limit in these two fixed slopes, then let $a \uparrow 1$, $b \downarrow 1$. This proves the assertion with the correct order of limits. Here (TR92) gives $V(t) = c_G t^{-8}$, so

$$U_s'(1) \longrightarrow -8c_G = b. \tag{TR94}$$

Smoothness of the selected separated moments identifies this derivative with the real first-position scalar derivative B_s^{sel}.

Apply physical conservation to (TR76), now pairing the resulting derivative identity with the selected scalar field. Derivatives of the unit vanish, exactly as in (TR77). Therefore

$$\kappa_s E_s^{\mathrm{sel}} + d_s B_s^{\mathrm{sel}} = 0.$$

Using $\kappa_s \to 1$, $d_s/g_W^2 \to d_2$, and (TR94) yields

$$\frac{\operatorname{Re} E_s^{\mathrm{sel}}}{g_W(s)^2} \longrightarrow (-d_2)b = c. \tag{TR95}$$

This proves the selected-vacuum diagnostic claimed for the same measured fields. The argument needs only pointwise scalar asymptotics, convexity and the exact conservation identity; it claims no higher-rate comparison between selected and unselected diagnostics.

16.7 The quantitative near-kernel vector for the trace sensor

The near-equation vector is an explicit linear combination of transported degree-five insertions. Its weighted sensor estimate is stronger than bare convergence to the equation subspace. The proof therefore constructs a funded local marker with a genuine finite density variation and replicated root bound. A formal integration by parts symbol, without its local representative and expense estimates, would not supply this weighted conclusion.

The near-kernel vector required in (TR50) is a specific combination of two transported degree-five insertions. Write $\widehat{\beta}_k$ for the prescribed clock and β_k for its extracted coupling; the common calibration gives $\widehat{\beta}_k/2 \le \beta_k$. Let $A_k(c) = \mathrm{AgeK}(c,k)$. For every fixed c,

$$A_k(c)/\beta_k \longrightarrow 0. \tag{TS1}$$

Indeed the clock estimate gives $A_k(c) \le \widehat{\beta}_k^{1/2}$ eventually, while $\beta_k^{-1} \le 2\widehat{\beta}_k^{-1}$. Their product is at most $2\widehat{\beta}_k^{-1/2} \to 0$.

Let P_k be the real orbit projection of the principal degree-five logarithmic insertion, and E_k the corresponding orbit projection of the degree-five divergence insertion associated with the uncorrected stress. They are expressed in the same transported coordinate chart and use the same extracted coupling. Define

$$v_k = \beta_k^{-1} P_k - E_k. \tag{TS2}$$

The principal bounds proved in the finite scalar calculation give $\|P_k\| \le C A_k(1/6)$ and $\mathrm{Lease}(P_k) \to -g$. The divergence insertion satisfies

$$E_k \longrightarrow -\,\mathrm{Stem5}, \qquad A_k(1/5)\,\mathrm{Lease}(E_k) \longrightarrow 0.$$

Here Lease(Stem5) = 0, because its bare vector is an equation-of-motion vector and the scalar sensor annihilates that equation subspace. The limit of E_k is obtained by projecting the same finite three-coordinate stress card; its coordinates tend to the negative equation vector. The scalar equality on these cards is exact, so its transported error has the displayed weighted bound.

Equations (TS1)–(TS2) consequently give

$$v_k \longrightarrow \mathrm{Stem5}, \qquad A_k(1/5)\,\mathrm{Lease}(v_k) \longrightarrow 0. \tag{TS3}$$

For the second assertion expand the linear functional in (TS2). The principal contribution is $(A_k(1/5)/\beta_k)\,\mathrm{Lease}(P_k) \to 0$ by (TS1) and its finite limit; the divergence contribution is the preceding weighted error.

The other required estimate concerns the measured vector $s_i(v_{k_i})$, where k_i is the common retained scaling sequence. Before transport, (TS2) is β_k^{-1} times the orbit of the coefficient difference

$$X_k - \beta_k Z_k.$$

The exact finite local divergence relation makes this difference a served equation insertion. Its coefficient charts and actual lattice functions have common limits by the already established logarithmic insertion comparison. The degree-five core projection commutes with the orbit operation, and the coordinate change is linear, so (TS2) is exactly its transformed, β_k^{-1}-normalized limit; this is an identity, not an asymptotic replacement.

The retained equation estimate on these coefficients gives, for every integer $p \geq 0$,

$$\widehat{\beta}_{k_i}^{\,p}\, s_i(v_{k_i}) \longrightarrow 0. \tag{TS4}$$

Only the first component needs to be estimated before the finite real orbit projection; the remaining components are its prescribed symmetry transforms. In particular take $p = 1$. The clock estimate $A_{k_i}(1/5) \leq \widehat{\beta}_{k_i}$ eventually and positivity give

$$A_{k_i}(1/5) s_i(v_{k_i}) \longrightarrow 0$$

by comparison of norms. Together with (TS3), this supplies both weighted leaks required by the finite-dimensional repair argument (TR50). That argument produces uniformly bounded functionals μ_i, convergence $s_i \to V_{\mathrm{base}}$, and

$$A_{k_i}(1/5)\|\mu_i \circ s_i - \mathrm{Lease}\,\|_{\mathrm{op}} \longrightarrow 0.$$

Thus the quantitative scalar sensor is obtained from an actual near-equation insertion, with its growing weight paid before the stress correction is solved.

The passage from the retained equation insertion to (TS4) is detected by degree-four spectators. At each retained horizon choose the finite local coefficient family for $X_k - \beta_k Z_k$ on the same actual region. Its locality window contains both the equation surgery and the divergence film: the first is contained in the inspected radius-canvasR tower, enlarged by eight; the second in the radius-townD tower. Use

$$\mathrm{Sliver} = \max(\mathrm{canvasR} + 8, \mathrm{townD} + 1).$$

This window is subpower in the clock. Its site count and the total uniform norm allowance are eventually at most the fourth power of the clock. Thus both local representations fit in the common chamber before applying the retained insertion comparison.

On the healthy region, the exact finite equation identity expresses the difference of the two films as the surgery marker plus its centered normalization. The surgery approximation error is

bounded by its small normalization times $e^{-\,\mathrm{prism}}$; after division by that normalization it is at most $e^{-\,\mathrm{prism}}$. Outside the healthy region use the sum of the uniform film bounds, covered by the same fourth-power allowance. These are the local and global bounds used to establish the equation probe estimate; no continuity of a hard healthy-set indicator is asserted.

For each fixed degree-four free spectator w, the resulting two-point estimate after taking the coefficient limit is

$$\Big|S\Big(x_i;\, a_i^5\,\mathrm{project}(X_{k_i}-\beta_{k_i}Z_{k_i}), U_i^{(4)}w\Big)\Big| \le \mathrm{Tau}(j_i,-100). \tag{TS5}$$

The left-hand side is an actual fixed-placement hierarchy limit of its lattice two-point moment. The two lattice positions are the prescribed near point and the origin. They are distinct and have a common nonwrapping placement on the eventual volume tail. The spectator is its retained corrected lift, so the coefficient and correlation limits use the same refinement.

The coordinate identity in (TS2) converts the first insertion of $\widehat{\beta}_{k_i}^{p}s_i(v_{k_i})$ under the scaling frame into

$$\widehat{\beta}_{k_i}^{p}b_i^5\,a_i^5\,\mathrm{project}(X_{k_i}-\beta_{k_i}Z_{k_i}),$$

where the physical scale is a_ib_i. It follows from linearity of the transported chart, the exact cancellation of $\beta_{k_i}^{-1}\beta_{k_i}$, and the definition of the degree-five scaling factor. Consequently (TS5) bounds each tested response by

$$\widehat{\beta}_{k_i}^{p}\,\mathrm{tiltCap}(j_i)^5\,\mathrm{Tau}(j_i,-100)\longrightarrow 0.$$

The clock and tilt factors have rate zero in the horizon variable, whereas the last factor has strictly negative rate. This proves the limit for every fixed p. Finally the real invariant-vector detector, obtained from these degree-four tests and the same frame comparison, converts their vanishing responses into convergence of the measured invariant vector to zero. This is (TS4); the estimate is proved on the exact degree-five near-kernel combination rather than inferred from its unweighted limit.

The marker estimate in this comparison is a finite local sum. At the summit scale $k=k(i)$, put $x=\widehat{\beta}_k$ and let $\rho(x)>0$ be the surgery parameter. For fixed surgery order m, the marker is

$$M_m(U)=\sum_{\alpha\in\mathcal{A}_{m,x}}\frac{w_{m,\alpha}}{\rho(x)}\,L_\alpha(U). \tag{TS6}$$

Here $\mathcal{A}_{m,x}$ consists of the surgery addresses in the prescribed finite box and their finite orbit indices; L_α is the local lens obtained from that surgery. The orbit census gives

$$\sum_{\alpha\in\mathcal{A}_{m,x}}|w_{m,\alpha}|\le \mathrm{cardB}(x)\,\mathrm{wheat}(m).$$

The box census has subpower growth in x, and the surgery parameter bound consequently gives

$$\mathrm{cardB}(x)\,\mathrm{wheat}(m)/\rho(x)\le x^2$$

eventually. On the chosen scale calendar this is at most T_i^{100000}. Each lens has support in the same window of radius $2\,\mathrm{coast}(x)$, that window fits the comparison chamber, and its pullback to the finest lattice has rooted test norm at most $\mathrm{T}_i^{-1900000}$. The local lens estimate applies at every address in the sum because its beacon is in the common enlarged box and within the allowed distance from the origin. The rooted test norm is a seminorm on these chamber-supported functions. Pullback commutes with this finite sum and its scalar coefficients. Thus (TS6) has rooted test norm at most

$$\mathrm{T}_i^{100000}\mathrm{T}_i^{-1900000}=\mathrm{T}_i^{-1800000}.$$

The surgery's pointwise bound also gives $|M_m(U)| \le x^2$ for every configuration. All these assertions hold eventually in i, uniformly in the finer cutoff and volume; m is fixed before that eventual threshold.

We spell out the label transfer as well. Let x_e, y_e be the two local films used in the surgery, and X_e, Y_e their billed degree-five labels. With fine comparison spacing a_i, summit spacing a_k, and actual coupling β, set

$$q = (a_i/a_k)^5, \qquad V_e = X_e - \beta Y_e, \qquad J_e = q(x_e - \beta y_e).$$

The billing approximation, physical normalization, and transport give

$$\| \operatorname{back}(J_e) - I(a_i^5 V_e)\|_{\mathrm{root}} \le \mathrm{T}_i^{-200}. \tag{TS7}$$

Indeed the two individual errors are at most T_i^{-430}; the triangle inequality multiplies their sum by at most $1+\beta$. Since $0 < \beta \le 2\widehat{\beta}_k$, the calendar bound $1+\beta \le \mathrm{T}_i^{230}$ pays this loss.

Let $\mathcal{D}_d$ denote the signed finite orbit sum defining the relevant divergence component. It commutes with the physical insertion map, with coefficient transport, and with pullback on the common window. Applying its rooted seminorm estimate to (TS7) gives the same stated error for

$$\operatorname{back}(\mathcal{D}_d J) - I(a_i^5 \mathcal{D}_d V).$$

The real orbit projection places $a_i^5 \mathcal{D}_d V$ in the required invariant label subspace. To check its coefficient budget, each billed input has transported coefficient size at most 2, hence V_e has size at most $2 + 2\beta$. If $c \ge 1$ is the fixed degree-five orbit cost, its divergence has size at most $c(2+2\beta)$. Eventually $6c \le \widehat{\beta}_k = t$, with $t \ge 1$, so

$$c(2+2\beta) \le 6ct \le t^2 \le t^4.$$

Multiplication by a_i^5 gives exactly the physical coefficient budget. The previously obtained crown bound for $\mathcal{D}_d J$, its chamber support, and (TS7) now give the required physical probe estimate for the corrected divergence label. In particular this transfer retains the coupling multiplying Y, its degree-five spacing factor, and the actual common readout.

Here is how the local lens estimate entering (TS6) is obtained. A single lens is the Haar average of a local surgery difference. Copy it into every chamber of a tiling based at a coarser reporting level J, applying the chamber's signed motion and conjugation. If the original surgery is based at z, each copied base lies at distance $\operatorname{dist}(0,z)$ from its chamber center. The embedded centers are separated by at least the block ratio between levels J and k. Choose that ratio larger than both

$$2\ \mathrm{pierSize} + 2 + 2\operatorname{dist}(0,z), \qquad 64B\,b + 2B + 2\operatorname{dist}(0,z).$$

The inspected surgery supports are then disjoint, the changed bonds are distinct, and the stronger separation needed for the marked expense bound holds. All rotations and optional conjugations preserve the local boundedness and gauge conditions. Integration in one component of a product Haar variable is its ordinary one-component Haar integral, so the marked coefficient at each chamber is exactly that chamber's copy of the original lens.

Introduce two marker slots per chamber and give each the same positive weight $w(k)$. The surgery expense inequalities have the following explicit form. Put s for the support exponent, l for the support-cover bound, F for the bound on the inserted local function, and m for the surgery truncation. With the ordinary and tail tariffs denoted by T_{wing} and T_{lease}, choose $c \ge 1$ and $r \ge 0$ such that

$$4ce^{(s-3)l} \le e^b, \qquad Fe^{(s-3)l+2}\frac{w(k)^2}{c} \le r, \qquad 4e^{-m}T_{\mathrm{lease}} \le 1, \qquad T_{\mathrm{wing}} \le 1.$$

The remaining expense is

$$t = r + 8e^{b}e^{-m}T_{\mathrm{lease}} + 2e^{b}b_{\mathrm{forced}}, \qquad (2B+1)^{4}t \le 1, \qquad e^{t}-1 \le u_{\mathrm{allowed}}.$$

The local surgery expansion and its support census then give a served marked insertion with support exponent $s-4$ and budget 1. The factor $w(k)^{2}/c$ is included in its coefficient bound. The tariff summation is performed separately at each surgery base; disjointness of the inspected supports permits their combination with no additional factor equal to the number of chambers.

Transport this insertion from k down to an intermediate level m_{0}, retaining the two markers per chamber and weights satisfying

$$w(j) \le \delta_{j} w(j+1), \qquad m_{0} \le j < k.$$

At each scale the embedded chamber anchors are distinct. Hence the occupancy of the doubled marker set is at most two, which is within the available multiplicity $r_{j} \ge 2$. This proves the admissibility of the marker transport at every step. The retained insertion estimate and positive scalar normalization bound its final, fully marked coefficient. If v is the number of reporting chambers and Q the physical Wilson moment of the product of the copied lenses, the result is

$$|Q|\, w(m_{0})^{2v} \le \exp[\,|\,\mathrm{Site}_{m_{0}}\,|\,\mathrm{treasury}(m_{0}) + 4v\,]\,.$$

This coefficient identity follows also directly by integrating the restored marked insertion: the partition identity transports the intermediate integral back to the finest Wilson integral, and the same positive Wilson partition factor occurs on both sides and cancels. The chamber-count identity is

$$|\,\mathrm{Site}_{m_{0}}\,|/v = (a_{J}/a_{m_{0}})^{4}.$$

Taking the v-th root therefore proves the local rooted estimate

$$|Q|^{1/v} \le \frac{\exp\big[(a_{J}/a_{m_{0}})^{4}\,\mathrm{treasury}(m_{0}) + 4\big]}{w(m_{0})^{2}}. \tag{TS8}$$

Thus the bound is independent of the number of copied chambers.

For the summit calendar take $J = i$ and m_{0} its prescribed basin. The surgery window radius and the necessary separation radius are subpower functions of $\widehat{\beta}_{k}$; the calendar's block ratio eventually exceeds them uniformly in the volume. Its treasury and weight choices make the right side of (TS8) at most $\mathrm{T}_{i}^{-1900000}$. The same window fits the reporting chamber. This establishes the local estimate used in (TS6), with the finite surgery expense conditions displayed above.

The expense conditions in this argument are consequences of the same large-coupling choice; they are not additional assumptions on the final theory. To see the numerical allocation, fix a stage and the integer l. Set

$$B = \lceil (\log x)^{2} \rceil, \qquad P = B^{40} + 1100, \qquad M = B^{60} + 1100, \qquad e_{x} = e^{-P/8}.$$

These are respectively the prism and surgery truncation used above, and $M \ge 4P$ eventually. For every fixed constant C and integer j,

$$Cx^{j} \le e^{P/32} \quad \text{eventually.} \tag{TS9}$$

Indeed $\log x \le B$ eventually, while the exponent P has degree forty in B; taking logarithms proves the claim. The lease tariff has an upper bound x^{3}, uniformly over the admissible representations at

this fixed stage. Its defining expression is a fixed finite sum of products of the coupling, threshold powers, path costs, signed-motion count, and the local site census. The coupling is at most $4x$; the other growing factors have subpower growth. The expression therefore has a rate bound of exponent 2, and is eventually at most x^3.

In the expense inequalities take the truncation $m = M$, the separation count $b = P$, and

$$c' = e^{P/2}, \qquad r = e_x, \qquad 1 \le w \le e^{P/32}.$$

The fixed factor $c_0 = e^{(s-3)l}$ satisfies $4c_0 \le e^{P/32}$ eventually. Thus $4c'c_0 \le e^P$. Likewise (TS9) gives $xc_0e^2 \le e^{P/32}$, and hence

$$xc_0e^2\, w^2/c' \le e^{3P/32-P/2} \le e_x.$$

Writing $u = e^{-M}T_{\mathrm{lease}}$, we have

$$8e^P u \le 8x^3 e^{P-M} \le e_x.$$

The forced allowance is the specified rapidly decreasing function $\mathrm{smoke}(x) = \exp[-\operatorname{Wave}(1/5, x)]$. Its decay dominates e^P, giving

$$2e^P \operatorname{smoke}(x) \le e_x$$

eventually. Consequently the total expense t is at most $3e_x$. The site factor $(2B+1)^4$ has subpower growth, so $(2B+1)^4 t \le 1$ eventually. Finally $e^t - 1 \le 1$, which is the required output allowance here. Also $4u \le 1$ follows from $8e^P u \le e_x \le 1$. The wing tariff is the already imposed local surgery tariff. This proves all the displayed expense conditions simultaneously for every weight in the stated interval.

For each fixed stage, the preceding bounds impose only one eventual lower bound on x. The stagewise threshold selection enlarges the original coupling-entry delay to exceed these bounds at every permitted stage transition. Taking the maximum with the other construction thresholds preserves them all; later delay increases preserve this choice. Thus the physical calendar used in (TS6)–(TS8) satisfies the expense bounds before the finer cutoff or volume is chosen.

16.7.1 Finite density surgery and the equation insertion

The finite lens changes one bond by ordered group multiplication. Health and chord inequalities preserve the exact flag set, so exponent differences only involve supports containing the modified source. The resulting finite subset expansion is exact. Only the explicitly stated finite-shelf approximation is an error, and its allowance is selected after division by the node step has been included. Haar invariance gives the zero-integral identity for the same stock density used throughout the marker construction.

Why the marker is a finite representative The equation insertion is constructed from a local change of density before the continuum quotient is taken. Its flags and replicated marks retain where that change is supported, and its expense calendar gives the root bounds needed by spectators. Thus the eventual equation-null class has an actual finite local representative on the chosen Wilson family. The surgery identity and the weighted estimates below establish this assertion together; a formally vanishing coefficient symbol would not supply the same conclusion.

The lens itself has a finite, exact definition. For an oriented bond issuing from z, let $T_{\nu(j)}U$ change its link by left multiplication by $\nu(j)$; for a negatively read bond the equivalent stored positive link is changed by the corresponding right inverse multiplication. Let $S_Q(U)$ denote the scalar stock of a support Q. Truncate the represented exponent by support-cover number:

$$H_M(U) = v|\,\mathrm{Site}\,| + \sum_{\mathrm{cov}_B(Q)\le M} S_Q(U).$$

The lattice is finite, so this sum is finite. For the bounded local function $J(U, j)$, the lens is

$$L(U) = \int_G \left[J(T_{\nu(j)}U, j) e^{H_M(T_{\nu(j)}U) - H_M(U)} - J(U, j) \right] d\nu_G(j). \tag{TS10}$$

The admissibility conditions on these data are concrete: $J(\,\cdot\,, j)$ reads the fixed support P; $J(gU, g_z j) = J(U, j)$; a nonzero value requires health at z with threshold $\theta/4$; $\nu(aj) = a\nu(j)a^{-1}$; ν is measurable; and its chord is at most D, with $8D \le \theta$. These conditions are preserved under the signed chamber copies used above.

To justify expansion of (TS10), put $V = T_{\nu(j)}U$ and suppose $J(V, j) \neq 0$. Every link changes in chord distance by at most D. The chord of a plaquette changes by at most $4D$, by telescoping its four unitary factors. A plaquette whose value changes touches z, so its chord at V is at most $\theta/4$. Its chord at U is therefore at most $\theta/4 + 4D \le 3\theta/4$. Both are below θ. Unchanged plaquettes have identical threshold comparisons. Hence the flag sets at U, V agree exactly. The same reasoning after any signed motion shows that their relaxed kept-plaquette tests agree, since each relaxed threshold is at least θ.

A stock whose support omits z is unchanged: the modified bond has an endpoint at z, all internal links of that support agree, and its flag data agree. Thus

$$H_M(V) - H_M(U) = \sum_{\substack{z \in Q \\ \mathrm{cov}_B(Q) \le M}} [S_Q(V) - S_Q(U)].$$

For this finite set of supports, set

$$t_Q(U, j) = \begin{cases} e^{S_Q(V) - S_Q(U)} - 1, & J(V, j) \neq 0, \\ 0, & J(V, j) = 0. \end{cases}$$

The distributive law and exponential addition give the exact identity

$$J(V, j) e^{H_M(V) - H_M(U)} = \sum_{A \subseteq \{Q : z \in Q,\ \mathrm{cov}_B(Q) \le M\}} J(V, j) \prod_{Q \in A} t_Q(U, j).$$

If $J(V, j) = 0$, both sides are zero; this also explains the zero branch in the definition of t_Q. It avoids any need to estimate exponent differences on configurations where the insertion is absent. This finite subset expansion is the local surgery expansion priced in the expense inequalities.

Every term reads only the inspected union of P and the supports Q in that formula. The cover-extent inequality bounds its distance from z by

$$\mathrm{pierSize} = (l + M)(4B + 2) + 10B$$

when $\mathrm{cov}_B(P) \le l$ and $0, z \in P$. The resulting lens reads the origin-centered window of radius $2\,\mathrm{pierSize}$. Its gauge invariance follows from the displayed covariance of J, ν, invariance of the stocks, and the change of Haar variable $j \mapsto g_z j$. Measurability and boundedness follow from the finite stock expansion and its envelope. These facts provide the locality, gauge invariance and integrability used in the copied-lens and root estimates.

For the equation marker used above, specialize (TS10) as follows. On a positive source link e let $v = \mathrm{Ad}_g e_c$ be a unit color direction, let $\nu(g) = \exp(t_i h v)$, and take

$$J(U, g) = 1_{\{\text{healthy at the source with threshold } \theta/4\}} \langle Y(U)_e, v \rangle.$$

For a reversed link use the equivalent reversed multiplication in (TS10). The current is gauge covariant at the source, so the covariance of J and ν is exactly the one required there. The step

$h = \theta x^{-1/24}$, with its fixed node multiples, eventually satisfies the stronger chart condition $64D \le \theta$; thus it also satisfies the $8D \le \theta$ flag condition of (TS10).

Let $\mathscr{L}_{e,t_i h}(U)$ be this lens with the finite stock exponent, and sum over the active links, color basis indices and nodes:

$$\mathfrak{M}(U) = \frac{1}{h} \sum_{e,c} \sum_{i<m} w_i \, \mathscr{L}_{e,t_i h}(U). \tag{TS11}$$

Haar integration is already part of each lens. The coefficient w_i/h is its actual normalization. For every node its lens integral is zero for the stock density, by the Haar change of variable; the matched course transports this identity to the corresponding finite Wilson equation expression. Expanding by (TS10), its empty subset supplies the main density variation and its nonempty subsets are exactly the secondary terms F in (LCE6). Taylor expansion with the node moments above gives the current divergence and action derivative, with the remainder E. The exact healthy-region identity is therefore

$$\mathfrak{M} = \mathcal{Z} - \beta \sum_{a} \mathscr{D}_a \widetilde{J}_{ab} + (\mathrm{div}_\nu \, Y_b)(1), \tag{TS12}$$

up to the explicitly retained finite-shelf approximation, uniformly at most $e^{-\mathrm{prism}(x)}$ after division by h. The shelves are exhausted only after choosing this error allowance, so division by h is paid in their selection.

The normalization and support costs in the rooted estimate can be read directly from (TS11):

$$\sum_{e,c,i} |w_i|/h \le \mathrm{card}_B(x) \Big(\sum_{i<m} |w_i| \Big)/h \le x^2$$

eventually at each fixed m. The active-link count and fixed node sum have rate zero, while h^{-1} has rate $1/2 + 1/24$, strictly below two. Each lens is supported in the window of radius $2\{(l + q)(4B + 2) + 10B\}$ obtained above. At the actual middle stop its copied root bound is at most $\mathrm{T}_i(-1900000)$; the preceding total coefficient expense is at most $\mathrm{T}_i(100000)$. The rooted triangle inequality and $\mathrm{T}_i(a)\mathrm{T}_i(b) = \mathrm{T}_i(a + b)$ give

$$\|\mathfrak{M}\|_{\mathrm{root}} \le \mathrm{T}_i(-1800000). \tag{TS13}$$

Here $\mathrm{T}_i(a) = \exp\{a(\log \widehat{\beta}_i)^{48}/10000\}$; all finite local radii and prices are absorbed before choosing the sufficiently late stop. The rescaling from that stop to the reporting scale has degree-five factor at most $\mathrm{T}_i(100000)$. The healthy error is exponentially smaller than any fixed negative T_i power, and the global polynomial bound and exceptional-event estimate pay its complement. Signed averaging removes the constant in (TS12) and preserves the root bound. This is the normalized equation marker used in (TS1)–(TS5), including the costs of the full joint replicated array; it is not a formal integration-by-parts symbol without a local representative.

17 Completion of the proof

The common witnesses assembled here establish the public theorem's quantifiers. All degree horizons, derivatives, label coordinates, operator products and trace comparisons are consumed on one Wilson family and one composite refinement. The inner mesh limit remains full at every suitable period; only the outer volume limit is subsequential. The detailed construction supplies additional stress and contact information on these common witnesses without changing the list of clauses in the public theorem.

We assemble the choices to verify the quantifiers of the theorem. Fix the given group and representation. Choose blocking penalty 1. The linear kernel construction gives its positive widths and stagewise lower bounds on block ratios. The nonlinear geometry, Gaussian and connected-activity estimates give further stagewise bounds, with block ratio chosen first, then the remaining fixed parameters, and finally the coupling-entry delay. The coefficient, detector, approximation and trace estimates impose bounds in this same order. At each stage only finitely many demands of its prescribed finite degree and derivative order occur. Their maximum is a finite demand. Increase the stage parameters to meet it. Increasing later block lower bounds and delays preserves the earlier requirements. Taking the increasing degree and derivative orders in this construction meets every fixed finite request eventually.

The exact one-step identities and the banked-history convergence then construct the calibrated Wilson family, including its common insertion tokens, before any volume subsequence is specified. Its represented action at the finest cutoff is the Wilson action in the theorem. Positivity of its partition function and the local insertion construction give the stated probability laws, continuity, gauge invariance and fixed finite stencil locality. The finite arcs (FCA7), with their exact Wilson startup and density identities, provide these calibrated histories. The summed drift argument (RUN1)–(RUN5), using the realized calendar (CAL1)–(CAL11) and the actual update (DR7), gives the positive logarithmic coupling slope uniformly in cutoff and period; its finest-cutoff specialization is the bare-coupling statement of the theorem.

Now fix the requested strictly increasing sequence of periods. Retain the coupling limits and perform the common hierarchy and coefficient extraction. Compose its successive refinement maps. The fixed-volume Cauchy argument and the smearing passage give the full mesh limit at every suitable fixed period; the refined period sequence gives its outer limit. Thus the auxiliary compact selections do not weaken the inner limit in the statement. The field quotient and the physical detector identify one finite-dimensional field filtration on this table.

The compression algebra constructs its compact metric phase space with full-support probability measure. Its character hierarchies, including the selected one, have the stated smooth correlations and barycenter identity. Apply the heat, tube, boundary and Gram constructions to these same hierarchies. They give the Hilbert fields on the common cyclic domain, the adjoint and locality identities, the strongly continuous spacetime action and its forward-cone momentum measure. The heat-ground argument gives the unique vacuum in every character. The label-separation argument gives field faithfulness and the nonzero unit field. The strict imaginary-time and boundary identities identify the reconstructed moments with the corresponding Euclidean hierarchy.

Choose the equivariant operator-product reader on this field filtration once. The separated expansion, its coefficient bounds and its tested remainder hold for every character and for the unselected hierarchy. The nested replacement and integration arguments retain this reader, and the collision-extension construction retains its basis and finite projections. Their outputs therefore refer to the same fields and correlations, including when several truncation degrees or collision patterns are requested together.

Finally, the Gaussian comparison and position normalization determine the physical stress tensor on that table. The finite equation-marker bounds and weighted near-kernel estimates identify its centered trace and fix its amplitude. The spectral and position-space calculation proves

$$\lim_{r\downarrow 0} r^{8}\log^{2}(1/r)\,S_{2}^{\chi}(\Theta(-r/2,\mathbf{0}),\Theta(r/2,\mathbf{0})) = \frac{3\dim_{\mathbb{R}}\mathfrak{g}}{\pi^{4}} > 0$$

in the selected phase, with the same scalar diagnostic under the phase identifications established above. The mass-shell asymptotic argument excludes the finite comparison class specified in the

theorem; the coherent-state argument gives the additional Maxwell countertests described in the trace chapter.

Every witness in these conclusions has been constructed on the one chosen Wilson family and retained on the one composite refinement. Passing to the public fields, operators, correlations and channel coefficients is therefore their identification on that common table. This proves the theorem.

18 Conclusion

The construction connects a finite Wilson probability law to a common continuum hierarchy through explicit kernels, nonlinear integrations, calibrated histories and physical coefficient spaces. The full mesh limit at each suitable fixed period precedes the outer volume refinement. Reflection positivity, the phase algebra and the tested boundary construction then give the fields on their common cyclic domain. The operator-product and contact constructions retain their finite horizons and chosen devices on that same hierarchy.

The stress calculation returns to actual finite local functions, controls their equation residuals and fixes the centered trace amplitude. Its logarithmic short-distance asymptotic yields the finite mass-shell and Maxwell comparison conclusions with the hypotheses stated above. No positive mass gap is asserted.

This paper provides the full natural-language mathematical exposition of the expanded source. The flagship in the same Zenodo release [2] develops Rodgers's source–readout realization and recovery of this construction. Both papers retain the definitions and derivations required for their own arguments.

A Editorial amendments and conventions

The mathematical source is the expanded `Full_pre_source_translation.md`, included in the release's `sources/` directory. Its 15,788 lines are counted with universal newlines. The coverage record locates every source heading in this paper. The following records the mathematical amendments and interpreted typography.

A1: principal predictor (source line 1017). The extracted expression is $m(q_0)^{-1}T_a/g_a - H/U$; the reconciled expression is $m(q_0)^{-1}(T_a/g_a - H/U)$. The direct subtraction in Section 5.22 shows $E_0 - r_0 = SG_{0,a}$ in the declared multiplier convention. The singular factor S^{-1} is cancelled before Fourier extraction, and the nonprincipal and dual branches retain their displayed definitions. The amended parenthesization implements the stated cancellation.

A2: Gaussian prefactor (source line 1317). The extracted exponent $-rh + d_v$ is replaced by $-(rh + d_v)$. The active, dormant and frozen-link counts, and the substitution $m_g = c(\sqrt{2\pi}g)^r$, are calculated in Section 6.2. The dormant Gaussian integrates to one, and its normalization remains present. The determinant factor is unchanged. This calculation fixes the normalization used by the chart integral.

Interpreted typography (E2). At source line 1575 use $C_K = (n_a + n_b)AJ$: the common row and column count multiplies their common estimate. At line 3218 use $4\,\mathrm{pcap}(4)(\sqrt{\mathrm{direct}_{\mathrm{top}}} + \mathrm{paths}_{\mathrm{top}})t$; the common factor t comes from both base-row estimates. At line 3428 use $HC(QR_{\mathrm{price}})^n$, consistent with the repeated bound at source line 4020. The source's damaged comma in (CAL21) and split `rm free` in (GL0) are restored from their displayed mathematical context.

Domains and interpretation. The combined startup allowances in Section 8.14 hold on $x/2 \leq \beta \leq 2x$. The forced estimate alone uses only the lower bound. The Gaussian comparison distinguishes C_0 and C_V, with their determinant and inverse factors (Sections 8.23–8.25). The primitive relation retains ordered group multiplication (Section 7.22). Exact radius matching is geometric, while moment comparisons keep their error bounds (Section 16.2). The full idempotent-kernel proof is included in Section 11.5. These conventions also govern the corresponding recovery maps in the source–readout paper; the constructions used in this exposition are defined and derived locally.

Scope of the results. The group and representation are given; no connectedness hypothesis is added. One Wilson family precedes the requested volume sequence. Its inner fixed-period mesh limit is full, and its outer volume limit is subsequential. Field assertions use the common cyclic invariant domain. Every OPE/contact statement retains its finite horizon, geometry and chosen subtraction device. The full-distance trace argument retains the same table and phase; the moving-mass estimates retain their positive-mass choice order; the real fourth-order remainder is distinguished from a formal jet and from leading transfer to weaker frames. The two comparison exclusions retain their specified hypotheses. No positive mass gap or Millennium-problem solution is asserted by this presentation.

Transcription typography in the companion. Short formulae written directly in source prose are set in mathematical mode; subscripts, superscripts, inequalities and function names are joined into coherent expressions. The Unicode aggregate signs Σ and indexed Π are printed as $\sum$ and $\prod$ where their context is a sum or product. The standalone heat-ground projection remains Π. These typography choices introduce no new operator or estimate. The following local interpretations make explicit the source's intended application or grouping:

Source line(s)	Typeset expression	Reason from the local argument
642	$f_\xi(p) - f_\xi(y)$	The immediately preceding sentence applies Cauchy's estimate to the function f_ξ at p and y.
815	$(d_2 B_1)^{\mathsf{T}} F (d_2 B_1)$	The transpose acts on the whole product: $V = B_1 M B_1$, $M = d_2^{\mathsf{T}} F d_2$, and B_1 is symmetric. This clarifies the extraction's misplaced transpose.
989	$L^4 (L^{-3})^2$	The Gram row squares the L^{-3} coefficient bound and sums over L^4 residues.
1431	$[(A\|_{W_j})^{-1}]_{ij}$	This is the ij entry of the inverse of the principal restriction, zero-padded outside W_j.
1699	$H_{p_0} = 1$	The distinguished patch point is p_0; the value of its frame matrix is fixed to the identity.
3121	$m_g(\beta) = c(\sqrt{2\pi}\,\beta^{-1/2})^r$	The damaged source exponent $\beta^{-1}/2$ is the Gaussian scale $g = \beta^{-1/2}$ already used in the active/dormant normalization.

4620, 4627	$\mathrm{logRead}_i(1)$	The supplied name is an operator label; its receipt and reset identities define its use locally.
5195	$\prod_{i\in p\setminus \mathrm{keys}(M)} t_i$	The text explicitly removes the marked team's keys before multiplying the remaining primitive-body tariffs.
9736	$(a_{p(k),m})_m$	This is a point in the countable product indexed by the retained level m, not a second subscript on one scalar.
9738	$(B_m - \widehat{\beta}_m)/t_m$	The preceding inequality bounds the normalized difference between the actual limiting coupling and its reference clock.
10044, 10050	Π	The standalone symbol denotes the heat-ground projection; it is distinct from indexed product notation.

In the reconstruction chapter, ASCII pairing brackets and postfix stars are set as Hilbert pairings $\langle \cdot, \cdot \rangle$ and adjoints on their declared domains. The zero-energy event is written $\{E = 0\}$ and the low-energy function as $\rho(E)$. These are the same pairings, spectral event and scalar function defined in the local proofs. Local internal source aliases identify provenance rather than supplying additional hypotheses or replacing a derivation.

B Source, attribution and reproducibility

The underlying Yang–Mills Lean work is credited to OpenAI. The source actually presented here is the expanded *Wilson Yang–Mills Construction: Mathematical Translation* [1], preserved in the release. Jeremy Rodgers is the author of this exposition, prepared with AI assistance. The related source–readout reconstruction develops the programme that Rodgers originated and directed [2].

The complete editable manuscript consists of the root LaTeX document, `sections/*.tex`, and `references.tex`. All mathematical content is contained in those files. The release includes the expanded source, coverage and amendment records, build instructions, and verification records. Source-unit comments connect the editable text to the original heading intervals. The verification record distinguishes the current build and focused calculations from retained numerical and structural checks; no new Lean/Lake execution is reported.

References

[1] *Wilson Yang–Mills Construction: Mathematical Translation.* Project source document: expanded natural-language extraction prepared for Jeremy Rodgers's Yang–Mills project with AI assistance and supplied by Jeremy Rodgers. Included unchanged as `sources/Full_pre_source_translation.md`. OpenAI is credited separately for the underlying Lean work; the supplied materials do not identify an exact upstream repository or release.

[2] J. Rodgers. *Shadow Theory and the Wilson Yang–Mills Construction: A Source–Readout Reconstruction.* Flagship paper in this shared Zenodo record. doi:10.5281/zenodo.24173681.

[3] J. Rodgers. *A Source-to-Readout Architecture for a Theory of Everything: Sector Physics as Coupled Projections from a Common Realization Carrier*, v1.0. Zenodo. DOI 10.5281/zenodo.21366204.

[4] K. G. Wilson. Confinement of quarks. *Phys. Rev. D* **10** (1974), 2445–2459. doi:10.1103/PhysRevD.10.2445.

[5] A. Haar. Der Massbegriff in der Theorie der kontinuierlichen Gruppen. *Ann. of Math.* (2) **34** (1933), 147–169. doi:10.2307/1968346.

[6] G. C. Wick. The evaluation of the collision matrix. *Phys. Rev.* **80** (1950), 268–272. doi:10.1103/PhysRev.80.268.

[7] I. Gelfand and M. Neumark. On the imbedding of normed rings into the ring of operators in Hilbert space. *Rec. Math. [Mat. Sbornik] N.S.* **12(54)**, no. 2 (1943), 197–217. `https://www.mathnet.ru/eng/sm6155`.

[8] I. E. Segal. Irreducible representations of operator algebras. *Bull. Amer. Math. Soc.* **53** (1947), 73–88. doi:10.1090/S0002-9904-1947-08742-5.

[9] M. H. Stone. The generalized Weierstrass approximation theorem. *Math. Mag.* **21**, no. 4 (1948), 167–184; continued, no. 5 (1948), 237–254. doi:10.2307/3029750; doi:10.2307/3029337.

[10] K. Osterwalder and R. Schrader. Axioms for Euclidean Green's functions. *Commun. Math. Phys.* **31** (1973), 83–112. doi:10.1007/BF01645738.

[11] K. Osterwalder and R. Schrader. Axioms for Euclidean Green's functions II. *Commun. Math. Phys.* **42** (1975), 281–305.

[12] K. Osterwalder and E. Seiler. Gauge field theories on a lattice. *Ann. Phys.* **110** (1978), 440–471.

[13] J. Fröhlich, R. Israel, E. H. Lieb and B. Simon. Phase transitions and reflection positivity. I. General theory and long range lattice models. *Commun. Math. Phys.* **62** (1978), 1–34.

[14] J. Glimm and A. Jaffe. *Quantum Physics: A Functional Integral Point of View*, 2nd ed. Springer, 1987.

[15] R. F. Streater and A. S. Wightman. *PCT, Spin and Statistics, and All That*. Benjamin, 1964.

[16] T. Bałaban. Renormalization group approach to lattice gauge field theories. I. Generation of effective actions in a small field approximation and a coupling constant renormalization in four dimensions. *Commun. Math. Phys.* **109** (1987), 249–301.

[17] D. C. Brydges. A short course on cluster expansions. In K. Osterwalder and R. Stora (eds.), *Critical Phenomena, Random Systems, Gauge Theories*, Part I, Les Houches Session XLIII (1984), North-Holland, Amsterdam, 1986, 129–183. `https://www.math.ubc.ca/~db5d/Seminars/les_houches_84.pdf`.

[18] E. Seiler. *Gauge Theories as a Problem of Constructive Quantum Field Theory and Statistical Mechanics*. Lecture Notes in Physics 159, Springer, 1982.

[19] D. J. Gross and F. Wilczek. Ultraviolet behavior of non-abelian gauge theories. *Phys. Rev. Lett.* **30** (1973), 1343–1346.

[20] H. D. Politzer. Reliable perturbative results for strong interactions? *Phys. Rev. Lett.* **30** (1973), 1346–1349.

[21] A. Jaffe and E. Witten. Quantum Yang–Mills theory. In J. Carlson, A. Jaffe and A. Wiles (eds.), *The Millennium Prize Problems*, Clay Mathematics Institute and American Mathematical Society, 2006, 129–152. `https://www.claymath.org/wp-content/uploads/2022/06/yangmills.pdf`.