

Geometric Realization of Completed Source Relations

Descent, Orbit Spaces, Invariant Relations, and
Variational Response in Shadow Theory

Jeremy Rodgers*
Independent Researcher

June 2026

Abstract

The minimal completion constructed in Paper 3 is an abstract quotient determined by a declared readout together with a set-indexed family of invariant source relations. This paper identifies the conditions under which such a completion possesses a genuine realization in differential geometry and classical field theory. The realization domain is restricted to *compatible* local metric, bundle, connection, matter, and boundary data: Čech cocycle and overlap equations are used to construct global fields, and uniqueness is asserted only up to bundle isomorphism, never up to a choice of coordinate chart or gauge representative. The physical geometric source is then defined as the orbit space of admissible global configurations under the groupoid of allowed principal-bundle isomorphisms covering allowed diffeomorphisms.

Gauge- and diffeomorphism-invariant global relations—holonomy conjugacy classes, Wilson observables, characteristic numbers, boundary and asymptotic data, and spectra of geometrically defined operators with specified domains—descend to functions on this orbit space. Their joint image with a declared Tier-1 readout realizes, on the restricted domain, the canonical minimal completion of Paper 3. We prove a restricted geometric realization theorem, together with an orbit reconstruction criterion characterizing when the completion map is injective. For flat $U(1)$ connections on an oriented circle we prove that holonomy classifies gauge orbits completely, yielding the exact realization $E_{\text{hol}} \cong U(1)$, and we exhibit a strictly coarser spectral completion $E_{\text{spec}} \cong [-1, 1]$ arising from a charged scalar relation.

The final sections embed completed relation data into variational dynamics. A gauge- and diffeomorphism-invariant relation-dependent action yields correctly normalized stress-energy, gauge-current, and matter responses; the Einstein, Yang–Mills, and matter equations follow by functional variation, and the Ward–Noether identities associated with diffeomorphism and gauge invariance enforce covariant conservation. Completed source relations therefore enter Tier-1 field equations through invariant geometric structure and derived response, not through a choice of coordinates or gauge. No claim is made that every abstract completion is geometrically realizable; the realization boundary is precisely the content of the results below.

Contents

1	Introduction and scientific position	2
1.1	Position in the sequence	2
1.2	Progression of the paper	3
1.3	Scope	3

*Website: everythingequation.com

2	The geometric source domain	4
2.1	Configurations	4
2.2	The redundancy groupoid	5
3	Local data and global descent	7
3.1	Local data and compatibility equations	7
3.2	The descent theorem	8
4	Curvature, holonomy, and invariant global relations	10
4.1	Curvature and the Bianchi identity	10
4.2	Invariant global relations	11
4.3	Holonomy and Wilson observables	12
4.4	Loops and diffeomorphism invariance	12
5	Geometric realization of the minimal completion	13
5.1	The geometric evaluation map	13
5.2	The restricted realization theorem	14
5.3	When does the completion reconstruct the source?	15
6	Classification of flat $U(1)$ connections on an oriented circle	15
6.1	The model and its conventions	15
6.2	The classification theorem	17
6.3	The realized completions	18
7	Variational realization of completed relations	21
7.1	Variational setting	21
7.2	The functional chain rule for relation functionals	22
7.3	Matter pairings and independent components	23
7.4	Response tensors, currents, and Euler operators	24
7.5	The coupled field equations	25
8	Ward–Noether identities and covariant consistency	26
8.1	Route (a): on a full solution, via the contracted Bianchi identity	27
8.2	Route (b): diffeomorphism invariance, independent of the Einstein equation	27
8.3	The gauge Ward identity	29
9	From completion to Tier-1 representation	30
10	Conclusion	31

1 Introduction and scientific position

1.1 Position in the sequence

Paper 1 [16] established source–readout non-equivalence and identified the descent obstructions that prevent a readout from being inverted. Paper 2 [17] characterized when the information lost by a readout obstructs a nominated target, and what any target-sufficient extension is obliged to distinguish. Paper 3 [18] constructed, for a set-indexed family $\{R_\alpha\}_{\alpha \in A}$ of invariant source relations and a declared readout $p: S \rightarrow T$, the canonical minimal completion

$$E_A = \text{Im}(p, R_A) \subseteq T \times \prod_{\alpha \in A} V_\alpha, \quad R_A = (R_\alpha)_{\alpha \in A}, \quad (1)$$

together with its quotient realization $S \rightarrow E_A$ and its terminal universal property among surjective extensions through which p and every R_α factor.

Construction (1) is purely set-theoretic. It is indifferent to whether the relation values $(t, (v_\alpha))$ can be produced by any geometry at all. Paper 4 supplies exactly the missing ingredient:

When do the abstract relation coordinates of Paper 3 arise from compatible global metrics, bundles, connections, matter fields, boundary data, and gauge-invariant observables?

The answer has the structure of a *realization boundary*. Not every point of the Cartesian product $T \times \prod_\alpha V_\alpha$ is attained; the attainable points are precisely the image of a geometric evaluation map, and that image is constrained by descent (Čech cocycle and overlap conditions), by invariance (the relations must be constant on orbits of the allowed symmetry groupoid), and by the admissibility class chosen for the model (regularity, signature, causal structure, boundary and asymptotic conditions). The abstract completion is then realized *relative to* the invariant geometric relations declared, and only on that restricted domain.

1.2 Progression of the paper

The paper follows a single line of development. Compatible local geometric data glue to global fields (Section 3, Theorem 3.4). Physical configurations are reduced by allowed bundle isomorphisms and diffeomorphisms (Section 2, Definition 2.9). Invariant global geometric relations become well-defined functions on the reduced orbit space (Section 4). Their joint image with a declared readout realizes the canonical completion of Paper 3 (Section 5, Theorem 5.3). Flat $U(1)$ connections on an oriented circle provide an exact classification example (Section 6, Theorem 6.5). Invariant relation-dependent actions generate stress-energy, currents, and matter response through functional variation (Section 7, Theorem 7.8). Gauge and diffeomorphism invariance impose the corresponding Ward–Noether identities (Section 8, Theorem 8.4 and Proposition 8.5).

1.3 Scope

The scope restrictions recorded here are load-bearing and are used explicitly in the statements of the theorems below. It is not claimed that every abstract completion of Paper 3 is geometrically realizable; realizability is a nontrivial condition, and Corollary 3.7 exhibits the obstruction. It is not claimed that any finite list of curvature scalars separates arbitrary Lorentzian geometries; separation is a hypothesis to be established on the restricted domain of a given model, and Remark 4.6 records why no general result of that kind is available. Gauge potentials and coordinate representatives are never treated as physical states: all realization relations are functions on the orbit space S_{geo} . A loop observable is not treated as diffeomorphism invariant unless the loop itself is included in the configuration data and transformed with it, or the observable is formed by an invariant construction over loops, or the declared boundary data pins the relevant homology class; Section 4.4 makes these alternatives precise and requires a model to declare which it uses. Global relation coordinates are not treated as independent dynamical fields unless a model explicitly promotes them to fields. No new physical interaction or empirical claim is introduced, and no claim is made about the physical status of our universe: the relation-dependent action of Section 7 is a template for admissible couplings, and its consequences are theorems about that template, not predictions. Finally, the projected-dynamics and memory machinery of Paper 5, and the Randall–Sundrum construction of Paper 7, are not used and are not assumed; where the present results terminate in an open dynamical question, this is stated as such (Section 9).

2 The geometric source domain

2.1 Configurations

Convention 2.1 (Smoothness, signature, and index conventions). Throughout the general theory, M is a smooth, connected, paracompact, oriented four-manifold, possibly with boundary ∂M . The choice $\dim M = 4$ is made for definiteness in the physical application and is used nowhere essentially: Sections 2–5 hold verbatim in any dimension. Section 6 is a deliberate one-dimensional specialization, declared as such in Convention 6.1, and the dimensional stipulation of the present convention is suspended there. All fields are smooth unless a weaker regularity class is explicitly declared as part of the admissibility data. A Lorentzian metric has signature $(-, +, +, +)$; a *time-oriented* Lorentzian metric is a pair $(g, \mathfrak{t})$ consisting of a Lorentzian metric g together with a choice $\mathfrak{t}$ of one of the two continuous global orientations of the timelike cone bundle of g (equivalently, the choice of a connected component of the set of g -timelike vector fields, when one exists). Not every Lorentzian metric admits such a choice; admitting one is a nontrivial global condition on g , and it is imposed below as part of admissibility rather than derived (see Remark 2.3). Greek indices run over $0, 1, 2, 3$ and are raised and lowered with g . G denotes a finite-dimensional Lie group with Lie algebra $\mathfrak{g}$, equipped with a fixed Ad-invariant nondegenerate bilinear form $\langle \cdot, \cdot \rangle$ on $\mathfrak{g}$ (for compact semisimple G one may take the negative Killing form; for $G = U(1)$ the standard normalization on $\mathfrak{u}(1) \cong i\mathbb{R}$). Finally $\rho: G \rightarrow \mathrm{GL}(W)$ is a fixed finite-dimensional representation; ρ is *not* assumed unitary, and where unitarity is needed it is hypothesized explicitly.

Definition 2.2 (Global geometric configuration). A *global geometric configuration* is a tuple

$$\Psi = (M, g, P, A, \Phi, \mathcal{B}), \quad (2)$$

in which

- (a) g is a Lorentzian metric on M ;
- (b) $\pi: P \rightarrow M$ is a principal G -bundle;
- (c) $A \in \Omega^1(P, \mathfrak{g})$ is a principal connection on P [6, 3];
- (d) Φ is a section of the associated bundle $P \times_{\rho} W \rightarrow M$;
- (e) $\mathcal{B}$ denotes the boundary, asymptotic, causal, and global data required to make the model well posed. Among these data we *always* include a choice of time orientation $\mathfrak{t}$ for g ; the remaining components of $\mathcal{B}$ are model-specific and may comprise, for instance, fixed induced data on ∂M , a decay class at infinity, a fixed conformal boundary structure, a global hyperbolicity requirement, a spin structure, or a fixed topological sector.

Remark 2.3 (Where time orientation lives). Time orientation is a property of the pair (M, g) , not of the bare smooth manifold M : it is meaningless to speak of M as “time-orientable” independently of a Lorentzian metric, since the timelike cone bundle whose components are being selected is defined by g . Accordingly the choice $\mathfrak{t}$ is carried in $\mathcal{B}$, and the arrows of the redundancy groupoid are required below to preserve it. This has a consequence for the descent theorem, recorded explicitly in Theorem 3.4 and Remark 3.6: descent reconstructs the *metric*, but membership of the reconstructed $(g, \mathfrak{t})$ in a time-oriented admissibility class is an additional *global* condition on the glued object, which either must be imposed after the fact or must be built into the descent data as compatible local time orientations.

Definition 2.4 (Admissibility class). An *admissibility class* $\mathcal{C}_{\mathrm{geo}}$ is a specified class of tuples (2) satisfying declared conditions on regularity, signature, causal structure, boundary and asymptotic behaviour, and topological sector, and in which every g is time-orientable with a distinguished time orientation $\mathfrak{t} \in \mathcal{B}$. The class $\mathcal{C}_{\mathrm{geo}}$ is part of the data of the model. All subsequent statements are relative to a fixed $\mathcal{C}_{\mathrm{geo}}$.

The insistence that $\mathcal{C}_{\text{geo}}$ be declared is not pedantry. Every separation, realization, and conservation statement below is sensitive to it: a relation family may separate orbits in one admissibility class and fail to separate them in a larger one, and a Ward identity may hold with compactly supported variations and fail without a boundary term.

2.2 The redundancy groupoid

Gauge and coordinate freedom in a bundle-theoretic setting is not, in general, a group acting on a fixed configuration space: different admissible configurations may carry non-isomorphic bundles, and the diffeomorphisms that lift to a given bundle need not lift to another. The correct structure is a groupoid.

Definition 2.5 (Ambient allowed automorphisms). Fix a principal G -bundle $P \rightarrow M$. Let $\text{Diff}_{\mathcal{B}}(M)$ denote the group of diffeomorphisms of M preserving the orientation of M and the non-metric structural data in $\mathcal{B}$ (boundary, asymptotic regime, topological sector), and let

$$\text{Aut}_{\mathcal{B}}(P) := \{\hat{F} \in \text{Aut}(P) : \hat{F} \text{ covers some } f \in \text{Diff}_{\mathcal{B}}(M)\} \quad (3)$$

be the group of G -equivariant diffeomorphisms of P covering such diffeomorphisms. Define

$$\text{Diff}_{P,\mathcal{B}}(M) := \text{image of } \text{Aut}_{\mathcal{B}}(P) \text{ in } \text{Diff}_{\mathcal{B}}(M), \quad (4)$$

that is, the subgroup of allowed base diffeomorphisms that admit an equivariant lift to P .

Proposition 2.6 (Structure of the ambient group). *With the definitions (3)–(4) there is a short exact sequence of groups*

$$1 \longrightarrow \text{Gau}(P) \longrightarrow \text{Aut}_{\mathcal{B}}(P) \longrightarrow \text{Diff}_{P,\mathcal{B}}(M) \longrightarrow 1, \quad (5)$$

where

$$\text{Gau}(P) = \{\hat{F} \in \text{Aut}(P) : \hat{F} \text{ covers } \text{id}_M\} \cong \Gamma(P \times_{\text{conj}} G)$$

is the gauge group. The sequence is exact by construction: $\text{Diff}_{P,\mathcal{B}}(M)$ is defined as the image of the projection, and the kernel of that projection is exactly the set of automorphisms covering the identity. The sequence need not split.

Proof. Exactness at $\text{Gau}(P)$ and at $\text{Diff}_{P,\mathcal{B}}(M)$ is immediate from the definitions of kernel and image. That the extension need not split is a statement about bundle topology: a splitting would furnish, for each allowed $f \in \text{Diff}_{P,\mathcal{B}}(M)$, an equivariant lift $\sigma(f) \in \text{Aut}_{\mathcal{B}}(P)$ depending on f homomorphically, and there is no reason for such a coherent choice to exist. Consequently $\text{Aut}_{\mathcal{B}}(P)$ is not in general a semidirect product $\text{Gau}(P) \rtimes \text{Diff}_{P,\mathcal{B}}(M)$, and nothing below assumes it is. $\square$

Remark 2.7 (The ambient group is not the isotropy group). Proposition 2.6 describes the *ambient* group of allowed bundle automorphisms over a fixed P . It does *not* describe the isotropy group of the groupoid at a configuration. The isotropy (stabilizer) group at $\Psi = (M, g, P, A, \Phi, \mathcal{B})$ is the subgroup

$$\text{Aut}(\Psi) := \{\hat{F} \in \text{Aut}_{\mathcal{B}}(P) \text{ covering some } f : f_*g = g, \quad f_*\mathfrak{t} = \mathfrak{t}, \quad (\hat{F}^{-1})^*A = A, \quad \hat{F}_\rho \circ \Phi = \Phi \circ f, \quad f_*\mathcal{B} = \mathcal{B}\} \leq \text{Aut}_{\mathcal{B}}(P), \quad (6)$$

that is, the stabilizer of the *whole tuple*: it preserves the metric, the chosen time orientation and orientation, the connection, the matter section, and the boundary data. $\text{Aut}(\Psi)$ is typically far smaller than $\text{Aut}_{\mathcal{B}}(P)$: for a generic (g, A, Φ) it is a small or even trivial group, its base part being contained in the isometry group of g . In particular the exact sequence (5) is a statement about $\text{Aut}_{\mathcal{B}}(P)$ and not about $\text{Aut}(\Psi)$.

Definition 2.8 (Redundancy groupoid). Let $\mathfrak{G}_{\text{geo}}$ be the groupoid whose objects are the elements of $\mathcal{C}_{\text{geo}}$. Given two objects

$$\Psi = (M, g, P, A, \Phi, \mathcal{B}), \quad \Psi' = (M, g', P', A', \Phi', \mathcal{B}'),$$

the arrows $F: \Psi \rightarrow \Psi'$ are the pairs $F = (\hat{F}, f)$ consisting of

- an allowed diffeomorphism $f \in \text{Diff}_{\mathcal{B}}(M)$, and
- a G -equivariant bundle isomorphism $\hat{F}: P \rightarrow P'$ covering f ,

subject to

$$f_*g = g', \quad f_*\mathbf{t} = \mathbf{t}', \quad (\hat{F}^{-1})^*A = A', \quad \hat{F}_\rho \circ \Phi = \Phi' \circ f, \quad f_*\mathcal{B} = \mathcal{B}', \quad (7)$$

where $\hat{F}_\rho$ is the induced isomorphism $P \times_\rho W \rightarrow P' \times_\rho W$. The second condition is the requirement that arrows preserve the chosen time orientation; it is not implied by the first, since $f_*g = g'$ alone permits f to exchange the two components of the timelike cone bundle. Composition is composition of bundle maps; identities are $(\text{id}_P, \text{id}_M)$. The isotropy group of $\mathfrak{G}_{\text{geo}}$ at Ψ is $\text{Aut}(\Psi)$ of (6).

The groupoid formulation—rather than a single group acting on a single configuration space—is forced by two features. First, the arrows relate objects carrying *different* bundle representatives, which is exactly what is needed when descent produces a bundle only up to isomorphism (Theorem 3.4). Second, isotropy jumps from configuration to configuration, since $\text{Aut}(\Psi)$ depends on Ψ ; there is no ambient group acting with constant stabilizer.

Definition 2.9 (Physical equivalence and the reduced source). Write

$$\Psi \simeq_{\text{geo}} \Psi' \iff \exists F: \Psi \rightarrow \Psi' \text{ in } \mathfrak{G}_{\text{geo}}. \quad (8)$$

This is an equivalence relation on $\mathcal{C}_{\text{geo}}$ because $\mathfrak{G}_{\text{geo}}$ is a groupoid. The *reduced geometric source* is the orbit space

$$S_{\text{geo}} := \mathcal{C}_{\text{geo}} / \simeq_{\text{geo}}, \quad (9)$$

with quotient map $q: \mathcal{C}_{\text{geo}} \rightarrow S_{\text{geo}}$, $q(\Psi) = [\Psi]$.

Remark 2.10 (Set-theoretic hygiene). S_{geo} is used below only as a set (equivalently, as the set of isomorphism classes of objects of the essentially small groupoid $\mathfrak{G}_{\text{geo}}$ obtained after fixing a set of bundle representatives). No smooth or stratified structure on S_{geo} is used or assumed; the orbit space of a gauge groupoid is in general not a manifold—precisely because isotropy jumps, as Remark 2.7 notes—and the arguments of Sections 5–6 do not require it to be. Where a topology is needed, as for the moduli space $\mathcal{M}_{\text{flat}}(S^1, U(1))$ in Section 6, it is supplied and identified explicitly.

Convention 2.11 (Status of relation coordinates). A relation coordinate r_α is a *function on* S_{geo} , not an independent field on M . It is not varied independently in Section 7 unless a model explicitly promotes it to a field by adjoining it to the configuration tuple (2) and enlarging $\mathcal{C}_{\text{geo}}$ accordingly. Where r_α is a functional of (g, A, Φ) its variation is computed by the functional chain rule of Section 7.2; where it labels a fixed topological or boundary sector it is held fixed under variations *within* that sector. Both cases are handled in Section 7; they are not conflated.

Every realization relation used in this paper is, by construction, a function on S_{geo} and never a function of an arbitrary coordinate chart or gauge potential.

3 Local data and global descent

We now establish the first leg of the progression: compatible local data glue to global fields. This is the theorem that fixes the *domain* of geometric realization. The material is classical [6, 11, 3]; what is new is only the use to which it is put and the care with which the invariance and admissibility hypotheses are tracked.

3.1 Local data and compatibility equations

Let $\{U_i\}_{i \in I}$ be an open cover of M . Suppose we are given, for each $i \in I$,

$$g_i \in \Gamma(S^2 T^* U_i), \quad A_i \in \Omega^1(U_i, \mathfrak{g}), \quad \Phi_i \in C^\infty(U_i, W), \quad (10)$$

where each g_i is an *intrinsic* smooth symmetric covariant 2-tensor field on the open set $U_i \subseteq M$, of Lorentzian signature at every point; together with smooth transition functions

$$h_{ij}: U_i \cap U_j \longrightarrow G. \quad (11)$$

If in addition each U_i is the domain of a chart $\psi_i: U_i \rightarrow \psi_i(U_i) \subseteq \mathbb{R}^4$, we may pass to *coordinate components*. These live on the coordinate images, not on M : setting

$$(\check{g}_i)_{ab} := ((\psi_i^{-1})^* g_i)_{ab} \quad \text{on} \quad \psi_i(U_i) \subseteq \mathbb{R}^4, \quad (12)$$

and writing $\psi_{ij} := \psi_i \circ \psi_j^{-1}$ for the coordinate transition map on $\psi_j(U_i \cap U_j)$, the components of a *single* intrinsic tensor in two charts are related by the usual pullback law

$$\check{g}_j = \psi_{ij}^* \check{g}_i \quad \text{on} \quad \psi_j(U_i \cap U_j), \quad \text{i.e.} \quad (\check{g}_j)_{cd} = \frac{\partial \psi_{ij}^a}{\partial x^c} \frac{\partial \psi_{ij}^b}{\partial x^d} (\check{g}_i)_{ab} \circ \psi_{ij}. \quad (13)$$

Equation (13) is a statement about *components on coordinate images*. It is distinct from the gluing condition (15) below, which is a relation among the intrinsic tensors g_i on their overlaps.

Definition 3.1 (Compatible local geometric data). The collection $\{g_i, h_{ij}, A_i, \Phi_i\}$ is *compatible* if, on every nonempty overlap (respectively triple overlap):

$$h_{ii} = 1, \quad h_{ji} = h_{ij}^{-1}, \quad h_{ij} h_{jk} = h_{ik}, \quad (14)$$

$$g_i|_{U_i \cap U_j} = g_j|_{U_i \cap U_j}, \quad (15)$$

$$A_j = \text{Ad}_{h_{ij}^{-1}} A_i + h_{ij}^* \theta_G, \quad (16)$$

$$\Phi_j = \rho(h_{ij}^{-1}) \Phi_i. \quad (17)$$

Here (14) is the Čech cocycle condition; (15) is the gluing condition (agreement of restrictions) for a section of the sheaf $S^2 T^* M$; and (16), (17) are the compatibility conditions for local connection forms and local matter representatives, θ_G denoting the left Maurer–Cartan form on G . For a matrix group $G \subseteq \text{GL}(n)$, equation (16) takes the familiar form

$$A_j = h_{ij}^{-1} A_i h_{ij} + h_{ij}^{-1} dh_{ij}.$$

Remark 3.2 (Čech H^1 for nonabelian G). A cocycle $\{h_{ij}\}$ satisfying (14) represents a class in $\check{H}^1(\{U_i\}, \underline{G})$, where $\underline{G}$ is the sheaf of smooth G -valued functions. For nonabelian G this object is a *pointed set* of cocycle classes—the distinguished point being the class of the trivial cocycle—and not a group; two cocycles are equivalent when they are cohomologous in the sense $h'_{ij} = \lambda_i^{-1} h_{ij} \lambda_j$ for smooth $\lambda_i: U_i \rightarrow G$. Passing to the direct limit over refinements gives the pointed set $\check{H}^1(M, \underline{G})$, which classifies principal G -bundles over M up to isomorphism [11]. Only for abelian G does this pointed set carry a natural group structure; we use it below only as a pointed set.

Assumption 3.3 (Regularity of the descent data). The cover $\{U_i\}$ is locally finite; the g_i , A_i , Φ_i , h_{ij} are smooth on their domains; and each g_i has the signature fixed in Convention 2.1. If a weaker regularity class is used in $\mathcal{C}_{\text{geo}}$, the same statements hold with “smooth” replaced uniformly by that class, provided the class is closed under restriction and under the gauge transformations appearing in (16)–(17).

3.2 The descent theorem

Theorem 3.4 (Geometric descent). *Let $\{g_i, h_{ij}, A_i, \Phi_i\}$ be compatible local geometric data in the sense of Definition 3.1, satisfying Assumption 3.3. Then these data determine:*

- (1) a global Lorentzian metric g on M with $g|_{U_i} = g_i$;
- (2) a principal G -bundle $\pi: P \rightarrow M$ with local trivializations over $\{U_i\}$ whose transition functions are $\{h_{ij}\}$;
- (3) a global principal connection A on P whose local representatives (pullbacks along the canonical local sections s_i) are the A_i ;
- (4) a global section Φ of $P \times_\rho W$ whose local representatives are the Φ_i .

The resulting global tuple (g, P, A, Φ) is unique up to the bundle isomorphisms induced by change of local trivialization. Conversely, every global tuple (g, P, A, Φ) over M arises in this way from compatible local data, for any cover trivializing P .

The theorem reconstructs the metric g ; it does not by itself produce a time orientation. Membership of (M, g) in the time-oriented admissibility class $\mathcal{C}_{\text{geo}}$ of Definition 2.4 is an additional global condition on the glued metric, discussed in Remark 3.6.

Proof. (1) *The metric.* Condition (15) says exactly that the intrinsic tensors g_i and g_j restrict to the same tensor field on $U_i \cap U_j$. Since $\Gamma(S^2 T^* -)$ is a sheaf on M , the gluing axiom yields a unique global section $g \in \Gamma(S^2 T^* M)$ with $g|_{U_i} = g_i$ for all i . Nondegeneracy and signature are pointwise properties of the symmetric bilinear form g_x on $T_x M$; every $x \in M$ lies in some U_i , where $g_x = (g_i)_x$ is Lorentzian by hypothesis. Hence g is Lorentzian. (If one prefers to present the data in coordinate components, the equivalent hypothesis is (13), which expresses that the $\check{g}_i$ are the component arrays of a single intrinsic tensor; the invariance of signature under the linear isomorphism $d\psi_{ij}(x)$ is then Sylvester’s law of inertia.)

(2) *The bundle.* Form the disjoint union $\coprod_{i \in I} U_i \times G \times \{i\}$ and impose

$$(x, a, i) \sim (x, h_{ji}(x)a, j), \quad x \in U_i \cap U_j, a \in G. \quad (18)$$

Reflexivity follows from $h_{ii} = 1$, symmetry from $h_{ij} = h_{ji}^{-1}$, and transitivity from $h_{kj}h_{ji} = h_{ki}$; hence (18) is an equivalence relation. Let P be the quotient with the quotient topology, $\pi[(x, a, i)] := x$.

We record the trivializations rigorously. Fix i and $x \in U_i$. Every equivalence class lying over x has exactly one representative whose third entry is i : existence, because if a class contains (x, b, j) then it contains $(x, h_{ij}(x)b, i)$; uniqueness, because $(x, a, i) \sim (x, a', i)$ forces $a' = h_{ii}(x)a = a$. Hence the map

$$\varphi_i: \pi^{-1}(U_i) \longrightarrow U_i \times G, \quad \varphi_i([(x, a, i)]) = (x, a) \quad (19)$$

is a well-defined bijection, with inverse $\varphi_i^{-1}(x, a) = [(x, a, i)]$; it is a homeomorphism by definition of the quotient topology, and we give P the smooth structure for which the φ_i are diffeomorphisms (consistent, since the change of trivialization $\varphi_i \circ \varphi_j^{-1}(x, a) = (x, h_{ij}(x)a)$ is smooth). The right G -action $[(x, a, i)] \cdot b := [(x, ab, i)]$ is well defined, because h_{ji} acts by left multiplication and hence commutes with right multiplication; it is free and transitive on fibres,

and each φ_i is G -equivariant. Thus $\pi: P \rightarrow M$ is a principal G -bundle with transition functions h_{ij} .

(3) *The connection.* Let $s_i: U_i \rightarrow P$ be the canonical local sections $s_i(x) := \varphi_i^{-1}(x, 1)$, so that $s_j = s_i \cdot h_{ij}$ on $U_i \cap U_j$. A principal connection $A \in \Omega^1(P, \mathfrak{g})$ is determined by its local representatives $A_i := s_i^* A$, and a family $\{A_i\}$ arises from a (necessarily unique) global A if and only if it obeys the law induced by $s_j = s_i \cdot h_{ij}$, namely

$$s_j^* A = \text{Ad}_{h_{ij}^{-1}}(s_i^* A) + h_{ij}^* \theta_G,$$

which is precisely (16). Given (16), define A over U_i by the pullback identity

$$(\varphi_i^{-1})^* A = \text{Ad}_{a^{-1}}(\pi_U^* A_i) + \theta_G \quad \text{on} \quad U_i \times G, \quad (20)$$

where $\pi_U: U_i \times G \rightarrow U_i$ and the projection $U_i \times G \rightarrow G$ are understood, a denotes the G -coordinate, and θ_G is pulled back along that projection. Formula (20) is an equality of $\mathfrak{g}$ -valued 1-forms on the manifold $U_i \times G$; it is not an evaluation at a point. Two such local definitions agree on $\pi^{-1}(U_i \cap U_j)$: transporting (20) for j through $\varphi_i \circ \varphi_j^{-1}(x, a) = (x, h_{ij}(x)a)$ and using the standard identity $(L_h \circ)^* \theta_G = \theta_G$ together with the Maurer–Cartan expansion of $d(h_{ij}a)$ reproduces (20) for i exactly when (16) holds. The resulting 1-form is G -equivariant, $R_b^* A = \text{Ad}_{b^{-1}} A$, and reproduces the fundamental vector fields, $A(\zeta_X) = X$; hence it is a principal connection [6].

(4) *The matter section.* Sections of $P \times_\rho W$ correspond bijectively to ρ -equivariant maps $\tilde{\Phi}: P \rightarrow W$ with $\tilde{\Phi}(p \cdot a) = \rho(a^{-1})\tilde{\Phi}(p)$, and such maps correspond to families $\Phi_i := \tilde{\Phi} \circ s_i$ obeying $\Phi_j = \tilde{\Phi} \circ (s_i \cdot h_{ij}) = \rho(h_{ij}^{-1})\Phi_i$, which is (17). Hence the Φ_i glue to a unique global section Φ .

Uniqueness up to isomorphism. Suppose two systems of local trivializations over a common cover (or over covers admitting a common refinement, after passing to it) produce compatible data $\{g_i, h_{ij}, A_i, \Phi_i\}$ and $\{g_i, h'_{ij}, A'_i, \Phi'_i\}$ with cohomologous cocycles, $h'_{ij} = \lambda_i^{-1} h_{ij} \lambda_j$ for smooth $\lambda_i: U_i \rightarrow G$, together with $A'_i = \text{Ad}_{\lambda_i^{-1}} A_i + \lambda_i^* \theta_G$ and $\Phi'_i = \rho(\lambda_i^{-1})\Phi_i$. Then the maps $\varphi_i'^{-1}(x, a) \mapsto \varphi_i^{-1}(x, \lambda_i(x)a)$ agree on overlaps—exactly because the cocycles are cohomologous—and assemble into a G -equivariant bundle isomorphism $P' \rightarrow P$ covering id_M , carrying A' to A and Φ' to Φ . That is, it is a vertical arrow of $\mathfrak{G}_{\text{geo}}$. Conversely, an arrow of $\mathfrak{G}_{\text{geo}}$ between global tuples, read in local trivializations, exhibits the two systems of local data as related in exactly this way. Hence the global tuple is well defined precisely as an object of $\mathfrak{G}_{\text{geo}}$ up to isomorphism.

Converse. Given a global tuple, choose any cover $\{U_i\}$ trivializing P (possible since M is paracompact and P is locally trivial); the induced $g_i := g|_{U_i}$, h_{ij} , $A_i = s_i^* A$, Φ_i satisfy (14)–(17) by the computations above. $\square$

Remark 3.5 (Changing the coordinate atlas is not a base diffeomorphism). The uniqueness clause above is stated for changes of *local trivialization*, which induce *vertical* bundle isomorphisms covering id_M . A change of *coordinate atlas* is a different operation: re-charting M changes the component arrays $\tilde{g}_i$ of (12) according to (13) while leaving the intrinsic tensor g —and hence the object Ψ and its class $[\Psi] \in S_{\text{geo}}$ —entirely unchanged. It does not automatically produce a nontrivial base diffeomorphism, and it does not move a point of S_{geo} . Nontrivial base diffeomorphisms enter only through the arrows of $\mathfrak{G}_{\text{geo}}$ with $f \neq \text{id}_M$, which act on the geometry itself and not merely on its coordinate presentation.

Remark 3.6 (Descent and time orientation). Theorem 3.4 glues the local metrics to a global Lorentzian g , but it does not produce the time orientation $\mathfrak{t}$ that Definition 2.4 demands. Whether g admits any time orientation at all is a global topological condition on (M, g) (it holds if and only if the timelike cone bundle of g has a global continuous section of one of its components), and no purely local compatibility condition can force it. Two consistent options are available, and a model must declare one.

Option (a): impose it after descent. Glue g as above, then require, as an additional global admissibility condition, that g be time-orientable and that a time orientation $\mathfrak{t}$ be chosen and recorded in $\mathcal{B}$. Under this option Theorem 3.4 reconstructs the metric, and membership of the reconstructed configuration in $\mathcal{C}_{\text{geo}}$ is a further global condition, verified separately.

Option (b): descend it. Adjoin to the local data a choice, for each i , of a g_i -timelike vector field $\tau_i \in \Gamma(TU_i)$, and impose the additional compatibility condition that on every nonempty overlap τ_i and τ_j lie in the *same* component of the timelike cone,

$$g_i(\tau_i, \tau_j) < 0 \quad \text{on } U_i \cap U_j \quad (21)$$

(with signature $(-, +, +, +)$, so that two timelike vectors have negative inner product exactly when they are causally co-oriented). Condition (21) is a genuine gluing condition: the relation “lies in the same cone component as” is an equivalence relation on timelike vectors at each point, and the local choices τ_i define, by (21), a coherent global selection of one component of the cone bundle at every $x \in M$. That selection is by definition a time orientation $\mathfrak{t}$ of the glued g , and it is unique given the τ_i . Descent then reconstructs the pair $(g, \mathfrak{t})$, and no further global condition on time orientation is needed. (A partition of unity subordinate to $\{U_i\}$ shows that $\mathfrak{t}$ is realized by an actual global timelike vector field, since the cone components are convex.)

Corollary 3.7 (Realizability boundary). *A collection of proposed local metric, connection, and matter expressions represents a global geometric source configuration only if the cocycle and overlap conditions (14)–(17) hold on all overlaps, and only if the resulting global tuple lies in the declared admissibility class $\mathcal{C}_{\text{geo}}$ —which, by Remark 3.6, includes the global time-orientation condition. Consequently, an abstract tuple of relation values that admits no compatible descent data lies outside the geometric realization domain: it is a point of the abstract product $T \times \prod_{\alpha} V_{\alpha}$ that is not in the image of any geometric evaluation map.*

Proof. Immediate from Theorem 3.4: any global configuration restricts to compatible local data, so failure of compatibility for *every* presentation of the proposed values is an obstruction to the existence of a configuration presenting them. Membership in $\mathcal{C}_{\text{geo}}$ is a further restriction, since Theorem 3.4 produces a global tuple but does not certify that it satisfies the declared causal, boundary, or topological conditions. $\square$

Corollary 3.7 is the first and most basic restriction on the abstract completion of Paper 3: *its coordinates must be jointly realizable by global geometric objects.* Two further restrictions—invariance (Section 4) and separation (Section 5)—follow.

Remark 3.8 (Obstructions are genuine). That the constraint is not vacuous is visible already in cohomological terms: the isomorphism class of P is the class of $\{h_{ij}\}$ in the pointed set $\check{H}^1(M, \underline{G})$ (Remark 3.2), and characteristic numbers computed from A are constrained by that class [4, 8]. A demand such as “a $U(1)$ bundle over S^2 with first Chern number $\frac{1}{2}$ ” fails at the cocycle stage; there is no compatible descent datum, hence no configuration, hence no point of $\mathcal{S}_{\text{geo}}$, hence no point of the realized completion. The Cartesian product of independently nominated relation values is, in general, strictly larger than what geometry attains.

4 Curvature, holonomy, and invariant global relations

4.1 Curvature and the Bianchi identity

The curvature of the connection A is the $\mathfrak{g}$ -valued 2-form

$$F_A = dA + A \wedge A \quad (\text{equivalently } F_A = dA + \tfrac{1}{2}[A \wedge A]), \quad (22)$$

in the matrix convention of Definition 3.1. Under a gauge transformation represented locally by $h: U \rightarrow G$,

$$A \mapsto A^h = h^{-1}Ah + h^{-1}dh, \quad F_A \mapsto F_{A^h} = h^{-1}F_A h, \quad (23)$$

so F_A transforms in the adjoint representation: it is a section of $\Lambda^2 T^*M \otimes \text{ad}P$, and is *not* itself gauge invariant, though its Ad-invariant polynomial functions are. The Bianchi identity is

$$D_A F_A = dF_A + [A, F_A] = 0, \quad (24)$$

an identity, not an equation of motion [6].

The Levi-Civita connection ∇ of g is the unique affine connection with

$$\nabla g = 0, \quad \text{Tor}(\nabla) = 0, \quad (25)$$

with Riemann tensor $R^\rho{}_{\sigma\mu\nu}$, Ricci tensor $R_{\mu\nu} = R^\rho{}_{\mu\rho\nu}$, scalar curvature $R = g^{\mu\nu} R_{\mu\nu}$, and Einstein tensor $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2}R g_{\mu\nu}$ [12].

4.2 Invariant global relations

Definition 4.1 (Invariant geometric relation). A map $J: \mathcal{C}_{\text{geo}} \rightarrow V$ into a set V is an *invariant geometric relation* if it is constant on the arrows of $\mathfrak{G}_{\text{geo}}$, i.e. $J(\Psi) = J(\Psi')$ whenever there is an arrow $\Psi \rightarrow \Psi'$. Equivalently, J factors uniquely through $q: \mathcal{C}_{\text{geo}} \rightarrow S_{\text{geo}}$, defining $\bar{J}: S_{\text{geo}} \rightarrow V$. We write J for both maps when no confusion arises.

Invariance under $\mathfrak{G}_{\text{geo}}$ means invariance under *both* gauge transformations and allowed diffeomorphisms (including preservation of orientation and time orientation). Both are required; neither alone suffices.

Example 4.2 (Integral invariants). Assume M compact without boundary, or that $\mathcal{C}_{\text{geo}}$ imposes decay/boundary conditions rendering the following integrals absolutely convergent. Then each of

$$J_{\text{EH}}(\Psi) = \int_M R_g \, \text{dvol}_g, \quad J_{\text{YM}}(\Psi) = \int_M \langle F_A \wedge *F_A \rangle, \quad J_{\text{top}}(\Psi) = \int_M \langle F_A \wedge F_A \rangle \quad (26)$$

is an invariant geometric relation. Gauge invariance of the latter two follows from (23) with the Ad-invariance of $\langle \cdot, \cdot \rangle$; diffeomorphism invariance follows because each integrand is a top-degree form built naturally from the fields and $\int_M f^* \omega = \int_M \omega$ for orientation-preserving f (the Hodge star is natural: $f^*(*_g \eta) = *_g(f^* \eta)$ when $f_* g = g'$ and f preserves orientation). Concerning normalization, we make no claim beyond the following, which is all that is used: for M compact oriented without boundary and G compact, the de Rham class of the Chern–Weil form $\langle F_A \wedge F_A \rangle$ is independent of A and represents (up to a normalization constant depending on the chosen Ad-invariant form $\langle \cdot, \cdot \rangle$ and on the conventions for the Chern–Weil homomorphism) a characteristic class of P in $H_{\text{dR}}^4(M; \mathbb{R})$ [4, 8, 7]. Hence J_{top} is constant on each topological sector. We do *not* assert a specific integrality normalization here; any such normalization is a convention that must be fixed together with $\langle \cdot, \cdot \rangle$, and none of the results below depend on the choice.

Example 4.3 (Spectral relations). Let $\mathcal{D}[\Psi]$ be an operator constructed naturally from the configuration data $(g, A, \Phi, \mathcal{B})$, together with a *specified domain* determined by $\mathcal{B}$. Then any function of $\mathcal{D}[\Psi]$ that is invariant under isomorphism is an invariant geometric relation. The spectrum, when it is well defined, is the leading example. The precise statement requires care, and we distinguish two cases.

(i) *Intertwining*. An arrow $F: \Psi \rightarrow \Psi'$ of $\mathfrak{G}_{\text{geo}}$ induces a linear isomorphism $\mathcal{U}_F$ from the space of sections over Ψ to the space of sections over Ψ' , carrying the domain of $\mathcal{D}[\Psi]$ onto the domain of $\mathcal{D}[\Psi']$ and satisfying the *operator intertwining* relation

$$\mathcal{D}[\Psi'] \circ \mathcal{U}_F = \mathcal{U}_F \circ \mathcal{D}[\Psi]. \quad (27)$$

This holds by naturality of the construction and requires no metric structure on W . It already implies that any relation depending only on the operator up to linear isomorphism (its spectrum as a set, when the operator is, say, closed with nonempty resolvent set) is $\mathfrak{G}_{\text{geo}}$ -invariant.

(ii) *Unitary equivalence.* If, in addition, the representation ρ is unitary (or, more generally, W carries a ρ -invariant Hermitian inner product), the associated bundle carries a bundle metric preserved by $\hat{F}_\rho$, the base measure dvol_g is preserved by the allowed f , and the declared boundary conditions define a self-adjoint realization of $\mathcal{D}$ on the resulting Hilbert space, then $\mathcal{U}_F$ is *unitary* and (27) is a unitary equivalence. In that case the spectrum as a *multiset* (with multiplicities), the heat trace $\text{tr} e^{-t\mathcal{D}}$, zeta values, and eigenvalue gaps are all invariant geometric relations.

An arbitrary finite-dimensional $\rho: G \rightarrow \text{GL}(W)$ is not unitary, so case (ii) requires its stated hypotheses; absent them, only the intertwining statement (i) is available. In either case the specification of the domain is essential: without it the spectrum is not a function of $[\Psi]$, and distinct self-adjoint extensions yield distinct invariant relations.

4.3 Holonomy and Wilson observables

Let γ be a piecewise-smooth closed curve in M based at $x \in M$, and let $p \in P_x$. Parallel transport of A along γ starting at p ends at $p \cdot a$ for a unique $a \in G$; writing $\text{Hol}_\gamma^p(A) := a$, a change of starting point $p \mapsto p \cdot b$ gives $\text{Hol}_\gamma^{p \cdot b}(A) = b^{-1} \text{Hol}_\gamma^p(A) b$. In a local trivialization with gauge h , this reads

$$\text{Hol}_\gamma(A^h) = h(x)^{-1} \text{Hol}_\gamma(A) h(x). \quad (28)$$

Proposition 4.4 (Gauge invariance of holonomy class and Wilson observables). *For a fixed based loop γ , the conjugacy class*

$$[\text{Hol}_\gamma(A)]_{\text{conj}} \in G/\sim_{\text{conj}} \quad (29)$$

is independent of the choice of $p \in P_x$ and is invariant under all gauge transformations. Consequently, for any finite-dimensional representation σ of G , the Wilson observable

$$W_{\gamma,\sigma}(A) = \text{tr}_\sigma \text{Hol}_\gamma(A) \quad (30)$$

is gauge invariant. If A is flat, $W_{\gamma,\sigma}$ depends on γ only through its free homotopy class.

Proof. Independence of p and gauge invariance are immediate from (28), since a conjugacy class is invariant under conjugation and tr_σ is a class function. The final statement holds because for flat A parallel transport depends only on the homotopy class of the path rel endpoints, and a free homotopy of based loops changes the holonomy by conjugation, which tr_σ does not see [6]. $\square$

4.4 Loops and diffeomorphism invariance

Proposition 4.4 concerns *gauge* invariance only. A Wilson observable attached to a *fixed* loop $\gamma \subset M$ is in general *not* invariant under the allowed diffeomorphisms, because a diffeomorphism f moves the loop: $W_{\gamma,\sigma}(f_*A) = W_{f^{-1} \circ \gamma, \sigma}(A)$, and $f^{-1} \circ \gamma \neq \gamma$ in general. Three constructions restore genuine $\mathfrak{G}_{\text{geo}}$ -invariance. They are mutually exclusive as declarations, and a model must state which it adopts; each is used somewhere below.

(L1) Include the loop in the configuration. Enlarge the tuple (2) to $(M, g, P, A, \Phi, \mathcal{B}, \gamma)$ and extend the arrows of $\mathfrak{G}_{\text{geo}}$ to carry γ to $f \circ \gamma$. Then $\Psi \mapsto W_{\gamma,\sigma}(A)$ is invariant on the enlarged groupoid by construction: $W_{f \circ \gamma, \sigma}(f_*A) = W_{\gamma,\sigma}(A)$, since parallel transport is natural under bundle isomorphisms. The loop is then part of the source data, not a background structure, and S_{geo} is correspondingly enlarged.

(L2) Use an invariant constructed from the whole loop space. Examples: the *set* of Wilson values $\{W_{\gamma,\sigma}(A) : \gamma \in \Omega(M, x)\} \subseteq \mathbb{C}$ (a subset of $\mathbb{C}$, hence unaffected by any relabelling of loops); the holonomy group $\text{Hol}_x(A) \leq G$ up to conjugacy; or, for flat A , the holonomy homomorphism $\pi_1(M, x) \rightarrow G$ taken modulo conjugation *and* modulo the outer action of f_*

on $\pi_1(M, x)$. Each is $\mathfrak{G}_{\text{geo}}$ -invariant without enlarging the configuration, precisely because the diffeomorphism permutes the loops over which the construction ranges.

(L3) Let the declared data pin a homology class. If $\mathcal{B}$ distinguishes a class $c \in H_1(M; \mathbb{Z})$ that every allowed diffeomorphism preserves, and if the observable depends on γ only through c , then it is invariant. This is what happens in Section 6: there $G = U(1)$ is abelian and A is flat, so $\text{Hol}_\gamma(A)$ depends on γ only through its class in $H_1(S^1; \mathbb{Z}) \cong \mathbb{Z}$; the allowed diffeomorphisms are declared to be the orientation-preserving ones, which act trivially on $H_1(S^1; \mathbb{Z})$; hence $\text{Hol}_{\gamma_0}(A)$ is a genuine function on S_{geo} . (Under an orientation-reversing diffeomorphism the generator is sent to its inverse and $\text{Hol} \mapsto \text{Hol}^{-1}$; this is exactly why the orientation is part of the declared data.)

Remark 4.5 (Curvature does not determine transport). Local curvature and global transport are related but not interchangeable. Curvature is the infinitesimal holonomy: the Ambrose–Singer theorem [2, 6] states that the Lie algebra of the holonomy group of a connection on a connected principal bundle is spanned by the curvature values $F_A(X, Y)$ at points reachable from the base point by horizontal curves, transported back along those curves. Nevertheless holonomy around noncontractible cycles is not determined by F_A : flat connections with $F_A \equiv 0$ can have arbitrary holonomy homomorphisms $\pi_1(M) \rightarrow G$ up to conjugacy. This is not a formal curiosity; it is the content of the Aharonov–Bohm effect [1] and of the Wu–Yang analysis of the global significance of gauge potentials [13], in which the physically operative datum is a holonomy in a region where the field strength vanishes identically. Section 6 is the sharpest mathematical illustration: there $F_A \equiv 0$ identically, and yet the reduced source is a full circle of physically inequivalent configurations, completely parametrized by holonomy.

Remark 4.6 (No finite separating list of curvature scalars). It is not assumed anywhere below that a finite family of curvature scalars—or any fixed finite family of local invariants—separates arbitrary Lorentzian geometries. It does not: in Lorentzian signature there exist non-flat metrics all of whose polynomial scalar curvature invariants vanish identically, and which therefore cannot be distinguished from flat space by any such family; the plane-wave and more general Kundt-type “VSI” spacetimes are the standard examples [10]. Separation is therefore never invoked as a general fact. Where it is needed it is proved on the restricted domain of the model, as in Theorem 6.5.

5 Geometric realization of the minimal completion

5.1 The geometric evaluation map

Definition 5.1 (Declared readout and relation family). Let

$$p_{\text{geo}}: S_{\text{geo}} \longrightarrow T \tag{31}$$

be a declared Tier-1 readout: a map, defined on the reduced source, recording exactly the data retained by the Tier-1 description of the model. Let A be an index *set* and let

$$J_\alpha: S_{\text{geo}} \longrightarrow V_\alpha, \quad \alpha \in A, \tag{32}$$

be invariant geometric relations in the sense of Definition 4.1. Write $J_A = (J_\alpha)_{\alpha \in A}$ and

$$j_A^{\text{geo}} = (p_{\text{geo}}, J_A): S_{\text{geo}} \longrightarrow T \times \prod_{\alpha \in A} V_\alpha. \tag{33}$$

Definition 5.2 (Geometrically realized completion). The *geometrically realized completion* of p_{geo} relative to $\{J_\alpha\}_{\alpha \in A}$ is

$$E_A^{\text{geo}} := \text{Im } j_A^{\text{geo}} \subseteq T \times \prod_{\alpha \in A} V_\alpha, \tag{34}$$

with corestriction $e_A^{\text{geo}}: S_{\text{geo}} \rightarrow E_A^{\text{geo}}$ and forgetful map

$$u_A: E_A^{\text{geo}} \rightarrow T, \quad u_A(t, (v_\alpha)) = t, \quad (35)$$

so that $u_A \circ e_A^{\text{geo}} = p_{\text{geo}}$ by construction.

Note that E_A^{geo} is the image, *not* the ambient product. This restriction performs the essential mathematical work: it enforces all the compatibility relations among metric, bundle, connection, matter, boundary, and invariant data that Section 3 exhibited. The Cartesian product of independently written coordinate values is in general *not* realizable (Corollary 3.7, Remark 3.8).

5.2 The restricted realization theorem

Theorem 5.3 (Restricted geometric realization). *Fix an admissibility class $\mathcal{C}_{\text{geo}}$, a readout p_{geo} , and a family $\{J_\alpha\}_{\alpha \in A}$ of invariant geometric relations on S_{geo} . Then:*

- (1) (Realizability.) *For every $\xi \in E_A^{\text{geo}}$ there exists $[\Psi] \in S_{\text{geo}}$ with $e_A^{\text{geo}}([\Psi]) = \xi$.*
- (2) (Factorization.) *$p_{\text{geo}} = u_A \circ e_A^{\text{geo}}$ and $J_\alpha = \text{pr}_\alpha \circ e_A^{\text{geo}}$ for every α .*
- (3) (Terminality.) *If $\pi: S_{\text{geo}} \rightarrow E$ is any surjection through which p_{geo} and every J_α factor—say $p_{\text{geo}} = u \circ \pi$ and $J_\alpha = \kappa_\alpha \circ \pi$ —then there is a unique map $\theta: E \rightarrow E_A^{\text{geo}}$ with $\theta \circ \pi = e_A^{\text{geo}}$; θ is surjective and satisfies $u_A \circ \theta = u$ and $\text{pr}_\alpha \circ \theta = \kappa_\alpha$.*

Thus E_A^{geo} is the canonical minimal completion of p_{geo} relative to $\{J_\alpha\}$, in the sense of Paper 3 [18], computed on the geometric source S_{geo} .

Proof. (1) is immediate: E_A^{geo} is by definition the image of j_A^{geo} . (2) holds by definition of j_A^{geo} and of the projections; well-definedness of p_{geo} and each J_α as maps on S_{geo} is exactly the invariance hypothesis of Definition 4.1, which is where gauge and diffeomorphism invariance are consumed.

(3) Define $\theta(e) := e_A^{\text{geo}}(s)$ for any $s \in \pi^{-1}(e)$; the fibre is nonempty since π is surjective. This is well defined: if $\pi(s) = \pi(s') = e$ then $p_{\text{geo}}(s) = u(\pi(s)) = u(\pi(s')) = p_{\text{geo}}(s')$ and likewise $J_\alpha(s) = \kappa_\alpha(\pi(s)) = \kappa_\alpha(\pi(s')) = J_\alpha(s')$ for every α ; hence $j_A^{\text{geo}}(s) = j_A^{\text{geo}}(s')$, i.e. $e_A^{\text{geo}}(s) = e_A^{\text{geo}}(s')$. By construction $\theta \circ \pi = e_A^{\text{geo}}$, and since e_A^{geo} is surjective onto E_A^{geo} , so is θ . Uniqueness: any θ' with $\theta' \circ \pi = e_A^{\text{geo}}$ agrees with θ on $\pi(S_{\text{geo}}) = E$, hence everywhere. The compatibilities follow by evaluating on $\pi(s)$ and using surjectivity of π .

This is the terminal universal property established in Paper 3 [18], applied to the source set S_{geo} , the readout p_{geo} , and the relation family $\{J_\alpha\}$; the additional content here is that S_{geo} is the *reduced* geometric source and that the J_α are *invariant* geometric relations, so that all maps involved are functions of physical orbits rather than of gauge or coordinate representatives. $\square$

Remark 5.4 (Relative, not absolute). Theorem 5.3 is *restricted* in two independent senses, both essential. First, it realizes the completion only *relative to the declared family* $\{J_\alpha\}$: a different family gives a different E_A^{geo} , generally neither finer nor coarser. Second, it realizes only what geometry attains: E_A^{geo} is an image, and Corollary 3.7 shows that this image can be a proper subset of the abstract product. Neither the theorem nor anything else in this paper asserts that an arbitrarily nominated abstract completion E_A of Paper 3 is geometrically realizable. It is realizable exactly when it is isomorphic, over T , to E_A^{geo} for some admissible $\mathcal{C}_{\text{geo}}$, readout, and invariant relation family—and establishing that in a given model is work.

5.3 When does the completion reconstruct the source?

Proposition 5.5 (Orbit reconstruction criterion). *The completion map $e_A^{\text{geo}}: S_{\text{geo}} \rightarrow E_A^{\text{geo}}$ is injective—equivalently, bijective onto E_A^{geo} —if and only if the family*

$$\{p_{\text{geo}}\} \cup \{J_\alpha : \alpha \in A\} \quad (36)$$

separates points of S_{geo} : for all $[\Psi], [\Psi'] \in S_{\text{geo}}$, equality of the readouts and of every J_α forces $[\Psi] = [\Psi']$. In that case $S_{\text{geo}} \cong E_A^{\text{geo}}$ as sets over T .

Proof. By construction $e_A^{\text{geo}}([\Psi]) = e_A^{\text{geo}}([\Psi'])$ if and only if the readouts agree and every J_α agrees. Injectivity of e_A^{geo} is therefore literally the separation property. Surjectivity onto E_A^{geo} always holds, by Theorem 5.3(1). The identification is over T because $u_A \circ e_A^{\text{geo}} = p_{\text{geo}}$. $\square$

Proposition 5.5 converts a question about completions into a question of invariant theory: *does the declared family separate the orbits?* No general answer exists, and none is asserted (Remark 4.6). Separation must be proved on the restricted domain relevant to the model. Section 6 carries out exactly such a proof in a case where the answer is affirmative and complete.

6 Classification of flat $U(1)$ connections on an oriented circle

This section supplies an exact classification example: a restricted domain on which the separation hypothesis of Proposition 5.5 can be *proved*, and on which the geometrically realized completion can be identified on the nose.

6.1 The model and its conventions

Convention 6.1 (Scope of the circle model). This section is a deliberate one-dimensional specialization, and two stipulations of the general framework are suspended for it. Both are recorded here, since leaving either implicit would create a genuine inconsistency with Section 2.

Dimension. The base is $M = S_L^1$, so $\dim M = 1$ rather than the $\dim M = 4$ of Convention 2.1, and the metric is Riemannian rather than Lorentzian. Nothing in Sections 2–5 used the dimension or the signature: the groupoid, the orbit space, Definition 4.1, Theorem 5.3, and Proposition 5.5 are all dimension- and signature-agnostic, and it is only these that the present section invokes. The descent theorem and the variational apparatus of Sections 7–8 are not used here.

The metric and the allowed maps. A tension must be resolved explicitly. The model fixes the standard metric dx^2 on S_L^1 while also admitting *every* orientation-preserving diffeomorphism as an allowed base map; under the groupoid of Definition 2.8, whose arrows satisfy $f_*g = g'$, a fixed metric would instead restrict the allowed base maps to its isometries, namely the rigid rotations. Of the three consistent ways to proceed—restricting the allowed diffeomorphisms to orientation-preserving isometries; transforming the metric along with the connection; or defining the circle source purely as the gauge quotient and treating the diffeomorphism question separately—we adopt the third, which is the formulation the arguments below actually use.

Accordingly, in this section the metric dx^2 and its measure dx are treated as *fixed background structure*, not as part of the configuration: the configuration is the connection A alone, and the reduced circle source is *defined* to be the gauge quotient

$$S_{\text{geo}} := \mathcal{C}_{\text{geo}}^\circ / \text{Gau}(P). \quad (37)$$

The background metric is what gives the operator $\mathcal{D}_A$ of (60) its meaning, and it is not transformed by the maps considered below. The role of the orientation-preserving diffeomorphisms is then a separate matter, treated separately: they are *not* arrows acting on a metric they would have to preserve, but are used solely through the fact that they fix the homology generator c ,

and hence preserve holonomy. Lemma 6.8(1) shows that they consequently identify no classes beyond gauge equivalence, so that (37) loses nothing by omitting them; and Lemma 6.8(2) shows that the spectral relation is invariant under them for the same reason. Definition 6.2 is to be read in this light.

Definition 6.2 (The circle model). Fix $L > 0$ and let $M = S_L^1 := \mathbb{R}/L\mathbb{Z}$ with its standard orientation and the global angular coordinate x (well defined modulo L). Let $G = U(1)$ and let $P = S_L^1 \times U(1)$ be the trivial bundle; this is no loss of generality, since every principal $U(1)$ -bundle over S^1 is trivial (the classifying space argument gives $[S^1, BU(1)] \cong H^2(S^1; \mathbb{Z}) = 0$; equivalently, every $U(1)$ -bundle over a one-complex is trivial [11]). Let $\mathcal{C}_{\text{geo}}^\circ$ be the class of smooth connections A on P , and let $\text{Gau}(P) = C^\infty(S_L^1, U(1))$ be the gauge group; by Convention 6.1 the configuration is the connection alone, and the reduced source is the gauge quotient (37). The declared data $\mathcal{B}$ consist of the orientation of S_L^1 , the resulting generator $c = [\gamma_0] \in H_1(S_L^1; \mathbb{Z}) \cong \mathbb{Z}$ of the fundamental cycle, $\gamma_0(s) = s$, $s \in [0, L]$, and the fixed background metric dx^2 with its measure. The orientation-preserving diffeomorphisms of S_L^1 fix c ; holonomy around c is therefore a legitimate function on orbits by mechanism (L3) of Section 4.4, and its invariance under those diffeomorphisms is established in Lemma 6.8.

Write a connection in the global trivialization as

$$A = a(x) dx, \quad a \in C^\infty(S_L^1, \mathbb{R}). \quad (38)$$

Because $\dim M = 1$, every 2-form on M vanishes; hence

$$F_A = dA \equiv 0 \quad \text{for every } A. \quad (39)$$

Every connection on the circle is flat. Local curvature is identically zero on this entire domain and therefore carries no information whatsoever; whatever separates orbits here must be global.

Convention 6.3 (Gauge action and charged covariant derivative). Identify $\mathfrak{u}(1) = i\mathbb{R}$ and write connections with the real potential a as in (38). A gauge transformation is a smooth map $u: S_L^1 \rightarrow U(1)$, acting on the connection by

$$A \mapsto A^u = A - i u^{-1} du, \quad (40)$$

which is real-valued because $-i u^{-1} du$ is real for $|u| = 1$. On a charged scalar $\phi \in C^\infty(S_L^1, \mathbb{C})$ of unit charge, the same u acts by

$$\phi \mapsto \phi^u = u \phi, \quad (41)$$

and the charged covariant derivative is

$$D_A \phi = (\partial_x - i a(x)) \phi dx. \quad (42)$$

These three conventions are mutually consistent. Writing $u = e^{i\varphi}$ locally, so that $-i u^{-1} du = d\varphi$ and hence $a^u = a + \varphi'$, we get

$$(\partial_x - i a^u) \phi^u = (\partial_x - i(a + \varphi'))(e^{i\varphi} \phi) = e^{i\varphi} (i\varphi' \phi + \partial_x \phi - i a \phi - i\varphi' \phi) = u (\partial_x - i a) \phi,$$

so $D_{A^u} \phi^u = u D_A \phi$ transforms covariantly, as required.

Definition 6.4 (Holonomy on the circle). For $A = a dx$ set

$$\text{Hol}(A) := \exp\left(i \int_0^L a(x) dx\right) \in U(1), \quad (43)$$

the holonomy around the declared fundamental cycle γ_0 .

6.2 The classification theorem

Theorem 6.5 (Holonomy classification). *Two smooth $U(1)$ connections $A = a \, dx$ and $A' = a' \, dx$ on S_L^1 are gauge equivalent if and only if*

$$\text{Hol}(A) = \text{Hol}(A'). \quad (44)$$

Consequently $[A] \mapsto \text{Hol}(A)$ is a bijection

$$\mathcal{M}_{\text{flat}}(S^1, U(1)) := \mathcal{C}_{\text{geo}}^\circ / \text{Gau}(P) \xrightarrow{\cong} U(1), \quad (45)$$

and it is a homeomorphism when the left-hand side carries the quotient of the C^∞ topology.

Proof. Step 1: holonomy is gauge invariant. Let $u \in C^\infty(S_L^1, U(1))$. Pull u back to the universal cover: the composite $\tilde{u} := u \circ \text{pr}: [0, L] \rightarrow U(1)$, where $\text{pr}(s) = s \bmod L$, lifts along the covering $\mathbb{R} \xrightarrow{t \mapsto e^{it}} U(1)$ to a smooth $\varphi: [0, L] \rightarrow \mathbb{R}$ with $u(x) = e^{i\varphi(x)}$ for $x \in [0, L]$, because $[0, L]$ is simply connected. Single-valuedness of u on S_L^1 forces

$$\varphi(L) - \varphi(0) = 2\pi n, \quad n = \deg(u) \in \mathbb{Z}. \quad (46)$$

Now $-i u^{-1} du = -i e^{-i\varphi} \cdot i \varphi' e^{i\varphi} dx = \varphi' dx = d\varphi$, so the declared action (40) gives

$$A^u = A - i u^{-1} du = A + d\varphi, \quad \text{i.e.} \quad a^u = a + \varphi'. \quad (47)$$

Hence the period changes by $+2\pi n$:

$$\int_0^L a^u = \int_0^L a + (\varphi(L) - \varphi(0)) = \int_0^L a + 2\pi n,$$

and therefore

$$\text{Hol}(A^u) = \exp\left(i \int_0^L a + 2\pi n\right) = \exp\left(i \int_0^L a\right) = \text{Hol}(A).$$

The gauge group can shift the period integral, but only by integer multiples of 2π , which $\exp(i \cdot)$ does not see. (The sign of the shift is $+2\pi n$ with the declared convention; either sign would of course leave the holonomy unchanged, but we record the correct one.)

Step 2: equal holonomy implies gauge equivalence. Suppose $\text{Hol}(A) = \text{Hol}(A')$. Then $\exp(i \int_0^L (a' - a)) = 1$, so there is $n \in \mathbb{Z}$ with

$$\int_0^L (a'(x) - a(x)) \, dx = 2\pi n. \quad (48)$$

Define

$$\chi(x) := \int_0^x (a'(y) - a(y)) \, dy, \quad x \in [0, L], \quad (49)$$

so that χ is smooth with

$$\chi' = a' - a \quad \text{and} \quad \chi(L) - \chi(0) = 2\pi n \quad (50)$$

by (48). The function χ itself need *not* descend to S_L^1 , but

$$u(x) := e^{i\chi(x)} \quad (51)$$

does: $u(L) = e^{i\chi(0)} e^{2\pi n} = e^{i\chi(0)} = u(0)$, and all derivatives match at the seam because a, a' are smooth and L -periodic, so χ' and all its derivatives are L -periodic. Hence $u \in C^\infty(S_L^1, U(1))$ is

a bona fide gauge transformation, of degree n . By exactly the computation of Step 1 applied to this u (with $\varphi = \chi$),

$$-i u^{-1} du = d\chi = (a' - a) dx = A' - A, \quad (52)$$

using (50). Therefore, directly from the declared action (40) in the form (47),

$$A^u = A - i u^{-1} du = A + (A' - A) = A'.$$

So u itself carries A to A' , and the two connections are gauge equivalent.

Step 3: the bijection. By Steps 1 and 2, Hol descends to an injection $\mathcal{C}_{\text{geo}}^\circ / \text{Gau}(P) \hookrightarrow U(1)$. It is surjective: for $\vartheta \in \mathbb{R}$ the constant connection

$$A_\vartheta := \frac{\vartheta}{L} dx \quad \text{satisfies} \quad \text{Hol}(A_\vartheta) = e^{i\vartheta}, \quad (53)$$

and every element of $U(1)$ is of the form $e^{i\vartheta}$. Hence (45) is a bijection.

Step 4: the homeomorphism. Give $\mathcal{C}_{\text{geo}}^\circ \cong C^\infty(S_L^1, \mathbb{R})$ its usual Fréchet (C^∞) topology and the quotient $\mathcal{M} := \mathcal{C}_{\text{geo}}^\circ / \text{Gau}(P)$ the quotient topology, with projection $\text{pr}_{\mathcal{M}}: \mathcal{C}_{\text{geo}}^\circ \rightarrow \mathcal{M}$. Write $\Theta: \mathcal{M} \rightarrow U(1)$ for the bijection (45).

Θ is continuous: the map $\mathcal{C}_{\text{geo}}^\circ \rightarrow U(1)$, $A \mapsto \exp(i \int_0^L a)$, is continuous (indeed it is continuous already in the L^1 topology, hence in C^∞), and it is constant on $\text{Gau}(P)$ -orbits by Step 1; the universal property of the quotient topology then makes Θ continuous.

For the inverse, we do *not* attempt to choose a global continuous argument ϑ on $U(1)$ —no such choice exists. Instead define

$$\sigma: \mathbb{R} \longrightarrow \mathcal{M}, \quad \sigma(\vartheta) := \text{pr}_{\mathcal{M}}(A_\vartheta) = [\frac{\vartheta}{L} dx]. \quad (54)$$

The map $\mathbb{R} \rightarrow \mathcal{C}_{\text{geo}}^\circ$, $\vartheta \mapsto A_\vartheta$, is continuous (it is linear into the constants, and the C^∞ seminorms of A_ϑ depend continuously on ϑ), so σ is continuous as a composite with $\text{pr}_{\mathcal{M}}$.

σ is 2π -periodic. Indeed $A_{\vartheta+2\pi} - A_\vartheta = \frac{2\pi}{L} dx$, and by Step 2 (with $a = \vartheta/L$, $a' = (\vartheta+2\pi)/L$, so that $\int_0^L (a' - a) = 2\pi$, i.e. $n = 1$) the connections A_ϑ and $A_{\vartheta+2\pi}$ are gauge equivalent; explicitly the gauge transformation is $u(x) = e^{2\pi i x/L}$, of degree 1. Hence $\sigma(\vartheta + 2\pi) = \sigma(\vartheta)$ for all ϑ .

Therefore σ factors through the quotient $\mathbb{R} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$, and since $\mathbb{R}/2\pi\mathbb{Z} \cong U(1)$ via $\vartheta \bmod 2\pi \mapsto e^{i\vartheta}$ (a homeomorphism), we obtain a continuous map

$$\bar{\sigma}: U(1) \longrightarrow \mathcal{M}, \quad \bar{\sigma}(e^{i\vartheta}) = \sigma(\vartheta), \quad (55)$$

well defined and continuous by the universal property of the quotient $\mathbb{R} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ (which is an open, hence quotient, map). Finally $\Theta \circ \bar{\sigma} = \text{id}_{U(1)}$ by (53), and $\bar{\sigma} \circ \Theta = \text{id}_{\mathcal{M}}$ because both sides are the identity on the classes $[A_\vartheta]$, which exhaust $\mathcal{M}$ by Step 3 together with the injectivity of Θ . Hence $\bar{\sigma} = \Theta^{-1}$ is continuous, and Θ is a homeomorphism. $\square$

Remark 6.6 (Where the hypotheses are used). Three hypotheses were used and cannot be dropped. *Smoothness and periodicity* of a, a' give smoothness of u across the seam in Step 2. *Connectedness of S^1* , with $\pi_1(S^1) \cong \mathbb{Z}$ generated by a single class, is what makes one period integral a complete invariant. *Orientation*, fixed in Definition 6.2, is what makes $\int_0^L$ well defined with a sign: reversing orientation sends $\text{Hol} \mapsto \text{Hol}^{-1}$, and it is precisely this ambiguity that reappears in Corollary 6.9. The declaration of the allowed diffeomorphism group as the orientation-preserving one is therefore part of the hypothesis of the theorem.

6.3 The realized completions

Corollary 6.7 (Exact geometric realization of the holonomy completion). *In the circle model, take the curvature readout*

$$p_F: S_{\text{geo}} = \mathcal{M}_{\text{flat}}(S^1, U(1)) \longrightarrow T = \{0\}, \quad p_F([A]) = F_A = 0, \quad (56)$$

which by (39) is constant and hence retains no information, and take the single invariant relation $J_{\text{hol}}([A]) = \text{Hol}(A) \in U(1)$. Then

$$E_{\text{hol}} = \text{Im}(p_F, J_{\text{hol}}) \cong \mathcal{M}_{\text{flat}}(S^1, U(1)) \cong U(1), \quad (57)$$

and the completion map $e_{\text{hol}}: S_{\text{geo}} \rightarrow E_{\text{hol}}$ is a bijection.

Proof. $\text{Im}(p_F, J_{\text{hol}}) = \{0\} \times \text{Im } J_{\text{hol}} = \{0\} \times U(1) \cong U(1)$, by the surjectivity established in Step 3 of Theorem 6.5. The family $\{p_F\} \cup \{J_{\text{hol}}\}$ separates orbits by the injectivity established there, so Proposition 5.5 gives injectivity of e_{hol} ; surjectivity is Theorem 5.3(1). $\square$

Corollary 6.7 is the sharpest illustration in this paper of the mechanism the sequence is about. The readout p_F is maximally degenerate—it is constant—so on its own it distinguishes nothing. Yet the reduced source is not a point: it is a circle of genuinely inequivalent physical configurations. The completion recovers exactly that circle, from a single global relation, on a domain where local curvature vanishes identically. Holonomy here is not one invariant among many: it separates *all* gauge orbits and therefore reconstructs the reduced geometric source completely. This is the mathematical content of the Aharonov–Bohm and Wu–Yang observations [1, 13].

Lemma 6.8 (The reduced circle source, and invariance of the charged spectrum). *In the circle model of Definition 6.2, with the scope fixed by Convention 6.1:*

- (1) *The orientation-preserving diffeomorphisms of S_L^1 identify no connection classes beyond gauge equivalence: for every such f and every A , the pushforward f_*A is gauge equivalent to A . Hence adjoining them to the gauge group would coarsen nothing, and the reduced source (37) may be written*

$$S_{\text{geo}} = \mathcal{C}_{\text{geo}}^\circ / \text{Gau}(P) = \mathcal{M}_{\text{flat}}(S^1, U(1)). \quad (58)$$

- (2) *The spectral multiset $\text{spec } \mathcal{D}_A$ of (60), computed with respect to the fixed background metric dx^2 and measure dx on S_L^1 , is invariant under gauge transformations and under orientation-preserving diffeomorphisms, and is therefore a well-defined invariant relation on S_{geo} .*

Proof. Let f be an orientation-preserving diffeomorphism of S_L^1 and let A be a connection. Since f preserves the orientation, it fixes the fundamental generator $c = [\gamma_0]$ of $H_1(S_L^1; \mathbb{Z})$. Because $G = U(1)$ is abelian and every connection on S_L^1 is flat by (39), the holonomy of a connection around a loop depends on that loop only through its class in $H_1(S_L^1; \mathbb{Z})$; the pushforward f_*A therefore has the same holonomy around c as A does:

$$\text{Hol}(f_*A) = \text{Hol}(A). \quad (59)$$

(Concretely, $\int_{S_L^1} f_*A = \int_{S_L^1} A$ by the change-of-variables formula for the orientation-preserving f , the integrand being a 1-form.)

(1). By (59) and the sufficiency direction of Theorem 6.5, equality of holonomy already implies gauge equivalence; hence f_*A is gauge equivalent to A for every orientation-preserving f and every A . The equivalence relation generated by the gauge group together with the orientation-preserving diffeomorphisms therefore coincides with gauge equivalence alone, so the quotient (37) is unchanged by adjoining them, and (58) holds.

(2). Let A and A' be connections related either by a gauge transformation or by the pushforward along an orientation-preserving diffeomorphism. In the second case part (1) supplies a gauge transformation achieving the same result, so in either case there is $u \in \text{Gau}(P)$ with $A' = A^u$. The operator $\mathcal{D}$ is defined with respect to the fixed background metric and measure of Convention 6.1, which are not transformed; the diffeomorphism therefore plays no further

role, and only the gauge covariance established in Convention 6.3 is needed. There we verified $D_{A^u}\phi^u = u D_A \phi$ for $\phi^u = u\phi$; iterating,

$$\mathcal{D}_{A^u}(u\phi) = u \mathcal{D}_A \phi \quad \text{for all } \phi \in C^\infty(S_L^1, \mathbb{C}).$$

Thus the multiplication operator $\mathcal{U}_u \phi := u\phi$ intertwines $\mathcal{D}_A$ with $\mathcal{D}_{A^u}$. Since $|u(x)| = 1$ pointwise, $\mathcal{U}_u$ is an isometry of $L^2(S_L^1, dx)$ with inverse $\mathcal{U}_{u^{-1}}$, hence unitary; and it preserves $C^\infty(S_L^1, \mathbb{C})$, hence carries the domain of $\mathcal{D}_A$ onto that of $\mathcal{D}_{A^u}$. Therefore $\mathcal{D}_A$ and $\mathcal{D}_{A'} = \mathcal{D}_{A^u}$ are unitarily equivalent, and their spectral multisets coincide. Since A and A' were an arbitrary pair related by a gauge transformation or by an allowed diffeomorphism, the spectral multiset is a well-defined function on S_{geo} , which is the assertion. $\square$

Corollary 6.9 (A strictly coarser spectral completion). *Retain the source architecture of Corollary 6.7: the source is $S_{\text{geo}} = \mathcal{M}_{\text{flat}}(S^1, U(1))$ and the declared Tier-1 readout remains the curvature p_F , which is constant. Adjoin, as the invariant relation, the spectrum of the charged-scalar operator*

$$\mathcal{D}_A := -(\partial_x - i a(x))^2 \quad \text{on} \quad C^\infty(S_L^1, \mathbb{C}) \subseteq L^2(S_L^1, dx), \quad (60)$$

computed with respect to the fixed standard metric dx^2 and measure dx on S_L^1 , with the conventions of Convention 6.3 and the standard self-adjoint domain (the closure of $\mathcal{D}_A$ on $C^\infty(S_L^1, \mathbb{C})$; S_L^1 is closed, so no boundary conditions are needed and the operator is essentially self-adjoint there). The unitary charge-one representation ensures that gauge transformations act unitarily on the charged-scalar Hilbert space, so the spectral multiset is gauge invariant and hence a well-defined function on the declared gauge quotient (37). Its invariance under the auxiliary orientation-preserving diffeomorphism action—which, as noted in Convention 6.1, is treated separately and does not consist of arrows of the fixed-metric groupoid, a general such diffeomorphism not preserving dx —is proved separately in Lemma 6.8(2). Set $J_{\text{spec}}([A]) := \text{spec } \mathcal{D}_A$. Writing $\text{Hol}(A) = e^{i\vartheta}$,

$$\text{spec } \mathcal{D}_A = \left\{ \left(\frac{2\pi k - \vartheta}{L} \right)^2 : k \in \mathbb{Z} \right\} \quad (\text{as a multiset}), \quad (61)$$

and the resulting realized completion is

$$E_{\text{spec}} \cong U(1)/(z \sim z^{-1}) \cong [-1, 1], \quad (62)$$

which is strictly coarser than E_{hol} : the completion map $S_{\text{geo}} \rightarrow E_{\text{spec}}$ is surjective but not injective, its nontrivial fibres being exactly the pairs $\{[A_\vartheta], [A_{-\vartheta}]\}$ with $\vartheta \not\equiv 0, \pi \pmod{2\pi}$.

Proof. The spectrum. By Theorem 6.5 every class has the representative $A_\vartheta = (\vartheta/L) dx$ of (53), and by Example 4.3(ii) the spectrum is unchanged along the gauge orbit, so we may compute with A_ϑ . With $a \equiv \vartheta/L$ constant, $\mathcal{D}_{A_\vartheta}$ has the complete orthonormal system of eigenfunctions $\psi_k(x) = L^{-1/2} e^{2\pi i k x/L}$, $k \in \mathbb{Z}$ (complete in $L^2(S_L^1, dx)$ by Fourier theory). By the convention (42),

$$\left(\partial_x - i \frac{\vartheta}{L} \right) \psi_k = \left(\frac{2\pi i k}{L} - \frac{i\vartheta}{L} \right) \psi_k = i \frac{2\pi k - \vartheta}{L} \psi_k,$$

and hence, squaring and inserting the overall minus sign of (60),

$$\mathcal{D}_{A_\vartheta} \psi_k = - \left(i \frac{2\pi k - \vartheta}{L} \right)^2 \psi_k = \left(\frac{2\pi k - \vartheta}{L} \right)^2 \psi_k.$$

Since the ψ_k are complete, the spectrum is purely discrete and is exactly the multiset (61). Note that ϑ is defined only modulo 2π , and (61) is invariant under $\vartheta \mapsto \vartheta + 2\pi$ via the relabelling $k \mapsto k + 1$; the multiset is therefore a well-defined function of $[A]$, as it must be.

“If”: inversion-related holonomies are spectrally equal. Replacing ϑ by $-\vartheta$ in (61) and relabelling $k \mapsto -k$ carries the multiset to itself, since $(2\pi(-k) - (-\vartheta))^2 = (-(2\pi k - \vartheta))^2 = (2\pi k - \vartheta)^2$.

Hence $[A_\vartheta]$ and $[A_{-\vartheta}]$ —which by Theorem 6.5 are distinct classes whenever $e^{i\vartheta} \neq e^{-i\vartheta}$, i.e. whenever $\vartheta \not\equiv 0, \pi \pmod{2\pi}$ —have the same spectrum.

“Only if”: *equal spectra force inversion-related holonomies.* We prove this directly rather than by citation. Suppose $\text{spec } \mathcal{D}_{A_\vartheta} = \text{spec } \mathcal{D}_{A_{\vartheta'}}$ as multisets. Normalize $\vartheta, \vartheta' \in [0, 2\pi)$, which is permissible by the periodicity just noted. By (61) every eigenvalue of $\mathcal{D}_{A_\vartheta}$ is a nonnegative real number, so we may form the set of both real square roots of the eigenvalues, rescaled by $L/2\pi$:

$$\begin{aligned} \Sigma(\vartheta) &:= \left\{ \pm \frac{L}{2\pi} \sqrt{\lambda} : \lambda \in \text{spec } \mathcal{D}_{A_\vartheta} \right\} \\ &= \left\{ k - \frac{\vartheta}{2\pi} : k \in \mathbb{Z} \right\} \cup \left\{ -k + \frac{\vartheta}{2\pi} : k \in \mathbb{Z} \right\} \\ &= \left(-\frac{\vartheta}{2\pi} + \mathbb{Z} \right) \cup \left(\frac{\vartheta}{2\pi} + \mathbb{Z} \right), \end{aligned}$$

which is determined by the spectral multiset alone. Thus equality of spectra gives

$$\left(\frac{\vartheta}{2\pi} + \mathbb{Z} \right) \cup \left(-\frac{\vartheta}{2\pi} + \mathbb{Z} \right) = \left(\frac{\vartheta'}{2\pi} + \mathbb{Z} \right) \cup \left(-\frac{\vartheta'}{2\pi} + \mathbb{Z} \right)$$

as subsets of $\mathbb{R}$. Intersecting both sides with the interval $[0, 1)$ and comparing, the unordered pair of residues $\{\frac{\vartheta}{2\pi} \bmod 1, -\frac{\vartheta}{2\pi} \bmod 1\}$ equals the pair $\{\frac{\vartheta'}{2\pi} \bmod 1, -\frac{\vartheta'}{2\pi} \bmod 1\}$. Hence $\vartheta' \equiv \vartheta$ or $\vartheta' \equiv -\vartheta \pmod{2\pi}$, that is, $\text{Hol}(A') = \text{Hol}(A)^{\pm 1}$, as claimed.

The completion. By the two directions just proved, the relation J_{spec} has exactly the level sets of the map $U(1) \rightarrow U(1)/(z \sim z^{-1})$. Since the readout p_F is constant and contributes nothing, Definition 5.2 gives $E_{\text{spec}} = \text{Im}(p_F, J_{\text{spec}}) \cong U(1)/(z \sim z^{-1})$. The map $z = e^{i\vartheta} \mapsto \frac{1}{2}(z + z^{-1}) = \cos \vartheta$ is a continuous bijection from this quotient onto $[-1, 1]$ (it is well defined and injective on the quotient, and surjective onto $[-1, 1]$), and it is a homeomorphism since the quotient is compact and $[-1, 1]$ is Hausdorff. This gives (62). Non-injectivity of the completion map and the description of its fibres now follow from Proposition 5.5 applied to the coarser relation family $\{p_F\} \cup \{J_{\text{spec}}\}$. $\square$

Remark 6.10 (The two corollaries together). Corollaries 6.7 and 6.9 are two halves of one lesson, and both are consequences of Theorem 5.3 and Proposition 5.5 applied to the *same* source S_{geo} with the *same* constant readout p_F ; only the declared invariant relation differs. The realized completion is relative to the declared relation family (Remark 5.4). Declare holonomy and the source is reconstructed exactly; declare instead the charged-scalar spectrum and a residual $z \sim z^{-1}$ degeneracy survives the completion—the shadow of the complex conjugation $\phi \mapsto \bar{\phi}$, equivalently of the orientation reversal excluded in Definition 6.2 (Remark 6.6). This is the Paper 2 phenomenon [17], realized geometrically: whether the lost information obstructs the nominated target depends on the target.

7 Variational realization of completed relations

We now embed the completed relation data into dynamics. The claim to be established is narrow and precise: an invariant relation-dependent action yields, by functional variation, correctly normalized stress-energy, gauge-current, and matter responses, and nothing is postulated by naming it.

7.1 Variational setting

Assumption 7.1 (Well-posed variational problem). The global configuration satisfies the descent conditions of Theorem 3.4 and lies in $\mathcal{C}_{\text{geo}}$. Furthermore, *one* of the following holds, and which one is part of the declaration of the model:

- (V1) all variations δg , δA , $\delta \Phi$ have compact support in the interior of M ; or
- (V2) boundary values of the fields, and for gravity the induced metric on ∂M , are held fixed, so that boundary contributions vanish identically; or
- (V3) the gravitational and gauge actions are supplemented by the boundary terms required for a well-posed variational problem. For gravity this is the Gibbons–Hawking–York term

$$\frac{1}{8\pi G_N} \oint_{\partial M} K \sqrt{|h|} \, d^3y$$

[15, 5], with the appropriate subtraction or counterterms in the asymptotic case; for the gauge sector, whatever boundary term the declared $\mathcal{B}$ requires.

All functional derivatives below are asserted to exist on the declared domain; where S_{rel} is nonlocal, this existence is a hypothesis on the model, not a theorem.

Definition 7.2 (Relation-dependent total action). Let $r = (r_\alpha)_{\alpha \in A}$ denote the completed invariant relation coordinates, i.e. the components of e_A^{geo} . Consider

$$S_{\text{tot}}[g, A, \Phi; r] = S_{\text{EH}}[g] + S_{\text{YM}}[g, A] + S_{\text{m}}[g, A, \Phi] + S_{\text{rel}}[g, A, \Phi; r], \quad (63)$$

with

$$S_{\text{EH}}[g] = \frac{1}{16\pi G_N} \int_M (R - 2\Lambda) \, \text{dvol}_g \quad (+ \text{ boundary term under (V3)}), \quad (64)$$

$$S_{\text{YM}}[g, A] = -\frac{1}{4g_{\text{YM}}^2} \int_M \langle F_{\mu\nu}, F^{\mu\nu} \rangle \, \text{dvol}_g, \quad (65)$$

and S_{m} a matter action for Φ [14]. The functional S_{rel} carries the model-specific dependence on r . It may be local, boundary-supported, or nonlocal, provided that

- (R1) it depends on the relation coordinates through a finite or countable dynamical subfamily $A_{\text{dyn}} \subseteq A$, and its Fréchet derivatives with respect to $g^{\mu\nu}$, A_μ , and the independent matter components exist on the stated domain and are represented by distributional kernels, as set out in Section 7.2; and
- (R2) it is invariant under the allowed bundle automorphisms and allowed diffeomorphisms—that is, S_{rel} is constant on the arrows of $\mathfrak{G}_{\text{geo}}$, exactly as the J_α are (Definition 4.1).

Condition (R2) is what makes S_{rel} a functional of the *physical orbit*, and it is what supplies the Ward identities of Section 8. An S_{rel} that fails to be $\mathfrak{G}_{\text{geo}}$ -invariant does not define a geometric realization of a completed relation.

7.2 The functional chain rule for relation functionals

Convention 7.3 (Two regimes for r , and the correct chain rule). By Convention 2.11, exactly one of the following applies to each component r_α , and the model declares which.

(a) r_α is a functional of the fields. Paper 3 permits an arbitrary index set A , and nothing in Sections 5–6 requires more. A differentiable action, however, needs additional analytic structure, and we impose it only where it is used. We therefore require that S_{rel} depend on the relation coordinates through a distinguished *dynamical subfamily*

$$A_{\text{dyn}} \subseteq A, \quad A_{\text{dyn}} \text{ finite or countable}, \quad (66)$$

so that $S_{\text{rel}} = S_{\text{rel}}[g, A, \Phi; (r_\alpha)_{\alpha \in A_{\text{dyn}}}]$ and $\partial S_{\text{rel}} / \partial r_\alpha = 0$ for $\alpha \notin A_{\text{dyn}}$. When A_{dyn} is countably infinite we require in addition that the derivative series in (68) below converge *unconditionally*

in the distributional-dual topology of the relevant space of kernels; when A_{dyn} is finite this is vacuous. The remaining relation coordinates, indexed by $A \setminus A_{\text{dyn}}$, continue to serve as completion coordinates in the sense of Sections 5–6; they simply do not enter the action.

Now suppose $r_\alpha = r_\alpha[g, A, \Phi]$, $\alpha \in A_{\text{dyn}}$, takes values in a Banach (or finite-dimensional) space V_α and is Fréchet differentiable on the declared domain. Then r_α is in general *nonlocal*: a Wilson observable integrates around a loop, an integral invariant integrates over M , a spectral quantity depends on the fields through a global eigenvalue problem. Its variation is accordingly given not by a pointwise partial derivative but by the pairing of distributional kernels with the variations:

$$\delta r_\alpha = \int_M \left[\frac{\delta r_\alpha}{\delta g^{\mu\nu}(y)} \delta g^{\mu\nu}(y) + \left\langle \frac{\delta r_\alpha}{\delta A_\mu(y)}, \delta A_\mu(y) \right\rangle + \left\langle \left\langle \frac{\delta r_\alpha}{\delta \Phi(y)}, \delta \Phi(y) \right\rangle \right\rangle \right] \sqrt{-g(y)} \, d^4 y, \quad (67)$$

where $\delta r_\alpha / \delta g^{\mu\nu}(y)$ etc. denote the V_α -valued distributional kernels representing the derivative, $\langle \cdot, \cdot \rangle$ is the Ad-invariant pairing on $\mathfrak{g}$, and $\langle \cdot, \cdot \rangle$ is the duality pairing between W^* and W discussed in Section 7.3. (For a Wilson loop r_α the kernel $\delta r_\alpha / \delta A_\mu(y)$ is supported on the loop and is a distribution, not a function; (67) is exactly the statement that must be used, and a pointwise $\partial r_\alpha / \partial A_\mu(y)$ does not exist.)

Then, writing $\partial S_{\text{rel}} / \partial r_\alpha \in V_\alpha^*$ for the derivative of S_{rel} with respect to its V_α -argument (a bounded linear functional, and an ordinary partial derivative only when V_α is finite-dimensional), the chain rule reads

$$\frac{\delta S_{\text{rel}}}{\delta g^{\mu\nu}(y)} = \left. \frac{\delta S_{\text{rel}}}{\delta g^{\mu\nu}(y)} \right|_{r \text{ fixed}} + \sum_{\alpha \in A_{\text{dyn}}} \left\langle \frac{\partial S_{\text{rel}}}{\partial r_\alpha}, \frac{\delta r_\alpha}{\delta g^{\mu\nu}(y)} \right\rangle_{V_\alpha^*, V_\alpha}, \quad (68)$$

and identically with $g^{\mu\nu}(y)$ replaced by $A_\mu(y)$ or by the matter components. The sum is over the dynamical subfamily (66), and converges by the unconditional-convergence hypothesis imposed there. Only in the special case where $V_\alpha = \mathbb{R}$ and r_α is a *local* scalar field does (68) reduce to an ordinary pointwise partial derivative; in general it does not.

(b) r_α labels a fixed topological or boundary sector. Then $\delta r_\alpha = 0$ under all variations within the sector—this is what “fixed sector” means—and r_α enters the field equations only as a parameter. The second term of (68) is absent.

In neither case is r_α varied as an independent dynamical field. A model wishing to promote r_α to a field must adjoin it to the configuration tuple, enlarge $\mathcal{C}_{\text{geo}}$, and supply its own kinetic term; that is a different model, and its Euler–Lagrange equation is an additional equation not derived here.

7.3 Matter pairings and independent components

Convention 7.4 (Matter variation and the dual pairing). The representation $\rho: G \rightarrow \text{GL}(W)$ is not assumed unitary (Convention 2.1), and W may be complex. Accordingly the matter Euler operator is valued in the dual, and the variational pairing must be specified. We adopt the following, and use it consistently.

(i) *Real independent components.* Suppose Φ denotes a collection of *real* independent field components. This is always achievable: decompose a complex field into its real and imaginary parts, or a section of a complex bundle into its underlying real bundle. Then $\delta S / \delta \Phi$ is valued in the real dual bundle $P \times_{\rho^*} W^*$, and the pairing

$$\langle \cdot, \cdot \rangle: W^* \times W \longrightarrow \mathbb{R}$$

is the canonical duality pairing. This pairing is G -invariant for the contragredient action $\rho^*(h) = \rho(h^{-1})^\top$, and requires no metric on W . There is then a single Euler equation $\mathcal{E}_\Phi = 0$.

(ii) *Complex fields treated holomorphically.* If instead Φ is treated as a complex field with Φ and $\bar{\Phi}$ varied independently—the usual convention for a charged scalar or a Dirac spinor—then there are *two* Euler equations,

$$\mathcal{E}_\Phi := \frac{1}{\sqrt{-g}} \frac{\delta S}{\delta \bar{\Phi}} = 0 \quad \text{and} \quad \overline{\mathcal{E}_\Phi} := \frac{1}{\sqrt{-g}} \frac{\delta S}{\delta \Phi} = 0, \quad (69)$$

which are conjugate to one another when the action is real, and the pairing used in the variation is the sesquilinear one furnished by a ρ -invariant Hermitian form on W —which exists precisely when ρ is unitary. If ρ is not unitary, no such Hermitian form is available and one must revert to convention (i).

For definiteness, and to avoid carrying two conjugate equations through every formula, *in the statements below Φ denotes real independent field components in the sense of (i)*, and $\mathcal{E}_\Phi$ is the corresponding single dual-valued Euler operator. Every result restates verbatim in convention (ii) upon replacing $\mathcal{E}_\Phi = 0$ by the conjugate pair (69) and $\langle\langle \cdot, \cdot \rangle\rangle$ by the invariant Hermitian pairing.

7.4 Response tensors, currents, and Euler operators

Definition 7.5 (Responses). For each nongravitational sector $X \in \{\text{YM}, \text{m}, \text{rel}\}$ define the stress-energy tensor

$$T_{\mu\nu}^X := -\frac{2}{\sqrt{-g}} \frac{\delta S_X}{\delta g^{\mu\nu}}, \quad (70)$$

the gauge current

$$J_{\text{m}}^\mu := -\frac{1}{\sqrt{-g}} \frac{\delta S_{\text{m}}}{\delta A_\mu}, \quad J_{\text{rel}}^\mu := -\frac{1}{\sqrt{-g}} \frac{\delta S_{\text{rel}}}{\delta A_\mu}, \quad (71)$$

each valued in $\mathfrak{g}$ (with $\mathfrak{g}$ -indices contracted using $\langle \cdot, \cdot \rangle$), and the matter Euler operators

$$\mathcal{E}_\Phi^X := \frac{1}{\sqrt{-g}} \frac{\delta S_X}{\delta \Phi}, \quad (72)$$

valued in the dual bundle and paired with $\delta\Phi$ by $\langle\langle \cdot, \cdot \rangle\rangle$ as in Convention 7.4(i). For $X = \text{rel}$ the functional derivatives in (70)–(72) are those computed by the chain rule (68), and thus already include the relation-mediated contributions.

Lemma 7.6 (Yang–Mills stress tensor). *With the normalization (65),*

$$T_{\mu\nu}^{\text{YM}} = \frac{1}{g_{\text{YM}}^2} \left(\langle F_{\mu\rho}, F_{\nu}{}^\rho \rangle - \frac{1}{4} g_{\mu\nu} \langle F_{\rho\sigma}, F^{\rho\sigma} \rangle \right), \quad (73)$$

which is symmetric, gauge invariant, and (in four dimensions) trace free.

Proof. Write $S_{\text{YM}} = -\frac{1}{4g_{\text{YM}}^2} \int_M \langle F_{\mu\rho}, F_{\nu\sigma} \rangle g^{\mu\nu} g^{\rho\sigma} \sqrt{-g} d^4x$; the components $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + [A_\mu, A_\nu]$ are independent of g . Using $\delta\sqrt{-g} = -\frac{1}{2}\sqrt{-g} g_{\mu\nu} \delta g^{\mu\nu}$ and varying the two inverse metrics,

$$\delta S_{\text{YM}} = -\frac{1}{4g_{\text{YM}}^2} \int_M \left(2\langle F_{\mu\rho}, F_{\nu}{}^\rho \rangle - \frac{1}{2} g_{\mu\nu} \langle F_{\rho\sigma}, F^{\rho\sigma} \rangle \right) \delta g^{\mu\nu} \sqrt{-g} d^4x,$$

and (70) gives (73). Gauge invariance follows from Ad-invariance of $\langle \cdot, \cdot \rangle$ and (23); tracelessness from $g^{\mu\nu} \langle F_{\mu\rho}, F_{\nu}{}^\rho \rangle = \langle F_{\rho\sigma}, F^{\rho\sigma} \rangle$ and $g^{\mu\nu} g_{\mu\nu} = 4$. $\square$

Lemma 7.7 (Yang–Mills variation with respect to the connection). *Under Assumption 7.1,*

$$\delta S_{\text{YM}} = -\frac{1}{g_{\text{YM}}^2} \int_M \langle F^{\mu\nu}, D_\mu \delta A_\nu \rangle \text{dvol}_g = \frac{1}{g_{\text{YM}}^2} \int_M \langle D_\mu F^{\mu\nu}, \delta A_\nu \rangle \text{dvol}_g, \quad (74)$$

up to the treated boundary term.

Proof. From $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu + [A_\mu, A_\nu]$ one has

$$\delta F_{\mu\nu} = D_\mu \delta A_\nu - D_\nu \delta A_\mu. \quad (75)$$

Hence, from (65),

$$\delta S_{\text{YM}} = -\frac{1}{4g_{\text{YM}}^2} \int_M 2 \langle F^{\mu\nu}, \delta F_{\mu\nu} \rangle \text{dvol}_g = -\frac{1}{2g_{\text{YM}}^2} \int_M \langle F^{\mu\nu}, D_\mu \delta A_\nu - D_\nu \delta A_\mu \rangle \text{dvol}_g.$$

The two terms in the bracket are *equal* after relabelling, by the antisymmetry of $F^{\mu\nu}$: swapping the dummy indices $\mu \leftrightarrow \nu$ in the second term gives $-\langle F^{\nu\mu}, D_\mu \delta A_\nu \rangle = +\langle F^{\mu\nu}, D_\mu \delta A_\nu \rangle$. Therefore the bracket contributes twice the first term, the factor $\frac{1}{2}$ cancels, and

$$\delta S_{\text{YM}} = -\frac{1}{g_{\text{YM}}^2} \int_M \langle F^{\mu\nu}, D_\mu \delta A_\nu \rangle \text{dvol}_g,$$

which is the first equality in (74). Covariant integration by parts—legitimate because $\langle \cdot, \cdot \rangle$ is Ad-invariant, hence D -compatible, and because the boundary term vanishes by Assumption 7.1—moves D_μ onto $F^{\mu\nu}$ and flips the sign, giving the second equality. $\square$

7.5 The coupled field equations

Theorem 7.8 (Coupled field equations). *Under Assumption 7.1 and Conventions 7.3 and 7.4(i), stationarity of S_{tot} under independent variations of $g^{\mu\nu}$, A_μ , and Φ is equivalent to the system*

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi G_N (T_{\mu\nu}^{\text{YM}} + T_{\mu\nu}^{\text{m}} + T_{\mu\nu}^{\text{rel}}), \quad (76)$$

$$\frac{1}{g_{\text{YM}}^2} D_\mu F^{\mu\nu} = J_{\text{m}}^\nu + J_{\text{rel}}^\nu, \quad (77)$$

$$\mathcal{E}_\Phi^{\text{m}} + \mathcal{E}_\Phi^{\text{rel}} = 0, \quad (78)$$

with representation indices and the pairings $\langle \cdot, \cdot \rangle$ and $\langle\langle \cdot, \cdot \rangle\rangle$ restored as required by G and ρ .

Proof. Metric variation. The standard variation of (64), after the boundary treatment stipulated in Assumption 7.1 (which is exactly what disposes of the total-derivative term $\int_M g^{\mu\nu} \delta R_{\mu\nu} \sqrt{-g}$; under (V3) it is the Gibbons–Hawking–York term that cancels it [15, 5]), gives

$$\delta S_{\text{EH}} = \frac{1}{16\pi G_N} \int_M (G_{\mu\nu} + \Lambda g_{\mu\nu}) \delta g^{\mu\nu} \text{dvol}_g. \quad (79)$$

By Definition 7.5,

$$\delta(S_{\text{YM}} + S_{\text{m}} + S_{\text{rel}}) \Big|_{\delta g} = -\frac{1}{2} \int_M (T_{\mu\nu}^{\text{YM}} + T_{\mu\nu}^{\text{m}} + T_{\mu\nu}^{\text{rel}}) \delta g^{\mu\nu} \text{dvol}_g, \quad (80)$$

where δS_{rel} is computed by the chain rule (68): in regime (a) of Convention 7.3 it includes the relation-mediated kernel contributions, and in regime (b) those are absent. In either case (80) defines its left side via (70) and adds no assumption. Setting $\delta S_{\text{tot}} = 0$ for arbitrary symmetric $\delta g^{\mu\nu}$ (compactly supported, or boundary-vanishing) and applying the fundamental lemma of the calculus of variations gives

$$\frac{1}{16\pi G_N} (G_{\mu\nu} + \Lambda g_{\mu\nu}) - \frac{1}{2} (T_{\mu\nu}^{\text{YM}} + T_{\mu\nu}^{\text{m}} + T_{\mu\nu}^{\text{rel}}) = 0,$$

which is (76) with the displayed $8\pi G_N$ normalization.

Gauge variation. By Lemma 7.7,

$$\delta S_{\text{YM}}|_{\delta A} = \frac{1}{g_{\text{YM}}^2} \int_M \langle D_\mu F^{\mu\nu}, \delta A_\nu \rangle \text{dvol}_g,$$

while by the definitions (71),

$$\delta(S_{\text{m}} + S_{\text{rel}})|_{\delta A} = - \int_M \langle J_{\text{m}}^\nu + J_{\text{rel}}^\nu, \delta A_\nu \rangle \text{dvol}_g.$$

Setting the sum of the two to zero for arbitrary δA_ν , and using the nondegeneracy of $\langle \cdot, \cdot \rangle$, gives (77).

Matter variation. S_{EH} and S_{YM} are independent of Φ , so by (72) $\delta_\Phi S_{\text{tot}} = \int_M \langle \mathcal{E}_\Phi^{\text{m}} + \mathcal{E}_\Phi^{\text{rel}}, \delta\Phi \rangle \text{dvol}_g$; arbitrariness of the real independent components $\delta\Phi$ and nondegeneracy of the duality pairing give (78). Under Convention 7.4(ii) this is replaced by the conjugate pair (69).

The three variations are independent, so stationarity of S_{tot} is equivalent to the simultaneous validity of (76)–(78). $\square$

Remark 7.9 (What has and has not been derived). The tensor $T_{\mu\nu}^{\text{rel}}$ is not an independent postulate. By (70) and the chain rule (68) it is the metric functional derivative of a specified $\mathfrak{G}_{\text{geo}}$ -invariant action satisfying (R1)–(R2). Different source completions, and different couplings of the same completion, produce different $T_{\mu\nu}^{\text{rel}}$, J_{rel}^μ , $\mathcal{E}_\Phi^{\text{rel}}$. Theorem 7.8 asserts only the *structure* of the response—that a completed relation enters Tier-1 dynamics through functional variation of an invariant action, with the normalizations displayed—and asserts nothing about the magnitude, sign, or empirical existence of any such term. No new interaction is proposed, and no empirical claim is made.

8 Ward–Noether identities and covariant consistency

Gauge and diffeomorphism invariance impose Noether identities [9], and those identities constrain every effective representation of a completed source relation. Two logically distinct routes to covariant conservation must be kept apart, and we present them separately.

Convention 8.1 (Infinitesimal allowed arrows and the induced field variations). A base vector field ξ on M does not by itself define a Lie derivative of a principal connection or of a section of an associated bundle: differentiating such objects requires a lift of the flow of ξ to the total space P . We therefore generate the variations from the groupoid itself.

Let $(\hat{F}_t, f_t)_{t \in (-\epsilon, \epsilon)}$ be a smooth one-parameter family of allowed arrows of $\mathfrak{G}_{\text{geo}}$ with $(\hat{F}_0, f_0) = (\text{id}_P, \text{id}_M)$, and let

$$\xi := \left. \frac{d}{dt} \right|_{t=0} f_t$$

be the generating base vector field, taken compactly supported in the interior of M . Write $\hat{\xi}$ for the induced G -invariant vector field on P generating $\hat{F}_t$, which projects to ξ . The induced infinitesimal variations of the fields are defined by

$$\begin{aligned} \delta_\xi g^{\mu\nu} &:= - \left. \frac{d}{dt} \right|_{t=0} (f_t^* g^{\mu\nu}) = -\mathcal{L}_\xi g^{\mu\nu}, \\ \delta_\xi A &:= - \left. \frac{d}{dt} \right|_{t=0} (\hat{F}_t^* A), \\ \delta_\xi \Phi &:= - \left. \frac{d}{dt} \right|_{t=0} (\hat{F}_{t,\rho}^{-1} \circ \Phi \circ f_t). \end{aligned} \tag{81}$$

The first is intrinsic to the base, and we record the inverse-metric identity we shall use: differentiating $g^{\mu\rho} g_{\rho\nu} = \delta_\nu^\mu$ gives

$$\mathcal{L}_\xi g^{\mu\nu} = -(\nabla^\mu \xi^\nu + \nabla^\nu \xi^\mu), \tag{82}$$

the sign being opposite to that for the covariant metric, $\mathcal{L}_\xi g_{\mu\nu} = \nabla_\mu \xi_\nu + \nabla_\nu \xi_\mu$. Hence, with the convention (81),

$$\delta_\xi g^{\mu\nu} = \nabla^\mu \xi^\nu + \nabla^\nu \xi^\mu. \quad (83)$$

The bundle variations require a further remark. The lift $\hat{\xi}$ of ξ is not unique: any two allowed lifts differ by a vertical, i.e. gauge, direction. Concretely, in a local trivialization over $U \subseteq M$, in which A is represented by A_μ and Φ by a W -valued map, a choice of lift yields

$$\delta_\xi A_\mu = -(\mathcal{L}_\xi A)_\mu + D_\mu \varepsilon, \quad \delta_\xi \Phi = -\mathcal{L}_\xi \Phi - \rho_*(\varepsilon) \Phi, \quad (84)$$

where $\mathcal{L}_\xi$ denotes the coordinate Lie derivative of the local representatives and $\varepsilon \in \Gamma(\text{ad}P|_U)$ is the vertical gauge parameter recording the choice of lift; changing the lift changes ε . Thus the induced connection variation agrees with the naive coordinate Lie derivative only up to a vertical term $D_\mu \varepsilon$, with the corresponding representation action $-\rho_*(\varepsilon)\Phi$ on the matter field.

This ambiguity is harmless for the Ward identity of Theorem 8.4, and we note why once here. The extra terms in (84) are precisely an infinitesimal gauge transformation, of the form (92). In the Ward computation they are contracted with the Euler coefficients of A and Φ , which vanish identically on the gauge and matter shell on which the theorem is evaluated. The vertical terms therefore multiply zero and drop out, and the conclusion is independent of the choice of lift.

Definition 8.2 (Total stress-energy).

$$T_{\mu\nu}^{\text{tot}} := T_{\mu\nu}^{\text{YM}} + T_{\mu\nu}^{\text{m}} + T_{\mu\nu}^{\text{rel}}. \quad (85)$$

8.1 Route (a): on a full solution, via the contracted Bianchi identity

Proposition 8.3 (Conservation on a full solution). *Suppose the Einstein equation (76) holds, with Λ constant. Then*

$$\nabla^\mu T_{\mu\nu}^{\text{tot}} = 0. \quad (86)$$

Proof. Take the covariant divergence of (76). The contracted second Bianchi identity gives $\nabla^\mu G_{\mu\nu} = 0$ [12], and metric compatibility (25) with constancy of Λ gives $\nabla^\mu (\Lambda g_{\mu\nu}) = 0$. Hence $8\pi G_N \nabla^\mu T_{\mu\nu}^{\text{tot}} = 0$, and (86) follows since $G_N \neq 0$. $\square$

This argument uses the Einstein equation: it is valid on a full solution of the coupled system, including the gravitational field equation, and it does not establish (86) from the gauge and matter equations alone. Route (b) below derives the same conclusion without the Einstein equation, as a Noether identity of diffeomorphism invariance.

8.2 Route (b): diffeomorphism invariance, independent of the Einstein equation

Theorem 8.4 (Diffeomorphism invariance and on-shell covariant conservation). *Let $S_{\text{ng}} := S_{\text{YM}} + S_{\text{m}} + S_{\text{rel}}$ denote the nongravitational part of the action (63), and assume:*

- (C1) (Invariance.) S_{ng} is invariant under every smooth one-parameter family of allowed arrows $(\hat{F}_t, f_t)$ of $\mathfrak{G}_{\text{geo}}$ with compactly supported generator ξ , the fields (g, A, Φ) transforming as in (81). This is the hypothesis (R2) of Definition 7.2, restricted to arrows near the identity;
- (C2) (Invariance of the fixed data.) all data held fixed during the variation are themselves invariant under that flow. Concretely: any background or external structure in $\mathcal{B}$ entering S_{ng} is either dragged along with the fields or is preserved by the flow; and every relation coordinate r_α in regime (b) of Convention 7.3 (a fixed topological or boundary sector label) satisfies $\delta_\xi r_\alpha = 0$, which holds because such labels are $\mathfrak{G}_{\text{geo}}$ -invariant by construction and ξ generates arrows of $\mathfrak{G}_{\text{geo}}$;

- (C3) (Boundary terms.) ξ has compact support in the interior of M , or the boundary terms generated by the integrations by parts below vanish under the treatment declared in Assumption 7.1;
- (C4) (Full Euler coefficients.) the coefficients of $\delta_\xi A_\mu$ and $\delta_\xi \Phi$ in $\delta_\xi S_{\text{ng}}$ are the full Euler expressions appearing on the left sides of (77) and (78)—that is, they include the YM, matter, and relation contributions, computed with the chain rule (68);
- (C5) (Availability of lifts.) for every smooth vector field ξ with compact support in the interior of M , the admissibility class and allowed-arrow structure of $\mathfrak{G}_{\text{geo}}$ contain a smooth allowed one-parameter family of arrows $(\hat{F}_t, f_t)$ whose base generator is ξ , with the fields (g, A, Φ) and the fixed data of $\mathcal{B}$ transformed as in (81). Thus every such ξ generates an allowed flow to which (C1) and (C2) apply.

Write the full Euler operators of the gauge and matter sectors as

$$\mathcal{G}^\nu := \frac{1}{g_{\text{YM}}^2} D_\mu F^{\mu\nu} - (J_{\text{m}}^\nu + J_{\text{rel}}^\nu), \quad \mathcal{E}_\Phi := \mathcal{E}_\Phi^{\text{m}} + \mathcal{E}_\Phi^{\text{rel}}, \quad (87)$$

so that the gauge equation (77) reads $\mathcal{G}^\nu = 0$ and the matter equation (78) reads $\mathcal{E}_\Phi = 0$.

Then diffeomorphism invariance gives the integrated variational identity displayed in the proof. On every configuration satisfying $\mathcal{G}^\nu = 0$ and $\mathcal{E}_\Phi = 0$, but not necessarily the Einstein equation, one has

$$\nabla^\mu T_{\mu\nu}^{\text{tot}} = 0. \quad (88)$$

If the sectors exchange energy–momentum, the individual divergences $\nabla^\mu T_{\mu\nu}^X$ need not vanish; only the sum does.

Proof. By (C1) and (C2), $\delta_\xi S_{\text{ng}} = 0$ identically in the fields—this is the off-shell content of diffeomorphism invariance, and (C2) is exactly what licenses it, since a non-invariant fixed datum would contribute an uncanceled term. Expanding the total variation in the three field arguments, with the full Euler coefficients of (C4),

$$0 = \delta_\xi S_{\text{ng}} = \int_M \left[\frac{\delta S_{\text{ng}}}{\delta g^{\mu\nu}} \delta_\xi g^{\mu\nu} + \left\langle \frac{\delta S_{\text{ng}}}{\delta A_\mu}, \delta_\xi A_\mu \right\rangle + \left\langle \left\langle \frac{\delta S_{\text{ng}}}{\delta \Phi}, \delta_\xi \Phi \right\rangle \right\rangle \right] d^4x, \quad (89)$$

We first identify the three coefficients in (89). By (71), (72) and Lemma 7.7, the coefficient of $\delta_\xi A_\mu$ is $\sqrt{-g}$ times $\mathcal{G}^\mu$, and the coefficient of $\delta_\xi \Phi$ is $\sqrt{-g}$ times $\mathcal{E}_\Phi$, in the notation (87). The coefficient of $\delta_\xi g^{\mu\nu}$ is, by (70) and (85), equal to $-\frac{1}{2}\sqrt{-g} T_{\mu\nu}^{\text{tot}}$. Hence (89) reads

$$0 = -\frac{1}{2} \int_M T_{\mu\nu}^{\text{tot}} \delta_\xi g^{\mu\nu} \text{dvol}_g + \int_M \left[\langle \mathcal{G}^\mu, \delta_\xi A_\mu \rangle + \langle \langle \mathcal{E}_\Phi, \delta_\xi \Phi \rangle \rangle \right] \text{dvol}_g. \quad (90)$$

Equation (90) is the integrated variational identity referred to in the statement. It holds for every compactly supported ξ generating a family of allowed arrows, and we now impose the gauge and matter equations on it.

Accordingly, let the configuration satisfy $\mathcal{G}^\mu = 0$ and $\mathcal{E}_\Phi = 0$. The second integral in (90) then vanishes term by term, whatever the variations $\delta_\xi A_\mu$ and $\delta_\xi \Phi$ happen to be. This is also the step announced in Convention 8.1: whatever lift $\hat{\xi}$ of ξ was chosen, and hence whatever vertical terms $D_\mu \varepsilon$ and $-\rho_*(\varepsilon)\Phi$ appear in (84), they enter only through $\delta_\xi A_\mu$ and $\delta_\xi \Phi$, which are here contracted with the vanishing $\mathcal{G}^\mu$ and $\mathcal{E}_\Phi$; the conclusion is therefore independent of the choice of lift. Hence (90) reduces to

$$0 = -\frac{1}{2} \int_M T_{\mu\nu}^{\text{tot}} \delta_\xi g^{\mu\nu} \text{dvol}_g.$$

Insert $\delta_\xi g^{\mu\nu} = \nabla^\mu \xi^\nu + \nabla^\nu \xi^\mu$ from (83), whose sign is fixed by (82), and use the symmetry of $T_{\mu\nu}^{\text{tot}}$:

$$0 = - \int_M T_{\mu\nu}^{\text{tot}} \nabla^\mu \xi^\nu \, \text{dvol}_g = \int_M (\nabla^\mu T_{\mu\nu}^{\text{tot}}) \xi^\nu \, \text{dvol}_g,$$

the last step by covariant integration by parts, whose boundary term vanishes by (C3). By (C5), every smooth compactly supported vector field ξ in the interior of M arises as the base generator of an allowed flow, so the identity above holds for all such ξ ; since ξ is thereby arbitrary among compactly supported interior fields, the fundamental lemma of the calculus of variations gives (88).

This argument nowhere uses the Einstein equation: (88) follows from diffeomorphism invariance [9] together with the gauge and matter equations alone. $\square$

8.3 The gauge Ward identity

Proposition 8.5 (Gauge-current identity). *Assume $S_{\text{m}} + S_{\text{rel}}$ is invariant under the vertical (gauge) subgroup $\text{Gau}(P)$; that any relation coordinate in regime (b) is gauge invariant (as it is, by construction); and that Assumption 7.1 holds. Then on solutions of the matter equation (78),*

$$D_\mu (J_{\text{m}}^\mu + J_{\text{rel}}^\mu) = 0. \quad (91)$$

Proof. Let $\epsilon \in \Gamma(\text{ad}P)$ be a compactly supported infinitesimal gauge parameter, acting by

$$\delta_\epsilon A_\mu = D_\mu \epsilon, \quad \delta_\epsilon \Phi = -\rho_*(\epsilon) \Phi, \quad (92)$$

where $\rho_*: \mathfrak{g} \rightarrow \mathfrak{gl}(W)$ is the induced Lie-algebra representation and $\delta_\epsilon \Phi$ is valued in W (no inner product on W is used). Gauge invariance gives $\delta_\epsilon (S_{\text{m}} + S_{\text{rel}}) = 0$ identically:

$$0 = \int_M \left(\left\langle \frac{\delta(S_{\text{m}} + S_{\text{rel}})}{\delta A_\mu}, D_\mu \epsilon \right\rangle + \left\langle \frac{\delta(S_{\text{m}} + S_{\text{rel}})}{\delta \Phi}, -\rho_*(\epsilon) \Phi \right\rangle \right) \text{d}^4x,$$

the second pairing being the duality pairing of Convention 7.4(i) between the dual-valued Euler operator and the W -valued variation. On the matter equation (78) the second term vanishes. Using (71), the first term is

$$0 = - \int_M \langle J_{\text{m}}^\mu + J_{\text{rel}}^\mu, D_\mu \epsilon \rangle \, \text{dvol}_g.$$

Covariant integration by parts—valid because $\langle \cdot, \cdot \rangle$ is Ad -invariant, hence D -compatible, and because the boundary term vanishes by Assumption 7.1—gives

$$0 = \int_M \langle D_\mu (J_{\text{m}}^\mu + J_{\text{rel}}^\mu), \epsilon \rangle \, \text{dvol}_g.$$

Since ϵ is arbitrary and $\langle \cdot, \cdot \rangle$ is nondegenerate, (91) follows. $\square$

Remark 8.6 (Consistency of the gauge equation with the current identity). Identity (91) is consistent with, and indeed implied on shell by, the gauge equation (77). Applying D_ν to (77) requires evaluating $D_\nu D_\mu F^{\mu\nu}$. Decomposing the second covariant derivative into its symmetric and anti-symmetric parts in (μ, ν) and contracting with the antisymmetric $F^{\mu\nu}$, only the antisymmetric part survives:

$$D_\nu D_\mu F^{\mu\nu} = \frac{1}{2} [D_\nu, D_\mu] F^{\mu\nu} = \frac{1}{2} [F_{\nu\mu}, F^{\mu\nu}],$$

by the Ricci identity for the gauge-covariant derivative on a section of $\text{ad}P$, namely $[D_\nu, D_\mu]X = [F_{\nu\mu}, X]$. (The Levi-Civita curvature terms cancel between the two spacetime indices of the antisymmetric $F^{\mu\nu}$, by the first Bianchi identity for R .) The remaining bracket vanishes: writing

$F_{\nu\mu} = -F_{\mu\nu}$ gives $[F_{\nu\mu}, F^{\mu\nu}] = -[F_{\mu\nu}, F^{\mu\nu}]$, and $[X, X] = 0$ for the $\mathfrak{g}$ -valued quantity $X = F_{\mu\nu}$ contracted with its own index-raised counterpart. Hence $D_\nu D_\mu F^{\mu\nu} = 0$, and (77) therefore forces $D_\mu (J_m^\mu + J_{\text{rel}}^\mu) = 0$ after relabelling the free index, in agreement with Proposition 8.5. The two derivations are independent: the proposition uses only gauge invariance and the matter equation, whereas this computation uses the gauge equation.

Remark 8.7 (Ward identities as a constraint on realizations). Theorem 8.4 and Proposition 8.5 function as selection rules. Every effective representation of a completed source relation must satisfy them, because they follow from nothing more than the invariance requirement (R2) that made S_{rel} a functional of the physical orbit in the first place. Consequently, a proposed correction to the Tier-1 field equations that violates the relevant Ward identity is *not* a geometric realization of an invariant completion. This is a sharp, checkable necessary condition on any candidate coupling of completed relations to Tier-1 dynamics.

9 From completion to Tier-1 representation

The completed geometric architecture established by Theorems 3.4, 5.3 and 7.8 is the factorization

$$S_{\text{geo}} \xrightarrow{e_A^{\text{geo}}} E_A^{\text{geo}} \xrightarrow{u_A} T, \quad u_A \circ e_A^{\text{geo}} = p_{\text{geo}}. \quad (93)$$

The first map retains the readout together with the selected invariant geometric relations; the second forgets those completion coordinates and returns the original Tier-1 state. Neither map chooses a coordinate chart, a gauge potential, or a preferred representative of a physical orbit: S_{geo} is already the orbit space, and e_A^{geo} is built from $\mathfrak{G}_{\text{geo}}$ -invariant functions. This is what it means for the completion to be geometric rather than notational.

When a relation-dependent action is specified in the sense of Definition 7.2, the completed state enters Tier-1 dynamics through the derived response fields

$$T_{\mu\nu}^{\text{rel}}, \quad J_{\text{rel}}^\mu, \quad \mathcal{E}_\Phi^{\text{rel}}. \quad (94)$$

A distinction must be drawn here, since the two halves of the construction carry invariance in different senses, and conflating them would misdescribe the responses.

Invariance. What is invariant are the *scalars* and the *orbit-defined information*: the action S_{rel} is a $\mathfrak{G}_{\text{geo}}$ -invariant functional by hypothesis (R2); the relation coordinates r_α and the readout p_{geo} are functions on the orbit space S_{geo} by Definition 4.1; and the completion (93) is assembled from these.

Covariance. The responses (94), by contrast, are *not* invariant pointwise, and are not scalars. They are the functional derivatives of an invariant action with respect to the fields, hence tensorial objects on M carrying spacetime and representation indices: $T_{\mu\nu}^{\text{rel}}$ is a symmetric 2-tensor, J_{rel}^μ an $\text{ad}P$ -valued spacetime vector field, and $\mathcal{E}_\Phi^{\text{rel}}$ a section of the dual matter bundle. Under an arrow $(\hat{F}, f)$ of $\mathfrak{G}_{\text{geo}}$ they transform *covariantly*—by pushforward of tensor indices along f and by the adjoint or dual representation along $\hat{F}$ —rather than being fixed. Thus, for instance, $J_{\text{rel}}^\mu \mapsto \text{Ad}_{h^{-1}} J_{\text{rel}}^\mu$ under a gauge transformation h , and $T_{\mu\nu}^{\text{rel}} \mapsto (f^{-1})^* T_{\mu\nu}^{\text{rel}}$ under an allowed diffeomorphism f .

The correct statement is therefore the following. What descends to the orbit space is the invariant action and the relation data; what the variation produces is a family of covariant response tensors, one per gauge and coordinate representative, transforming into one another under the arrows of $\mathfrak{G}_{\text{geo}}$. It is precisely this covariance—not pointwise invariance—that makes the field equations (76)–(78) well posed as equations between tensors, and that makes the Ward identities of Section 8 meaningful: $\nabla^\mu T_{\mu\nu}^{\text{tot}} = 0$ is a covariant equation, and is preserved under the arrows of $\mathfrak{G}_{\text{geo}}$ rather than being invariant under them in the sense of Definition 4.1. By Theorem 7.8 and Section 8, the responses (94) are correctly normalized, Ward-consistent, covariant objects obtained by functional variation of a $\mathfrak{G}_{\text{geo}}$ -invariant action.

This is the whole of what the present paper establishes about dynamics, and the boundary of the argument should be stated precisely. Paper 4 constructs covariant response fields from $\mathfrak{G}_{\text{geo}}$ -invariant actions. It does not determine whether the resulting Tier-1 evolution closes autonomously on T , which completion coordinates (if any) must be retained for a closed description, what happens when unresolved variables are eliminated, or whether such elimination produces memory or other history-dependent effective behaviour. In particular, no inference of Tier-1 autonomy may be drawn merely from the observation that some response functional happens to descend to a function of t at the kinematic level considered here: descent of a response at one level does not establish closure of an evolution, and the two questions are not the same. Autonomous closure, retained variables, elimination, memory, and projected evolution require the separate analysis of Paper 5, and are not treated here.

10 Conclusion

Canonical source completion acquires physical meaning only on a domain where its relation coordinates are jointly realized by compatible geometric data. Four results establish where that domain lies and what it supports.

The geometric descent theorem (Theorem 3.4) constructs the global metric, principal bundle, connection, and matter section from compatible local data, uniquely up to bundle isomorphism, with the global time-orientation condition treated explicitly (Remark 3.6); its corollary (Corollary 3.7) exhibits the realizability boundary, since relation values admitting no compatible descent data are simply outside the geometric realization domain. Passage to the orbit space of the bundle-automorphism groupoid (Definition 2.9) removes coordinate and gauge redundancy *before* completion, so that no gauge representative is taken for a physical state; the isotropy of a configuration is its stabilizer, which is in general a proper subgroup of the ambient automorphism group (Remark 2.7). The restricted geometric realization theorem (Theorem 5.3) then identifies the joint image of the readout and the declared invariant relations with the canonical minimal completion of Paper 3—relative to that family, and only on that domain—while the orbit reconstruction criterion (Proposition 5.5) reduces exact reconstruction to a separation question in invariant theory, answered case by case and never by fiat.

The flat-circle classification (Theorem 6.5 and Corollaries 6.7–6.9) shows in a fully solvable case that a global relation can completely parametrize a reduced source even when local curvature vanishes identically, and that a coarser declared relation leaves a strictly coarser completion behind. Finally, invariant relation-dependent actions (Theorem 7.8) convert completed structure into correctly normalized Einstein, Yang–Mills, and matter equations, with the associated covariant response fields derived rather than postulated, and the Ward–Noether identities (Proposition 8.3, Theorem 8.4, Proposition 8.5) enforce covariant conservation and supply a sharp necessary condition on any candidate realization.

This supplies the geometric bridge from abstract completion to observable Tier-1 structure. It does not assert that every abstract completion is geometrically realizable, that any finite family of local invariants separates Lorentzian geometries, or anything about the physical status of our universe. The dynamical questions of closure, retention, elimination, and memory belong to Paper 5.

References

- [1] Y. Aharonov and D. Bohm, *Significance of Electromagnetic Potentials in the Quantum Theory*, Physical Review **115** (1959), 485–491.
- [2] W. Ambrose and I. M. Singer, *A theorem on holonomy*, Transactions of the American Mathematical Society **75** (1953), 428–443.

- [3] D. Bleecker, *Gauge Theory and Variational Principles*, Addison–Wesley, Reading, MA, 1981.
- [4] S.-S. Chern, *Complex Manifolds Without Potential Theory*, 2nd ed., Springer, New York, 1979.
- [5] G. W. Gibbons and S. W. Hawking, *Action integrals and partition functions in quantum gravity*, Physical Review D **15** (1977), 2752–2756.
- [6] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry*, Volume I, Interscience Publishers, New York, 1963.
- [7] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry*, Volume II, Interscience Publishers, New York, 1969.
- [8] J. W. Milnor and J. D. Stasheff, *Characteristic Classes*, Annals of Mathematics Studies 76, Princeton University Press, Princeton, NJ, 1974.
- [9] E. Noether, *Invariante Variationsprobleme*, Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse (1918), 235–257.
- [10] V. Pravda, A. Pravdová, A. Coley, and R. Milson, *All spacetimes with vanishing curvature invariants*, Classical and Quantum Gravity **19** (2002), 6213–6236.
- [11] N. Steenrod, *The Topology of Fibre Bundles*, Princeton University Press, Princeton, NJ, 1951.
- [12] R. M. Wald, *General Relativity*, University of Chicago Press, Chicago, 1984.
- [13] T. T. Wu and C. N. Yang, *Concept of nonintegrable phase factors and global formulation of gauge fields*, Physical Review D **12** (1975), 3845–3857.
- [14] C. N. Yang and R. L. Mills, *Conservation of Isotopic Spin and Isotopic Gauge Invariance*, Physical Review **96** (1954), 191–195.
- [15] J. W. York, Jr., *Role of Conformal Three-Geometry in the Dynamics of Gravitation*, Physical Review Letters **28** (1972), 1082–1085.
- [16] J. Rodgers, *Source–Readout Non-Equivalence: Descent and Equivariant Reconstruction Obstructions*, Zenodo, 2026. [doi:10.5281/zenodo.21369967](https://doi.org/10.5281/zenodo.21369967).
- [17] J. Rodgers, *Target-Relative Necessity of Completion: When Readout Loss Obstructs, and What a Sufficient Extension Must Retain*, Zenodo, 2026. [doi:10.5281/zenodo.21370501](https://doi.org/10.5281/zenodo.21370501).
- [18] J. Rodgers, *Canonical Minimal Source Completion: The Coarsest Readout Extension on which a Nominated Family of Source Relations Becomes Well Defined*, Zenodo, 2026. [doi:10.5281/zenodo.21370763](https://doi.org/10.5281/zenodo.21370763).