

CANONICAL MINIMAL SOURCE COMPLETION: THE COARSEST READOUT EXTENSION ON WHICH A NOMINATED FAMILY OF SOURCE RELATIONS BECOMES WELL DEFINED

JEREMY RODGERS

ABSTRACT. Let $p: S \rightarrow T$ be a surjective source-to-readout map on a reduced source space, and let $\{r_\alpha: S \rightarrow V_\alpha\}_{\alpha \in A}$ be a set-indexed family of physically invariant source relations that need not descend along p . We construct the coarsest extension of the readout on which every r_α becomes well defined. The construction is explicit: form the joint map $j_A = (p, (r_\alpha)_{\alpha \in A}): S \rightarrow T \times \prod_{\alpha \in A} V_\alpha$ and set $E_A = \text{Im } j_A$; equivalently, quotient S by equality of the readout together with every selected relation value. We prove that E_A is terminal in the category of surjective relation-sufficient extensions: every sufficient extension E admits a unique surjection $E \rightarrow E_A$. This yields existence, uniqueness up to unique isomorphism, and a precise minimality theorem, without nominating a preferred representative of any readout fiber. We characterize redundant and indispensable relation coordinates by an exact separation criterion, construct canonical refinement maps for nested families and prove their cocycle law, prove invariance under faithful reparameterization of value spaces, and construct the completion determined directly by a family of target-answer maps, together with the canonical surjection from the relation completion onto it. We give the gauge-reduced form of the theorem (completion cannot resurrect redundancy already quotiented away) and the topological form, distinguishing carefully between the quotient topology on E_A and the subspace topology inherited from the product. For flat $U(1)$ connections on a circle, whose curvature readout is identically zero, oriented transport completion returns $U(1)$, whereas the minimal completion for the charged-scalar spectral target is the inversion quotient $U(1)/(z \sim z^{-1}) \cong [-1, 1]$; the canonical refinement $e^{i\theta} \mapsto \cos \theta$ exhibits the general relation-to-target surjection concretely.

1. INTRODUCTION

1.1. The problem inherited from Papers 1 and 2. This is the third paper of a sequence on source-readout non-equivalence. The first companion manuscript [1] established that a readout map $p: S \rightarrow T$ need not admit descent for a given source-defined quantity: a function $r: S \rightarrow V$ factors as $r = \bar{r} \circ p$ if and only if r is constant on every fiber $F_t = p^{-1}(t)$, and equivariant refinements of this statement produce genuine obstructions. The second companion manuscript [2] sharpened the significance of that failure. Information destroyed by a readout is not by itself a defect; it becomes consequential only *relative to a nominated target*. Paper 2 characterized exactly when a target with correct-answer map $a: S \rightarrow A_Q$ is solvable from the readout alone (precisely when a is $\sim_p$ -invariant), and, when it is not, constructed the coarsest extension of T sufficient to determine that single answer map.

What Paper 2 does not supply is the family-level construction. Physical practice almost never nominates a single answer map. One nominates a *family*: a collection of invariant source relations r_α , indexed by a set A , each entering some independently specified mathematical

Date: June 2026.

2020 Mathematics Subject Classification. Primary 18A30, 54B15; Secondary 53C29, 58J50, 81Q70.

Key words and phrases. Quotient construction, universal property, terminal object, relation-sufficient extension, quotient topology, gauge reduction, holonomy, flat connections, spectral invariants, descent.

Website: everythingequation.com.

or physical target. Assembling the single-target completions of Paper 2 one at a time does not answer the family question, because the completions do not simply compose: the object one wants is a single extension of the readout on which *all* the nominated relations become simultaneously well defined, and one wants it to be canonical—adding no distinctions beyond those the nominated family forces. Neither existence, nor minimality, nor uniqueness of such an object is formal from Paper 2.

1.2. What this paper proves. Fix the reduced source S , the surjective readout $p: S \rightarrow T$, and a set-indexed family $r_\alpha: S \rightarrow V_\alpha$, $\alpha \in A$. Write $V_A = \prod_{\alpha \in A} V_\alpha$, $R_A = (r_\alpha)_{\alpha \in A}: S \rightarrow V_A$, and

$$j_A = (p, R_A): S \longrightarrow T \times V_A, \quad E_A := \text{Im}(j_A). \quad (1.1)$$

Let $e_A: S \rightarrow E_A$ be the induced surjection and let $u_A: E_A \rightarrow T$, $\rho_\alpha: E_A \rightarrow V_\alpha$ be the coordinate projections. Our results are as follows.

- (1) *Explicit construction and quotient realization* (Section 3). We have $p = u_A \circ e_A$ and $r_\alpha = \rho_\alpha \circ e_A$ for all α , and $E_A \cong S/\sim_A$ where $s \sim_A s'$ iff $p(s) = p(s')$ and $r_\alpha(s) = r_\alpha(s')$ for every $\alpha \in A$. The information ordering $\Delta_S \subseteq \sim_A \subseteq \sim_p$ holds, and the fiber $u_A^{-1}(t)$ is canonically the set $R_A(F_t)$ of *realized relation signatures* over t .
- (2) *Terminality* (Sections 4 and 5). We define the category of surjective relation-sufficient extensions and prove that $\mathcal{E}_A$ is *terminal* in it: for every sufficient extension $\mathcal{E} = (E, e, u, (\rho_\alpha^E))$ there is a *unique* morphism $\phi_{\mathcal{E}}: E \rightarrow E_A$, and it is surjective. The arrow runs $E \rightarrow E_A$, never $E_A \rightarrow E$. Elementarily, terminality is the statement $\sim_e \subseteq \sim_A$ for every sufficient extension: E_A is the *coarsest* sufficient extension, not the finest.
- (3) *Coordinate economy* (Section 6). A nominated relation r_α is indispensable if and only if some pair s, s' agrees on the readout and on all r_β , $\beta \neq \alpha$, yet is separated by r_α ; and deleting r_α strictly coarsens the completion exactly in that case.
- (4) *Functoriality in the family* (Section 7). For $A \subseteq B$ there is a unique surjection $\pi_{BA}: E_B \rightarrow E_A$ over S , and these satisfy $\pi_{CA} = \pi_{BA} \circ \pi_{CB}$. Faithful reparameterization of value spaces leaves the completion canonically unchanged over T .
- (5) *Target completion* (Section 8). The completion E_Q determined by a family of answer maps is constructed, and when the relations mediate the answers there is a canonical surjection $\chi: E_A \rightarrow E_Q$, an isomorphism precisely when the two induced equivalence relations coincide. Target completion may be *strictly coarser* than relation completion.
- (6) *Gauge boundary and topology* (Sections 9 and 10). Completion occurs *after* gauge reduction and cannot undo it. In the topological setting, E_A with the quotient topology is terminal among topological sufficient extensions; the canonical injection $E_A \hookrightarrow T \times V_A$ is an embedding exactly when quotient and subspace topologies agree, for which compactness of S and Hausdorffness of $T \times V_A$ suffice.
- (7) *Worked example* (Section 11) and *exact reconstruction scope* (Section 12): e_A is injective iff $\{p\} \cup \{r_\alpha\}$ separates points of S .

1.3. What this paper does not prove. The theorem is *model-relative*. Once a reduced source space, a readout, and a nominated family of invariant relations or targets are specified, there is a canonical coarsest sufficient quotient. Nothing here shows that the source space is empirically real, that observed physics is a projection of a hidden source, that every lost distinction must be restored, that the full source is in general recoverable, or that any particular relation family is physically privileged absent independent justification. In particular, “minimal completion” is not “reconstruction of the source”: the two coincide only under the separation criterion of Theorem 12.1. We return to these limits in Section 13.

1.4. Notation. We keep the frozen notation of the companion manuscripts. S is the reduced source space, T the readout space, $p: S \rightarrow T$ the surjective readout, and $F_t = p^{-1}(t)$ the readout fiber over $t \in T$. Throughout, A denotes a *set* indexing the nominated relation family; products $V_A = \prod_{\alpha \in A} V_\alpha$ are ordinary set-theoretic products, so that a point of V_A is a choice function $\alpha \mapsto v_\alpha \in V_\alpha$, and no cardinality restriction on A is needed. We write $\Delta_S = \{(s, s) : s \in S\}$ for the equality relation on S and $\sim_p$ for equality of the readout. For equivalence relations, $\sim \subseteq \approx$ means that $\sim$ is *finer* (identifies fewer pairs, retains more information) and $\approx$ is *coarser*.

2. REDUCED SOURCE, READOUT, AND THE NOMINATED RELATION FAMILY

2.1. The reduced source. We assume, as in Paper 1, that S has already been reduced: all declared gauge redundancy and all coordinate-labelling redundancy have been quotiented out. Concretely, one begins with a raw configuration space Ω carrying an equivalence relation $\simeq_{\text{phys}}$ of physical indistinguishability—gauge equivalence, diffeomorphism relabelling, choice of chart—and forms

$$q: \Omega \longrightarrow S = \Omega / \simeq_{\text{phys}} . \quad (2.1)$$

The order of operations matters and is not a convention we are free to reverse: *reduction precedes completion*. Section 9 states the consequence precisely.

Definition 2.1 (Eligible relation). A raw quantity $\tilde{r}: \Omega \rightarrow V$ is *eligible* for source completion when it is invariant on physical-equivalence classes,

$$\omega \simeq_{\text{phys}} \omega' \implies \tilde{r}(\omega) = \tilde{r}(\omega'), \quad (2.2)$$

so that there is a unique $r: S \rightarrow V$ with $\tilde{r} = r \circ q$. All relations nominated below are assumed eligible, i.e. genuinely defined on the reduced source S .

Eligibility is exactly the condition under which the raw quantity is a function of the physical state rather than of a representative of it. A quantity failing (2.2)—a gauge potential, a component in a chosen frame, a coordinate label—is not a candidate relation, and no completion of the readout can make it one.

2.2. External nomination. Fix a set A and a family of eligible relations

$$r_\alpha: S \longrightarrow V_\alpha, \quad \alpha \in A. \quad (2.3)$$

We stress the epistemic status of the family, because the mathematics below is only as meaningful as the nomination that feeds it. The relations r_α are nominated *externally*: each is selected because it enters an independently specified mathematical or physical target—an amplitude, a spectrum, a transport law, an interference pattern. Physical relevance is **not** inferred from fiber variation. That a function happens to be nonconstant on F_t shows only that the readout does not determine it; it does not show that the distinctions it makes are physical. Confusing the two would convert every artifact of a redundant description into a physical observable. The theorems below are therefore conditional in exactly one place, and it is the right place: they take the family as given and tell one, completely, what follows from it.

2.3. The completion problem. Write

$$V_A := \prod_{\alpha \in A} V_\alpha, \quad R_A: S \rightarrow V_A, \quad R_A(s) := (r_\alpha(s))_{\alpha \in A}. \quad (2.4)$$

Definition 2.2 (The family completion problem). An extension of the readout *sufficient* for the family (2.3) is a space E through which S maps onto, such that

- (1) the original readout is retained: p factors through E ; and

- (2) every nominated relation becomes a well-defined function of the completed state: each r_α factors through E .

The problem is to construct the *coarsest* such extension: one that makes no distinction between source states beyond those forced by requirements (1) and (2), and to prove it canonical.

A sufficient extension is emphatically *not* required to distinguish source states on which both the readout and every selected relation agree. This is the entire content of minimality, and it is what separates completion from reconstruction.

Remark 2.3 (The empty family). If $A = \emptyset$, the ordinary set-theoretic convention gives $V_\emptyset = \{*\}$, the singleton empty product, and $j_\emptyset = (p, *)$, so that $E_\emptyset = \text{Im}(j_\emptyset) \cong T$ and $u_\emptyset$ is a bijection. Nominating nothing completes nothing: the completion of the readout relative to the empty family is the readout. This is the degenerate boundary case of every theorem below, and we will not comment on it again.

3. THE JOINT-IMAGE CONSTRUCTION AND ITS QUOTIENT REALIZATION

Definition 3.1 (Minimal relation completion). Define the *joint map*

$$j_A: S \longrightarrow T \times V_A, \quad j_A(s) = (p(s), R_A(s)), \quad (3.1)$$

and define the *minimal relation completion*

$$E_A := \text{Im}(j_A) \subseteq T \times V_A. \quad (3.2)$$

Let $e_A: S \twoheadrightarrow E_A$, $e_A(s) := j_A(s)$, be the induced surjection, and let

$$u_A: E_A \longrightarrow T, \quad \rho_\alpha: E_A \longrightarrow V_\alpha \quad (\alpha \in A) \quad (3.3)$$

be the restrictions to E_A of the coordinate projections of $T \times V_A$.

Proposition 3.2 (Factorization). *The maps of Theorem 3.1 satisfy*

$$p = u_A \circ e_A, \quad r_\alpha = \rho_\alpha \circ e_A \quad (\alpha \in A), \quad (3.4)$$

and e_A is surjective. In particular every nominated relation descends to E_A .

Proof. Surjectivity of e_A is immediate, since E_A was defined as the image of j_A and e_A is j_A with codomain restricted to that image. For $s \in S$ we have $e_A(s) = (p(s), (r_\alpha(s))_{\alpha \in A})$. Applying the first coordinate projection gives $u_A(e_A(s)) = p(s)$, and applying the projection onto the α -th factor of V_A gives $\rho_\alpha(e_A(s)) = r_\alpha(s)$. As s was arbitrary, (3.4) holds as an identity of functions on S . $\square$

Thus $(E_A, e_A, u_A, (\rho_\alpha)_{\alpha \in A})$ solves the *existence* half of Theorem 2.2. The content of the paper is that it solves the *minimality* and *uniqueness* halves as well.

3.1. The induced equivalence relation.

Definition 3.3. Define a relation $\sim_A$ on S by

$$s \sim_A s' \iff [p(s) = p(s') \text{ and } r_\alpha(s) = r_\alpha(s') \text{ for every } \alpha \in A]. \quad (3.5)$$

This is an equivalence relation, being an intersection of pullbacks of equality relations along functions: reflexivity, symmetry, and transitivity for $\sim_A$ follow coordinatewise from the corresponding properties of equality in T and in each V_α .

Proposition 3.4 (Quotient realization). *The assignment*

$$\kappa: S/\sim_A \longrightarrow E_A, \quad \kappa([s]_A) := j_A(s) \quad (3.6)$$

is a well-defined bijection, and it satisfies $\kappa \circ \pi_A = e_A$, where $\pi_A: S \rightarrow S/\sim_A$ is the canonical projection. Hence

$$E_A \cong S/\sim_A \quad (3.7)$$

canonically.

Proof. We verify the three assertions separately, rather than invoking a generic isomorphism principle for quotients, since the well-definedness step is precisely where the definition of the relation $\sim_A$ earns its keep.

Well-definedness. Suppose $[s]_A = [s']_A$, i.e. $s \sim_A s'$. By (3.5) this means $p(s) = p(s')$ and $r_\alpha(s) = r_\alpha(s')$ for every $\alpha \in A$. Two elements of the product $T \times V_A$ are equal precisely when all of their coordinates agree; and here the coordinates of $j_A(s)$ and $j_A(s')$ are $p(s), p(s')$ in the T -factor and $r_\alpha(s), r_\alpha(s')$ in the α -th factor of V_A , respectively. All agree, so $j_A(s) = j_A(s')$. Hence the value assigned by (3.6) does not depend on the chosen representative, and κ is a well-defined function.

Injectivity. Suppose $\kappa([s]_A) = \kappa([s']_A)$, i.e. $j_A(s) = j_A(s')$. Comparing coordinates in $T \times V_A$ gives $p(s) = p(s')$ and $r_\alpha(s) = r_\alpha(s')$ for every $\alpha \in A$, which is exactly (3.5). Hence $s \sim_A s'$ and $[s]_A = [s']_A$.

Surjectivity. Let $x \in E_A$. By Theorem 3.1, $E_A = \text{Im}(j_A)$, so there is $s \in S$ with $x = j_A(s) = \kappa([s]_A)$.

Finally, for every $s \in S$ we have $\kappa(\pi_A(s)) = \kappa([s]_A) = j_A(s) = e_A(s)$, so $\kappa \circ \pi_A = e_A$. $\square$

So the joint image and the quotient by “agreement of the readout and of every nominated relation” are the same object, viewed extensionally and intensionally. The first description makes existence manifest; the second makes minimality manifest.

3.2. The information ordering.

Proposition 3.5 (Information ordering). *As relations on S ,*

$$\Delta_S \subseteq \sim_A \subseteq \sim_p, \quad (3.8)$$

and consequently there is a factorization of the readout,

$$S \xrightarrow{e_A} E_A \xrightarrow{u_A} T. \quad (3.9)$$

Proof. The first inclusion is reflexivity: $s \sim_A s$ for every s , since $p(s) = p(s)$ and $r_\alpha(s) = r_\alpha(s)$. The second holds because the defining condition (3.5) contains the clause $p(s) = p(s')$ as a conjunct; hence $s \sim_A s'$ implies $s \sim_p s'$. The factorization (3.9) is (3.4). $\square$

The reading of (3.8) must be got right, because the direction of the inclusions encodes the direction of every universal arrow in the paper. A *smaller* equivalence relation identifies *fewer* pairs, hence corresponds to a *finer* partition and to *more* retained information. Thus (3.8) says: equality on S retains everything; $\sim_A$ retains at most that, and at least the readout; $\sim_p$ retains the least of the three. The completion E_A refines the readout partition, but—this is the whole point of minimality—it need not refine it all the way down to individual source states.

3.3. What completion adds above a readout value.

Proposition 3.6 (Completed fibers). *For every $t \in T$ there is a canonical bijection*

$$u_A^{-1}(t) \cong R_A(F_t) = \{ R_A(s) : s \in F_t = p^{-1}(t) \} \subseteq V_A. \quad (3.10)$$

Proof. Consider the second-coordinate projection $\text{pr}_{V_A} : T \times V_A \rightarrow V_A$ and restrict it to $u_A^{-1}(t)$.

Its image is $R_A(F_t)$. An element $x \in u_A^{-1}(t)$ lies in $E_A = \text{Im}(j_A)$, so $x = (p(s), R_A(s))$ for some $s \in S$; and $u_A(x) = t$ forces $p(s) = t$, i.e. $s \in F_t$. Hence $\text{pr}_{V_A}(x) = R_A(s) \in R_A(F_t)$. Conversely, given $v = R_A(s)$ with $s \in F_t$, the point $x = (t, v) = j_A(s)$ lies in E_A and satisfies $u_A(x) = t$, so v is attained.

It is injective on $u_A^{-1}(t)$. If $x, x' \in u_A^{-1}(t)$ have the same V_A -coordinate, then they also share the T -coordinate, namely t ; being elements of a product with all coordinates equal, $x = x'$.

Thus pr_{V_A} restricts to a bijection $u_A^{-1}(t) \rightarrow R_A(F_t)$, and this bijection was constructed without any choice. $\square$

Theorem 3.6 answers exactly the question that “completion” invites one to ask badly. What sits above a readout value t in E_A is *not* a preferred source state in F_t , and not an arbitrary enrichment of t by extra labels. It is the set of *realized relation signatures* over t : the distinct values that the nominated family actually takes on the fiber. Two source states in F_t that agree on every r_α contribute the *same* completed state; nothing in the construction selects between them, and nothing in the construction can. We record this as a standing prohibition, because it is the failure mode that a naive completion invites.

Remark 3.7 (No selector, no section, no representative). The construction of E_A uses no selection function on the fibers of p , no section of p , and no canonical representative of F_t ; it requires no choice principle beyond forming the image of a function, and it produces none of these objects either. This should be stated precisely, since the point is about what the construction *supplies*, not about what is set-theoretically possible.

If one separately chooses a representative of each $\sim_A$ -class—a choice that in general requires an appeal to a choice principle, and which is available when such a choice can be made—then the resulting set of chosen representatives is in bijection with E_A , by **Theorem 3.4**. Such a set is a *noncanonical presentation* of the completion. It carries no invariant information beyond what E_A already carries: the bijection depends on the choice, different choices give different presentations of the same quotient, and nothing in the completion distinguishes one of them. Crucially, no such presentation is supplied or privileged by the universal construction. The object E_A is defined without any choice whatever, as an image; and terminality (**Theorem 5.1**) yields only the surjections $\phi_{\mathcal{E}} : E \rightarrow E_A$ onto the completion. It furnishes no inverse to e_A , no canonical splitting, and no canonical representative selector. Whatever a representative-based presentation may be used for, it is extra structure imposed from outside, not a consequence of the theorem.

4. THE CATEGORY OF RELATION-SUFFICIENT EXTENSIONS

We now fix the data $(S, T, p, \{r_\alpha\}_{\alpha \in A})$ for the remainder of the paper and make **Theorem 2.2** precise as a category.

Definition 4.1 (Relation-sufficient extension). A *relation-sufficient extension* of p (relative to the nominated family) is a tuple

$$\mathcal{E} = (E, e, u, (\rho_\alpha^E)_{\alpha \in A}) \quad (4.1)$$

consisting of a set E , a **surjection** $e: S \twoheadrightarrow E$, and maps $u: E \rightarrow T$ and $\rho_\alpha^E: E \rightarrow V_\alpha$ ($\alpha \in A$) satisfying

$$p = u \circ e, \quad r_\alpha = \rho_\alpha^E \circ e \quad (\alpha \in A). \quad (4.2)$$

Two features of [Theorem 4.1](#) deserve comment.

First, the factorizations (4.2) are exactly requirements (1) and (2) of [Theorem 2.2](#): the readout survives, and every nominated relation is a function of the completed state.

Second, surjectivity of e is not a technical convenience but a substantive part of the definition, and it is what makes the universal property come out uniquely. A completed state that no source state realizes is not a state of the system; it is a spurious element. If unrealized elements were permitted, then given any sufficient E one could adjoin an arbitrary point $*$ with arbitrary values $u(*)$ and $\rho_\alpha^E(*)$ (subject to nothing, since (4.2) constrains only the image of e), and the map to E_A would no longer be determined: the value $\phi(*)$ would be unconstrained by $e_A = \phi \circ e$. Uniqueness in [Theorem 5.1](#) would fail. Surjectivity excludes exactly this, and it does so for the honest reason: a completion consists of states *realized* by source states.

Definition 4.2 (Morphism). Let $\mathcal{E} = (E, e, u, (\rho_\alpha^E))$ and $\mathcal{E}' = (E', e', u', (\rho_\alpha^{E'}))$ be relation-sufficient extensions. A *morphism* $f: \mathcal{E} \rightarrow \mathcal{E}'$ is a function $f: E \rightarrow E'$ satisfying

$$e' = f \circ e. \quad (4.3)$$

Thus a morphism is a map *under* S : it is compatible with the two completion maps out of the source. We now check that this single condition already forces compatibility with all the structure maps.

Lemma 4.3 (Structure maps are automatically preserved). *If $f: \mathcal{E} \rightarrow \mathcal{E}'$ is a morphism in the sense of [Theorem 4.2](#), then*

$$u' \circ f = u, \quad \rho_\alpha^{E'} \circ f = \rho_\alpha^E \quad (\alpha \in A). \quad (4.4)$$

Proof. Let $x \in E$. Since e is surjective, choose $s \in S$ with $x = e(s)$. Then, using (4.3) and the factorizations (4.2) for $\mathcal{E}'$ and then for $\mathcal{E}$,

$$u'(f(x)) = u'(f(e(s))) = u'(e'(s)) = p(s) = u(e(s)) = u(x),$$

and likewise, for each $\alpha \in A$,

$$\rho_\alpha^{E'}(f(x)) = \rho_\alpha^{E'}(e'(s)) = r_\alpha(s) = \rho_\alpha^E(e(s)) = \rho_\alpha^E(x).$$

Both computations are independent of the chosen preimage s , since their left-hand sides do not mention s . As x was arbitrary, (4.4) holds. $\square$

We display (4.4) rather than merely deducing it, because it records what a morphism of completions actually preserves: the readout and every nominated relation value. What is *not* imposed is any compatibility beyond that—no metric, no ordering, no choice of labels.

Proposition 4.4 (This is a category). *Relation-sufficient extensions of p relative to $\{r_\alpha\}_{\alpha \in A}$, with the morphisms of [Theorem 4.2](#), form a category, which we denote $\mathbf{Suff}(p; \{r_\alpha\})$.*

Proof. For $\mathcal{E} = (E, e, u, (\rho_\alpha^E))$, the identity function id_E satisfies $e = \text{id}_E \circ e$, so it is a morphism $\mathcal{E} \rightarrow \mathcal{E}$. If $f: \mathcal{E} \rightarrow \mathcal{E}'$ and $g: \mathcal{E}' \rightarrow \mathcal{E}''$ are morphisms, then $(g \circ f) \circ e = g \circ (f \circ e) = g \circ e' = e''$, so $g \circ f: \mathcal{E} \rightarrow \mathcal{E}''$ is a morphism. Associativity of composition and the unit laws are inherited from composition of functions. $\square$

We add no further categorical apparatus; none is needed.

5. THE CANONICAL MINIMAL-COMPLETION THEOREM

Theorem 5.1 (Canonical minimal completion). *The tuple*

$$\mathcal{E}_A = (E_A, e_A, u_A, (\rho_\alpha)_{\alpha \in A}) \quad (5.1)$$

*of Theorem 3.1 is an object of $\mathbf{Suff}(p; \{r_\alpha\})$, and it is **terminal**: for every relation-sufficient extension $\mathcal{E} = (E, e, u, (\rho_\alpha^E))$ there exists a **unique** morphism*

$$\phi_{\mathcal{E}}: \mathcal{E} \longrightarrow \mathcal{E}_A, \quad \text{that is,} \quad \phi_{\mathcal{E}}: E \longrightarrow E_A. \quad (5.2)$$

It is given by the explicit formula

$$\phi_{\mathcal{E}}(x) = \left(u(x), (\rho_\alpha^E(x))_{\alpha \in A} \right), \quad x \in E, \quad (5.3)$$

and $\phi_{\mathcal{E}}$ is surjective.

Before the proof, we emphasize the direction, since reversing it inverts the meaning of the theorem. The universal map runs **from an arbitrary sufficient extension to E_A** , not from E_A to it. E_A is therefore the *terminal*—the *coarsest*—sufficient extension: every sufficient extension maps *onto* it, and so retains at least as much information as it does. E_A is not an initial object, not a universal cover, and not the finest extension; the finest sufficient extension is S itself (with $e = \text{id}_S$), which is precisely the object one is trying not to be forced into.

Proof. $\mathcal{E}_A$ is an object. e_A is surjective and the factorizations (3.4) hold, by Theorem 3.2.

Now fix a relation-sufficient extension $\mathcal{E} = (E, e, u, (\rho_\alpha^E))$ and define $\phi_{\mathcal{E}}$ by the formula (5.3). We establish the required properties in turn.

(i) *The formula is well defined as a function $E \rightarrow T \times V_A$.* For each $x \in E$ the value $u(x) \in T$ is determined, and for each $\alpha \in A$ the value $\rho_\alpha^E(x) \in V_\alpha$ is determined; the assignment $\alpha \mapsto \rho_\alpha^E(x)$ is therefore an element of the set-theoretic product $V_A = \prod_{\alpha \in A} V_\alpha$. Hence (5.3) assigns to each $x \in E$ exactly one element of $T \times V_A$. No choice is involved, and no representative of x is used: the formula depends only on x through the given structure maps of $\mathcal{E}$.

(ii) *The image lies in E_A , and $e_A = \phi_{\mathcal{E}} \circ e$.* Let $x \in E$. By surjectivity of e there is $s \in S$ with $x = e(s)$. Then, using (4.2),

$$\phi_{\mathcal{E}}(x) = \phi_{\mathcal{E}}(e(s)) = \left(u(e(s)), (\rho_\alpha^E(e(s)))_{\alpha \in A} \right) = \left(p(s), (r_\alpha(s))_{\alpha \in A} \right) = j_A(s) = e_A(s). \quad (5.4)$$

The right-hand side lies in $\text{Im}(j_A) = E_A$ by Theorem 3.1. Since $x \in E$ was arbitrary, $\phi_{\mathcal{E}}(E) \subseteq E_A$, so $\phi_{\mathcal{E}}$ may be regarded as a map $E \rightarrow E_A$; and reading (5.4) from the ends, $\phi_{\mathcal{E}}(e(s)) = e_A(s)$ for every $s \in S$, i.e.

$$e_A = \phi_{\mathcal{E}} \circ e. \quad (5.5)$$

Note that the computation (5.4) is legitimate for *every* choice of preimage s of x , and its left-hand side does not depend on s ; so no consistency question arises.

(iii) *It is a morphism, and preserves all structure maps.* Equation (5.5) is precisely condition (4.3) of Theorem 4.2, so $\phi_{\mathcal{E}}$ is a morphism $\mathcal{E} \rightarrow \mathcal{E}_A$. Preservation of the structure maps then follows from Theorem 4.3, but in the present case it is also visible directly from (5.3): applying the coordinate projections u_A and ρ_α to $\phi_{\mathcal{E}}(x) = (u(x), (\rho_\beta^E(x))_\beta)$ gives

$$u_A \circ \phi_{\mathcal{E}} = u, \quad \rho_\alpha \circ \phi_{\mathcal{E}} = \rho_\alpha^E \quad (\alpha \in A). \quad (5.6)$$

(iv) *Surjectivity.* Let $y \in E_A$. By Theorem 3.2 the map e_A is surjective, so $y = e_A(s)$ for some $s \in S$. By (5.5), $y = \phi_{\mathcal{E}}(e(s))$, and $e(s) \in E$. Hence $\phi_{\mathcal{E}}$ is onto E_A .

(v) *Uniqueness.* Suppose $\psi: \mathcal{E} \rightarrow \mathcal{E}_A$ is any morphism, i.e. any function $\psi: E \rightarrow E_A$ with $e_A = \psi \circ e$. Let $x \in E$ and choose $s \in S$ with $x = e(s)$. Then

$$\psi(x) = \psi(e(s)) = e_A(s) = \phi_{\mathcal{E}}(e(s)) = \phi_{\mathcal{E}}(x),$$

using $e_A = \psi \circ e$ for the second equality and (5.5) for the third. Since x was arbitrary, $\psi = \phi_{\mathcal{E}}$. (This is exactly where surjectivity of e is consumed: it guarantees that the condition $e_A = \psi \circ e$ pins down ψ at *every* point of E , not merely on a subset.) $\square$

Corollary 5.2 (Uniqueness up to unique isomorphism). *Any two terminal objects of the category $\mathbf{Suff}(p; \{r_\alpha\})$ are isomorphic by a unique isomorphism, and this isomorphism commutes with the completion maps from S and with all of the structure maps. In particular, any relation-sufficient extension that satisfies the universal property of Theorem 5.1 is canonically isomorphic to $\mathcal{E}_A$.*

Proof. Let $\mathcal{E}, \mathcal{E}'$ be terminal. Terminality of $\mathcal{E}'$ gives a unique morphism $f: \mathcal{E} \rightarrow \mathcal{E}'$; terminality of $\mathcal{E}$ gives a unique morphism $g: \mathcal{E}' \rightarrow \mathcal{E}$. Then $g \circ f: \mathcal{E} \rightarrow \mathcal{E}$ is a morphism, as is $\text{id}_{\mathcal{E}}$; by terminality of $\mathcal{E}$ there is only one morphism $\mathcal{E} \rightarrow \mathcal{E}$, so $g \circ f = \text{id}_{\mathcal{E}}$. Symmetrically $f \circ g = \text{id}_{\mathcal{E}'}$. Hence f is an isomorphism, and it is unique because morphisms into a terminal object are unique. Compatibility with the completion maps is (4.3); compatibility with the structure maps is Theorem 4.3. $\square$

Corollary 5.3 (Coarsest sufficiency, elementary form). *Let $\mathcal{E} = (E, e, u, (\rho_\alpha^E))$ be any relation-sufficient extension, and define its induced equivalence relation by*

$$s \sim_e s' \quad :\Longleftrightarrow \quad e(s) = e(s'). \quad (5.7)$$

Then

$$\sim_e \subseteq \sim_A. \quad (5.8)$$

Equivalently: every sufficient extension separates at least every pair that E_A separates, while E_A separates no pair that sufficiency does not force it to separate.

Proof. Suppose $e(s) = e(s')$. Applying u and using $p = u \circ e$,

$$p(s) = u(e(s)) = u(e(s')) = p(s').$$

Applying ρ_α^E and using $r_\alpha = \rho_\alpha^E \circ e$, for each $\alpha \in A$,

$$r_\alpha(s) = \rho_\alpha^E(e(s)) = \rho_\alpha^E(e(s')) = r_\alpha(s').$$

By (3.5) this is precisely $s \sim_A s'$. $\square$

Theorem 5.3 and Theorem 5.1 are two faces of one statement. Terminality is the categorical form; the inclusion (5.8) is the elementary, checkable form. The inclusion also settles the direction question once and for all: since $\sim_e$ is contained in $\sim_A$, the partition induced by e *refines* the partition induced by e_A , so there is a well-defined map from the finer quotient onto the coarser one—from E onto E_A , and not back.

Remark 5.4 (The sense in which the completion is a construction). Theorem 5.1 constructs the coarsest relation-sufficient extension in the paper's category $\mathbf{Suff}(p; \{r_\alpha\})$ of set-theoretic source quotients: the object E_A exists, is terminal, and is unique up to unique isomorphism, and those claims are unconditional. What the theorem builds it builds from the nominated relations themselves. The joint map $j_A = (p, R_A)$ takes the r_α as given input, so the construction presupposes the relation data rather than producing it; in particular it furnishes *no* procedure for recovering the r_α from the readout p when they fail to descend—if such a procedure existed the relations would already be functions of p , and the completion would be T itself. Nor does

terminality establish that E_A is experimentally accessible, measurable, smooth, geometric, or realized as a physical state space independent of its defining coordinates; E_A is, as constructed, a quotient of S recording exactly the nominated values. A concrete realization must independently supply and justify the relation data, and any additional geometric or physical structure carried by E_A , on its own terms. That realization question is not settled here; it is deferred to the companion manuscript on geometric realization discussed in [Section 13](#).

Remark 5.5 (What “minimal” means, and what it does not). Minimality here means: *least retained distinction among sufficient extensions*. It is a statement about the ordering of equivalence relations by inclusion, equivalently about the ordering of partitions by refinement. It does *not* mean minimum cardinality selected without regard to structure—the one-point set is smaller than E_A and is not, in general, sufficient—and it does *not* mean that E_A is a small or economical description of S in any other sense. Above all, minimal completion is *not* reconstruction of S : by [Theorem 3.5](#) the map e_A need not be injective, and [Section 12](#) determines exactly when it is.

6. REDUNDANT AND INDISPENSABLE RELATION COORDINATES

The nominated family may be wasteful: one of its members may already be determined by the readout together with the others, in which case its presence in the list changes nothing. We characterize this exactly.

Throughout this section fix $\alpha \in A$ and write

$$A_{-\alpha} := A \setminus \{\alpha\}, \quad R_{A_{-\alpha}} = (r_\beta)_{\beta \in A_{-\alpha}} : S \rightarrow V_{A_{-\alpha}}. \quad (6.1)$$

Definition 6.1 (Redundancy and indispensability). The relation r_α is *redundant relative to p and the remaining relations* if it factors through their joint map, i.e. if there exists a function $f_\alpha : T \times V_{A_{-\alpha}} \rightarrow V_\alpha$ with

$$r_\alpha = f_\alpha \circ (p, R_{A_{-\alpha}}). \quad (6.2)$$

It is *indispensable* if no such factorization exists.

Theorem 6.2 (Indispensability criterion). *The following are equivalent.*

- (1) r_α is indispensable.
- (2) There exist $s, s' \in S$ with

$$p(s) = p(s'), \quad r_\beta(s) = r_\beta(s') \quad \text{for every } \beta \in A_{-\alpha}, \quad r_\alpha(s) \neq r_\alpha(s'). \quad (6.3)$$

Proof. We prove the contrapositive equivalence: r_α is redundant if and only if no pair (6.3) exists, i.e. if and only if r_α is constant on every fiber of the joint map $g := (p, R_{A_{-\alpha}}) : S \rightarrow T \times V_{A_{-\alpha}}$. This is the descent statement of Paper 1 [1] applied to g ; we give the argument in full.

Redundant $\Rightarrow$ constant on fibers of g . Suppose (6.2) holds and suppose $g(s) = g(s')$, i.e. $p(s) = p(s')$ and $r_\beta(s) = r_\beta(s')$ for all $\beta \neq \alpha$. Applying f_α to the equal values $g(s) = g(s')$ gives $r_\alpha(s) = f_\alpha(g(s)) = f_\alpha(g(s')) = r_\alpha(s')$. Hence no pair as in (6.3) can exist.

Constant on fibers of $g \Rightarrow$ redundant. Suppose that $g(s) = g(s')$ implies $r_\alpha(s) = r_\alpha(s')$. Define f_α on the image of g by

$$f_\alpha(g(s)) := r_\alpha(s) \quad (s \in S), \quad (6.4)$$

The prescription (6.4) is well defined on $\text{Im}(g)$ precisely by the assumed fiber-constancy: if $g(s) = g(s')$ then the two prescribed values $r_\alpha(s)$ and $r_\alpha(s')$ agree. Every point of $\text{Im}(g)$ is of the form $g(s)$, so f_α is defined on all of $\text{Im}(g)$. It remains to extend f_α to the whole of $T \times V_{A_{-\alpha}}$, and we treat the two cases explicitly.

Case $S \neq \emptyset$. Fix any $s_0 \in S$ and define $f_\alpha(y) := r_\alpha(s_0)$ for every $y \in (T \times V_{A-\alpha}) \setminus \text{Im}(g)$. This is legitimate because $r_\alpha(s_0)$ is an element of V_α , so $V_\alpha \neq \emptyset$ and a value is available. The extension changes nothing on $\text{Im}(g)$, where (6.4) already fixes f_α .

Case $S = \emptyset$. Then $T = \emptyset$: the readout $p: S \rightarrow T$ is surjective by standing hypothesis, and the image of the empty map is empty, so $T = p(S) = \emptyset$. Hence the product $T \times V_{A-\alpha} = \emptyset$ as well, and the unique map $\emptyset \rightarrow V_\alpha$ —the empty map—serves as f_α . It satisfies $r_\alpha = f_\alpha \circ g$ vacuously as an identity of maps out of $S = \emptyset$, both sides being the empty map. So r_α is redundant in this case, which is consistent: no witness pair (6.3) can exist in an empty source, so r_α is not indispensable, exactly as the equivalence requires. (The two sides of the criterion agree here; the case is not merely set aside.)

In either case f_α is defined on all of $T \times V_{A-\alpha}$ and satisfies $r_\alpha = f_\alpha \circ g$, which is (6.2).

Finally, “no pair as in (6.3)” is by inspection the same statement as “ $g(s) = g(s')$ implies $r_\alpha(s) = r_\alpha(s')$ ”, since the first two clauses of (6.3) say $g(s) = g(s')$ and the third says $r_\alpha(s) \neq r_\alpha(s')$. Negating both sides of the equivalence just proved yields $1 \Leftrightarrow 2$. $\square$

The criterion is a genuine separation test, and it is the reason a long list of nominated relations is not thereby an economical one. A coordinate earns its place only by separating some pair of source states that the readout and all the *other* coordinates identify.

Corollary 6.3 (Strict coarsening on deletion). *Let $\sim_A$ and $\sim_{A-\alpha}$ be the equivalence relations (3.5) induced by the full family and by the family with r_α deleted, and let $\pi: E_A \rightarrow E_{A-\alpha}$ be the canonical surjection of Theorem 7.1 below (with $B = A$). Then always $\sim_A \subseteq \sim_{A-\alpha}$, and the following are equivalent:*

- (1) *the inclusion is strict: $\sim_A \subsetneq \sim_{A-\alpha}$;*
- (2) *π is not injective, i.e. deleting r_α strictly coarsens the completion;*
- (3) *r_α is indispensable.*

Proof. The inclusion $\sim_A \subseteq \sim_{A-\alpha}$ is immediate from (3.5), since the defining condition for $\sim_A$ is the condition for $\sim_{A-\alpha}$ conjoined with the extra clause $r_\alpha(s) = r_\alpha(s')$.

(1) $\Leftrightarrow$ (3). Strictness means: there exist s, s' with $s \sim_{A-\alpha} s'$ but not $s \sim_A s'$. Unwinding (3.5), this says $p(s) = p(s')$ and $r_\beta(s) = r_\beta(s')$ for all $\beta \neq \alpha$, while $r_\alpha(s) \neq r_\alpha(s')$ —which is exactly the witness (6.3). By Theorem 6.2 this holds iff r_α is indispensable.

(1) $\Leftrightarrow$ (2). By Theorem 3.4, $E_A \cong S/\sim_A$ and $E_{A-\alpha} \cong S/\sim_{A-\alpha}$, and under these identifications π is the canonical map $S/\sim_A \rightarrow S/\sim_{A-\alpha}$ sending $[s]_A \mapsto [s]_{A-\alpha}$, which is well defined by the inclusion just proved. This map is injective iff $[s]_{A-\alpha} = [s']_{A-\alpha}$ implies $[s]_A = [s']_A$, i.e. iff $\sim_{A-\alpha} \subseteq \sim_A$, i.e. (given the reverse inclusion) iff the two relations are equal. Its failure is precisely strictness. $\square$

Remark 6.4 (Irredundancy is not uniqueness of coordinates). A family in which every member is indispensable is *irredundant relative to the nomination*. This is a statement about the given list, and it confers no uniqueness on the list. Different irredundant families may induce the very same equivalence relation $\sim_A$ and hence, by Theorem 3.4, the very same completion—Theorem 7.3 below exhibits an entire class of such rewritings. An irredundant list of relation coordinates must therefore not be confused with a preferred coordinate system on the completion.

The correct statement of the coordinate situation is a relative one, and we make it precisely. *Relative to the fixed nominated family $\{r_\alpha\}_{\alpha \in A}$* , the completion E_A does carry canonical structure maps: the readout projection $u_A: E_A \rightarrow T$ and the relation coordinates $\rho_\alpha: E_A \rightarrow V_\alpha$ of (3.3), which are part of the data of the terminal object and are determined by it. What E_A does *not* possess is a family-independent, intrinsically unique coordinate presentation: a system of coordinates singled out by the quotient alone, without reference to a nomination.

Faithful reparameterizations ([Theorem 7.3](#)) and distinct irredundant families may present one and the same quotient in different coordinates, and the quotient does not choose among them. The invariant is E_A together with its structure maps for the given nomination; the coordinates are a presentation relative to that nomination.

7. NESTED FAMILIES, REFINEMENT MAPS, AND REPARAMETERIZATION

7.1. Nested families. Let $A \subseteq B$ be sets of indices, the A -relations being a subfamily of the B -relations (i.e. the relations indexed by A are the same functions in both families). Then, directly from (3.5),

$$\sim_B \subseteq \sim_A, \quad (7.1)$$

since agreement of *all* B -relations (and the readout) entails agreement of the sub-collection indexed by A (and the readout). In words: a larger nomination is at least as demanding, hence yields an at least as fine completion.

Proposition 7.1 (Canonical refinement maps). *Let $A \subseteq B$. There is a unique map*

$$\pi_{BA}: E_B \longrightarrow E_A \quad \text{with} \quad e_A = \pi_{BA} \circ e_B, \quad (7.2)$$

and it is surjective. It satisfies $u_A \circ \pi_{BA} = u_B$ and $\rho_\alpha^{(A)} \circ \pi_{BA} = \rho_\alpha^{(B)}$ for every $\alpha \in A$, where the superscripts indicate the completion in which the structure map lives. Moreover $\pi_{AA} = \text{id}_{E_A}$, and for $A \subseteq B \subseteq C$,

$$\pi_{CA} = \pi_{BA} \circ \pi_{CB}. \quad (7.3)$$

Proof. Existence and uniqueness. We exhibit E_B as a relation-sufficient extension for the A -family and invoke [Theorem 5.1](#). Indeed, $e_B: S \twoheadrightarrow E_B$ is surjective; $u_B: E_B \rightarrow T$ satisfies $p = u_B \circ e_B$; and for each $\alpha \in A \subseteq B$ the structure map $\rho_\alpha^{(B)}: E_B \rightarrow V_\alpha$ satisfies $r_\alpha = \rho_\alpha^{(B)} \circ e_B$. Hence $(E_B, e_B, u_B, (\rho_\alpha^{(B)})_{\alpha \in A})$ is an object of $\mathbf{Suff}(p; \{r_\alpha\}_{\alpha \in A})$. [Theorem 5.1](#) supplies a unique morphism from it to the terminal object $\mathcal{E}_A$, i.e. a unique map $\pi_{BA}: E_B \rightarrow E_A$ with $e_A = \pi_{BA} \circ e_B$, and asserts that it is surjective. Compatibility with u and with the ρ_α , $\alpha \in A$, is [Theorem 4.3](#) (equivalently (5.6)).

Identity. For $B = A$, the identity id_{E_A} satisfies $e_A = \text{id}_{E_A} \circ e_A$; by the uniqueness just established, $\pi_{AA} = \text{id}_{E_A}$.

Cocycle law. Let $A \subseteq B \subseteq C$. Both π_{CA} and $\pi_{BA} \circ \pi_{CB}$ are maps $E_C \rightarrow E_A$. Compose each with the surjection e_C :

$$\pi_{CA} \circ e_C = e_A, \quad (\pi_{BA} \circ \pi_{CB}) \circ e_C = \pi_{BA} \circ (\pi_{CB} \circ e_C) = \pi_{BA} \circ e_B = e_A.$$

Two maps out of E_C that agree after precomposition with the surjection e_C agree on all of E_C : given $x \in E_C$, write $x = e_C(s)$ and evaluate. Hence (7.3). (Alternatively, both maps are morphisms $\mathcal{E}_C \rightarrow \mathcal{E}_A$ in $\mathbf{Suff}(p; \{r_\alpha\}_{\alpha \in A})$, and [Theorem 5.1](#) says there is only one.) $\square$

Nested nominations therefore generate a coherent tower of completions,

$$S \longrightarrow E_C \xrightarrow{\pi_{CB}} E_B \xrightarrow{\pi_{BA}} E_A \xrightarrow{u_A} T \quad (A \subseteq B \subseteq C), \quad (7.4)$$

in which each arrow forgets exactly the distinctions that the smaller nomination does not require. This is a hierarchy of *information content*. Its arrows run from finer to coarser, and by [Theorem 5.1](#) each is uniquely determined by the pair of nominations.

Remark 7.2 (Strictness is not automatic). The completion for a larger family is at least as fine as that for a smaller one, and *may* be strictly finer, but need not be. If every relation indexed by $B \setminus A$ is redundant relative to p and the A -relations, then $\sim_B = \sim_A$ by the argument of [Theorem 6.3](#), and π_{BA} is a bijection. Adding relations to the nomination adds information to

the completion precisely to the extent that the added relations separate pairs that the existing nomination identifies.

7.2. Faithful reparameterization. The completion depends on the *distinctions* carried by the nominated relations, not on the labels in which their values happen to be recorded. We make this precise.

Proposition 7.3 (Invariance under faithful reparameterization). *For each $\alpha \in A$ let $f_\alpha: V_\alpha \rightarrow W_\alpha$ be a map whose restriction to the realized subset $r_\alpha(S) \subseteq V_\alpha$ is injective, and set*

$$r'_\alpha := f_\alpha \circ r_\alpha: S \rightarrow W_\alpha. \quad (7.5)$$

Let $\sim_A$ and $\sim'_A$ be the equivalence relations induced by $\{r_\alpha\}$ and by $\{r'_\alpha\}$ (with the same readout p), and let E_A, E'_A be the corresponding completions. Then

$$\sim_A = \sim'_A, \quad (7.6)$$

and consequently there is a canonical bijection $E_A \cong E'_A$ commuting with the completion maps from S and with the projections to T :

$$u'_A \circ (E_A \xrightarrow{\cong} E'_A) = u_A. \quad (7.7)$$

Proof. $\sim_A \subseteq \sim'_A$. Suppose $s \sim_A s'$. Then $p(s) = p(s')$, and for each α , $r_\alpha(s) = r_\alpha(s')$; applying the function f_α to equal arguments gives $r'_\alpha(s) = f_\alpha(r_\alpha(s)) = f_\alpha(r_\alpha(s')) = r'_\alpha(s')$. Hence $s \sim'_A s'$. (Only well-definedness of f_α is used here; injectivity is not.)

$\sim'_A \subseteq \sim_A$. Suppose $s \sim'_A s'$. Then $p(s) = p(s')$, and for each α , $f_\alpha(r_\alpha(s)) = f_\alpha(r_\alpha(s'))$. Both $r_\alpha(s)$ and $r_\alpha(s')$ lie in the realized subset $r_\alpha(S)$, on which f_α is injective by hypothesis; therefore $r_\alpha(s) = r_\alpha(s')$. As this holds for every α , $s \sim_A s'$.

This proves (7.6). By Theorem 3.4 applied to each family, $E_A \cong S/\sim_A = S/\sim'_A \cong E'_A$, and the composite bijection sends $e_A(s) \mapsto e'_A(s)$, so it commutes with the completion maps. Since $u_A(e_A(s)) = p(s) = u'_A(e'_A(s))$ for every s , and e_A is surjective, (7.7) follows. $\square$

Remark 7.4 (The precise form of the invariance). The value spaces V_α and W_α are in general different sets, so $E_A \subseteq T \times V_A$ and $E'_A \subseteq T \times W_A$ are subsets of different products. Theorem 7.3 therefore asserts an isomorphism of the *induced quotients over T* —equivalently, of the two objects as completions of the readout—and not an identity, nor an isomorphism internal to a single fixed category of relation signatures valued in one fixed product. The relation-signature coordinates are a presentation; the quotient is the invariant. Faithful relabelling of the presentation does not move the invariant.

8. COMPLETION DETERMINED BY A FAMILY OF TARGET ANSWERS

Paper 2 [2] works with a single target Q and its correct-answer map $a: S \rightarrow A_Q$, and constructs the coarsest extension of the readout that determines that answer; we will denote that single-target object by E_Q , following Paper 2. We now treat a *family* of targets, and we choose notation that cannot collide with the single-target object.

Let $\mathcal{Q}$ be a set of targets, each equipped with a correct-answer map

$$a_Q: S \rightarrow A_Q, \quad Q \in \mathcal{Q}, \quad (8.1)$$

and define

$$j_{\mathcal{Q}}: S \rightarrow T \times \prod_{Q \in \mathcal{Q}} A_Q, \quad j_{\mathcal{Q}}(s) = \left(p(s), (a_Q(s))_{Q \in \mathcal{Q}} \right), \quad E_{\mathcal{Q}} := \text{Im}(j_{\mathcal{Q}}), \quad (8.2)$$

with induced surjection $e_Q: S \twoheadrightarrow E_Q$, projection $u_Q: E_Q \rightarrow T$, and structure maps $\sigma_Q: E_Q \rightarrow A_Q$. Write $\sim_Q$ for the induced equivalence relation: $s \sim_Q s'$ iff $p(s) = p(s')$ and $a_Q(s) = a_Q(s')$ for every $Q \in \mathcal{Q}$.

Proposition 8.1 (Target-family completion). *The data $(E_Q, e_Q, u_Q, (\sigma_Q)_{Q \in \mathcal{Q}})$ is the terminal object of the category $\mathbf{Suff}(p; \{a_Q\}_{Q \in \mathcal{Q}})$ of surjective extensions of the readout on which every answer map a_Q becomes well defined. In particular $E_Q \cong S / \sim_Q$, and E_Q is the coarsest extension of the readout determining every answer in the family.*

Proof. The answer maps $a_Q: S \rightarrow A_Q$ form a set-indexed family of functions on S , exactly as the r_α do. Theorems 3.1, 4.1 and 4.2 and Theorems 3.2, 3.4, 5.1 and 5.3 were stated and proved for an arbitrary such family and use no property of the r_α beyond their being functions on S . Substituting $\mathcal{Q}$ for A , a_Q for r_α , and A_Q for V_α throughout yields the assertions verbatim. $\square$

8.1. Mediation and the canonical surjection. Now suppose the nominated relation family is *sufficient for the targets*, in the precise sense that each answer is computable from the readout together with the relation values.

Definition 8.2 (Mediation). The family $\{r_\alpha\}_{\alpha \in A}$ *mediates* $\mathcal{Q}$ if for every $Q \in \mathcal{Q}$ there is a map

$$\Phi_Q: T \times V_A \longrightarrow A_Q \quad \text{with} \quad a_Q(s) = \Phi_Q(p(s), R_A(s)) \quad \text{for all } s \in S. \quad (8.3)$$

Theorem 8.3 (Relation completion versus target completion). *Suppose $\{r_\alpha\}_{\alpha \in A}$ mediates $\mathcal{Q}$. Then there is a unique map*

$$\chi: E_A \longrightarrow E_Q \quad \text{with} \quad e_Q = \chi \circ e_A, \quad (8.4)$$

and χ is surjective; explicitly, $\chi(x) = (u_A(x), (\Phi_Q(x))_{Q \in \mathcal{Q}})$ for $x \in E_A \subseteq T \times V_A$. Moreover, χ is a bijection if and only if

$$\sim_Q = \sim_A, \quad (8.5)$$

that is, if and only if equality of the readout together with all target answers induces exactly the same equivalence relation on S as equality of the readout together with all selected relation values. In general one has only

$$\sim_A \subseteq \sim_Q, \quad (8.6)$$

so the target completion may be **strictly coarser** than the relation completion.

Proof. E_A is sufficient for $\mathcal{Q}$. Define $\sigma_Q^A := \Phi_Q|_{E_A}: E_A \rightarrow A_Q$ (the restriction of Φ_Q to $E_A \subseteq T \times V_A$). For $s \in S$,

$$\sigma_Q^A(e_A(s)) = \Phi_Q(p(s), R_A(s)) = a_Q(s)$$

by (8.3). Together with the identity $p = u_A \circ e_A$ and the surjectivity of the map e_A , this exhibits the tuple $(E_A, e_A, u_A, (\sigma_Q^A)_{Q \in \mathcal{Q}})$ as an object of the category $\mathbf{Suff}(p; \{a_Q\})$.

Existence, uniqueness, surjectivity of χ . By Theorem 8.1, E_Q is terminal in that category, so Theorem 5.1 supplies a unique morphism $\chi: E_A \rightarrow E_Q$, necessarily satisfying $e_Q = \chi \circ e_A$ and necessarily surjective; and by the explicit formula (5.3) it is given on $x \in E_A$ by $\chi(x) = (u_A(x), (\sigma_Q^A(x))_{Q \in \mathcal{Q}}) = (u_A(x), (\Phi_Q(x))_{Q \in \mathcal{Q}})$.

The ordering (8.6). Suppose $s \sim_A s'$. Then $p(s) = p(s')$ and $R_A(s) = R_A(s')$, so by (8.3), $a_Q(s) = \Phi_Q(p(s), R_A(s)) = \Phi_Q(p(s'), R_A(s')) = a_Q(s')$ for every Q . Hence $s \sim_Q s'$. (This is Theorem 5.3 applied to the sufficient extension E_A of the $\mathcal{Q}$ -problem.)

Injectivity criterion. By Theorem 3.4 applied to both families, χ is identified with the canonical map $S / \sim_A \rightarrow S / \sim_Q$, $[s]_A \mapsto [s]_Q$, which is well defined by (8.6). Such a map is injective iff $\sim_Q \subseteq \sim_A$; combined with (8.6) this is exactly (8.5). Since χ is always surjective, it is a bijection precisely under (8.5). $\square$

Remark 8.4 (Two different things are being retained). The distinction encoded in [Theorem 8.3](#) is substantive and is easy to lose. The *relation* completion E_A retains the nominated relation *values* themselves. The *target* completion E_Q retains only those distinctions among relation values that *affect the nominated answers*. If two source states carry different relation signatures but the mediating maps Φ_Q send both signatures to the same answers for every $Q \in \mathcal{Q}$, then E_A separates them and E_Q does not. Hence the strict-coarsening possibility in (8.6) is not an artifact; [Section 11](#) exhibits it concretely, where the relation completion retains an orientation that no answer in the nominated spectral family can detect.

9. THE GAUGE-REDUCTION BOUNDARY

We return to the raw configuration space Ω and the physical quotient (2.1), $q: \Omega \rightarrow S = \Omega / \simeq_{\text{phys}}$. Recall [Theorem 2.1](#): only quantities invariant under $\simeq_{\text{phys}}$ descend to S , and only such quantities are eligible for nomination.

Corollary 9.1 (Completion does not restore gauge coordinates). *For every nominated family $\{r_\alpha\}_{\alpha \in A}$ of eligible relations, the composite $e_A \circ q: \Omega \rightarrow E_A$ is constant on raw physical-equivalence classes:*

$$\omega \simeq_{\text{phys}} \omega' \implies e_A(q(\omega)) = e_A(q(\omega')). \quad (9.1)$$

Proof. If $\omega \simeq_{\text{phys}} \omega'$ then $q(\omega) = q(\omega')$ by definition of the quotient map, and applying the function e_A to a single element of S yields a single value. $\square$

Mathematically this is a triviality, and we present it as such: it is a corollary of the order of operations, not a theorem with content of its own. Its *physical* significance is nevertheless worth stating precisely, because completion is easily misread as a licence to reintroduce structure that was deliberately discarded. The precise content of (9.1) is that the composite $e_A \circ q$ is constant on physical-equivalence classes. Consequently, no completion, however large the nomination, can

- (1) distinguish two gauge-equivalent raw configurations $\omega \simeq_{\text{phys}} \omega'$;
- (2) recover which raw representative ω was originally supplied, given only the completed state;
- (3) restore discarded gauge or coordinate data as *physical* information—that is, as an invariant of the completed state; or
- (4) canonically generate a gauge section $\Omega \leftarrow S$ from the completion theorem.

We are careful not to overstate this. We do *not* claim that a representative of a physical-equivalence class can never be selected in any sense whatever: a separately chosen gauge-fixing convention, or a section of q , may well assign a representative when one exists, and such devices are routine in practice. The point is about provenance and status. A conventional gauge-fixing section is *extra, noncanonical structure supplied from outside*. It is not produced by the completion, is not privileged by terminality, and—this is the substantive claim—is not a recovery of lost raw information: the information it appears to supply was chosen by the convention, not extracted from the physics. Two different conventions give two different representatives of the same physical state, and nothing in E_A adjudicates between them.

Completion therefore refines the readout partition, but only *within* the reduced source S ; it cannot resurrect a distinction that the reduction q was defined to destroy. Reduction precedes completion, and the universal completion does not undo physical equivalence. If one wants such a distinction back as physics, one must argue that the reduction was wrong—i.e. that the quantity was physical after all, hence eligible in the sense of [Theorem 2.1](#)—and that is an argument about physics, not a construction available in this paper.

10. THE TOPOLOGICAL UNIVERSAL PROPERTY

Assume now that S , T , and each V_α are topological spaces, and that p and every r_α are continuous. Give $V_A = \prod_{\alpha \in A} V_\alpha$ the product topology and $T \times V_A$ the product topology; then $j_A = (p, R_A)$ is continuous, by the universal property of the product.

Definition 10.1 (Topological completion). Equip E_A with the *quotient topology* induced by the surjection $e_A: S \twoheadrightarrow E_A$: a subset $U \subseteq E_A$ is open iff $e_A^{-1}(U)$ is open in S . A *topological relation-sufficient extension* is a relation-sufficient extension $\mathcal{E} = (E, e, u, (\rho_\alpha^E))$ in which E is a topological space, e is a *quotient map*, and u and every ρ_α^E are continuous. Morphisms are continuous maps satisfying (4.3).

Proposition 10.2 (Continuity of the structure maps). *With the quotient topology of Theorem 10.1, the maps $u_A: E_A \rightarrow T$ and $\rho_\alpha: E_A \rightarrow V_\alpha$ are continuous, so $\mathcal{E}_A$ is a topological relation-sufficient extension.*

Proof. By the characteristic property of the quotient topology, a map out of E_A is continuous iff its composite with e_A is continuous. Now $u_A \circ e_A = p$ and $\rho_\alpha \circ e_A = r_\alpha$ by (3.4), and p and r_α are continuous by hypothesis. Finally e_A is a quotient map by construction. $\square$

Theorem 10.3 (Topological universal property). *$\mathcal{E}_A$, with the quotient topology, is terminal in the category of topological relation-sufficient extensions: for every such $\mathcal{E}$, the unique set map $\phi_\mathcal{E}: E \rightarrow E_A$ of Theorem 5.1 is continuous, and it is the unique continuous morphism $\mathcal{E} \rightarrow \mathcal{E}_A$. It is a continuous surjection.*

Proof. Theorem 5.1 supplies the unique function $\phi_\mathcal{E}: E \rightarrow E_A$ satisfying $e_A = \phi_\mathcal{E} \circ e$, and asserts its surjectivity. For continuity: the composite $\phi_\mathcal{E} \circ e = e_A$ is continuous (it is the quotient map onto E_A), and e is a quotient map by Theorem 10.1; by the characteristic property of quotient maps, a function out of E whose composite with the quotient map e is continuous is itself continuous. Hence $\phi_\mathcal{E}$ is continuous. Uniqueness among *continuous* morphisms is inherited from uniqueness among all set maps, proved in Theorem 5.1 (v). Finally $\mathcal{E}_A$ is itself an object of the category by Theorem 10.2. $\square$

10.1. Quotient topology versus subspace topology. E_A carries two a priori distinct topologies: the quotient topology of Theorem 10.1, and the subspace topology inherited from $T \times V_A$ via the canonical injection

$$\iota: E_A \hookrightarrow T \times V_A. \quad (10.1)$$

These must not be conflated. The universal property of Theorem 10.3 is a statement about the *quotient* topology, and it is the quotient topology that the completion carries by definition.

Proposition 10.4 (Comparison, and the embedding criterion). *Equip E_A with the quotient topology.*

- (1) *The injection ι of (10.1) is continuous. Consequently the quotient topology is finer than or equal to the subspace topology on $\text{Im}(j_A)$.*
- (2) *ι is a topological embedding if and only if the quotient topology coincides with the subspace topology; equivalently, if and only if the corestriction $j_A: S \rightarrow \text{Im}(j_A)$ is a quotient map when $\text{Im}(j_A)$ carries the product subspace topology.*

Proof. 1 We have $\iota \circ e_A = j_A$, which is continuous; since e_A is a quotient map, ι is continuous by the characteristic property. Continuity of ι says exactly that every subspace-open subset of E_A (i.e. every $\iota^{-1}(W)$ with W open in $T \times V_A$) is quotient-open; hence the quotient topology contains the subspace topology.

2 ι is injective and continuous, so it is an embedding iff it is open onto its image, i.e. iff every quotient-open subset of E_A is subspace-open. Combined with 1, this holds iff the two topologies are equal.

For the reformulation, write $I := \text{Im}(j_A)$ with the subspace topology and let $\bar{j}_A: S \rightarrow I$ be the corestriction of j_A , so that as functions $\bar{j}_A = \iota \circ e_A$ with ι now regarded as the identity of underlying sets $E_A \rightarrow I$. Suppose the topologies agree. Then ι is a homeomorphism and e_A is a quotient map, so $\bar{j}_A$ is a quotient map. Conversely, suppose $\bar{j}_A$ is a quotient map. A subset $U \subseteq E_A$ is quotient-open iff $e_A^{-1}(U)$ is open in S ; and $e_A^{-1}(U) = \bar{j}_A^{-1}(U)$ as subsets of S , since e_A and $\bar{j}_A$ agree as functions into the common underlying set. Because $\bar{j}_A$ is a quotient map, $\bar{j}_A^{-1}(U)$ open in S implies U open in I , i.e. subspace-open. Hence quotient-open implies subspace-open, and with 1 the two topologies coincide. $\square$

Corollary 10.5 (Compact–Hausdorff sufficient condition). *Suppose S is compact and $T \times V_A$ is Hausdorff. Then the quotient and subspace topologies on E_A agree, and $\iota: E_A \hookrightarrow T \times V_A$ is a topological embedding; in particular E_A is a compact Hausdorff space.*

Proof. Write $I = \text{Im}(j_A) \subseteq T \times V_A$ with the subspace topology, and let $\bar{j}_A: S \rightarrow I$ be the (continuous, surjective) corestriction of j_A . As a subspace of a Hausdorff space, I is Hausdorff. We show $\bar{j}_A$ is a closed map, hence a quotient map, and then apply Theorem 10.4.2.

Let $C \subseteq S$ be closed. Since S is compact, C is compact; since $\bar{j}_A$ is continuous, $\bar{j}_A(C)$ is a compact subset of I ; since I is Hausdorff, the compact subset $\bar{j}_A(C)$ is closed in I . Thus $\bar{j}_A$ is a continuous closed surjection, and every continuous closed surjection is a quotient map: if $\bar{j}_A^{-1}(U)$ is open then $S \setminus \bar{j}_A^{-1}(U)$ is closed, its image $\bar{j}_A(S \setminus \bar{j}_A^{-1}(U)) = I \setminus U$ is closed by surjectivity, so U is open. By Theorem 10.4, the topologies agree and ι is an embedding. Finally $E_A \cong I$ is compact (continuous image of a compact space) and Hausdorff (subspace of a Hausdorff space). $\square$

Remark 10.6. Without hypotheses of this kind the two topologies may genuinely differ, and the abstract completion does not require them to agree: Theorem 10.3 holds for the quotient topology regardless. One should not assert the agreement, nor treat E_A as a subspace of $T \times V_A$ in the topological sense, without verifying Theorem 10.4.2 or a sufficient condition such as Theorem 10.5.

11. WORKED EXAMPLE: FLAT $U(1)$ CONNECTIONS ON THE CIRCLE

We now exhibit every phenomenon of the general theory in a single, completely explicit gauge-theoretic example. It is the extreme case of readout loss: the readout is *constant*, so it distinguishes nothing at all, and the completion must supply everything.

11.1. The reduced source and the readout. Consider $U(1)$ connections on the trivial bundle over the circle S^1 , restricted to the flat ones. Gauge reduction (2.1) identifies gauge-equivalent connections; on the circle, the resulting reduced moduli space of flat connections is classified by the holonomy around the loop, so that

$$S_{\text{flat}} \cong U(1), \quad z = e^{i\theta}, \quad \theta \in \mathbb{R}/2\pi\mathbb{Z}. \quad (11.1)$$

This is the standard classification of flat connections on the circle by their holonomy (see, e.g., [5, 7]); we take it as given, and stress that S_{flat} is already the *reduced* source in the sense of Section 2: gauge potentials and gauge transformations have been quotiented away, and it is not available to us to reintroduce them.

The readout is the curvature. For flat connections the curvature vanishes identically, so

$$p_F: S_{\text{flat}} \longrightarrow T = \{0\}, \quad p_F(z) = 0. \quad (11.2)$$

The readout fiber is the whole source: $F_0 = S_{\text{flat}} \cong U(1)$. Every nontrivial distinction on S_{flat} therefore fails to descend along p_F , and $\sim_{p_F}$ is the total relation. Note carefully what is *not* being claimed: we do not claim that curvature is the complete physical observable content of a gauge field. On the contrary, the example exists precisely because it is not—the Aharonov–Bohm effect [6] and the Wu–Yang analysis [7] are the physical statement that holonomy carries information the field strength does not. The readout p_F is a readout, not an ontology.

11.2. Transport completion. Nominate the single relation

$$r_{\text{hol}}: S_{\text{flat}} \rightarrow U(1), \quad r_{\text{hol}}(z) = z, \quad (11.3)$$

the oriented holonomy itself. It is eligible in the sense of [Theorem 2.1](#): holonomy around a fixed loop is gauge-invariant (for the abelian group $U(1)$ it is invariant outright, not merely up to conjugacy), so it is a function on the reduced source. It is nominated because oriented parallel transport is what enters targets involving directed transport of charge around the loop.

By [Theorem 3.1](#), the completion is

$$E_{\text{hol}} = \text{Im}(p_F, r_{\text{hol}}) = \{(0, z) : z \in U(1)\} \cong U(1), \quad (11.4)$$

and $e_{\text{hol}}: S_{\text{flat}} \rightarrow E_{\text{hol}}$ is a bijection. Here S is compact and $T \times V = \{0\} \times U(1)$ is Hausdorff, so [Theorem 10.5](#) applies: the quotient and subspace topologies agree and (11.4) is a homeomorphism. By [Theorem 12.1](#) below, the nomination $\{p_F, r_{\text{hol}}\}$ separates points of S_{flat} , so this is a case in which completion happens to reconstruct the whole reduced source. It does *not* restore gauge potentials or gauge representatives—by [Theorem 9.1](#) it cannot—and it restores nothing beyond the reduced source.

11.3. Spectral completion. Now nominate instead the charged-scalar spectral target of Paper 2 [2]. For a scalar of unit charge on S^1 minimally coupled to the flat connection with holonomy $e^{i\theta}$, let H_θ denote the associated Laplace-type operator. The result we import from Paper 2 is

$$\text{Spec}(H_\theta) = \text{Spec}(H_{\theta'}) \iff \theta' \equiv \pm\theta \pmod{2\pi}, \quad (11.5)$$

with multiplicities. We do not reprove (11.5) here; it is exactly the Paper 2 input, and the new content of the present section is the family-level completion and refinement structure built on top of it.

Take $\mathcal{Q} = \{Q_{\text{spec}}\}$ with answer map $a_{\text{spec}}(z) = \text{Spec}(H_\theta)$ for $z = e^{i\theta}$ (a well-defined function of z , since $\theta \mapsto H_\theta$ depends on θ only mod 2π). By (11.5), the induced equivalence relation on S_{flat} is

$$z \sim_{\text{spec}} z' \iff z' = z \text{ or } z' = z^{-1}, \quad (11.6)$$

the readout clause $p_F(z) = p_F(z')$ being vacuous. Hence, by [Theorem 8.1](#),

$$E_{\text{spec}} \cong S_{\text{flat}} / \sim_{\text{spec}} = U(1) / (z \sim z^{-1}). \quad (11.7)$$

Proposition 11.1 (The spectral completion is an interval). *Give $U(1)/(z \sim z^{-1})$ the quotient topology. Then the assignment*

$$c: U(1)/(z \sim z^{-1}) \longrightarrow [-1, 1], \quad c([z]) = \text{Re } z = \cos \theta \quad (z = e^{i\theta}), \quad (11.8)$$

is a well-defined homeomorphism. Hence $E_{\text{spec}} \cong [-1, 1]$.

Proof. Well defined. $\text{Re}(z^{-1}) = \text{Re}(\bar{z}) = \text{Re } z$ for $z \in U(1)$, so the formula does not depend on the representative of the class.

Injective. Suppose $\cos \theta = \cos \theta'$ with $\theta, \theta' \in \mathbb{R}/2\pi\mathbb{Z}$. Then $\theta' \equiv \pm\theta \pmod{2\pi}$, i.e. $e^{i\theta'} = e^{\pm i\theta}$, i.e. $z' \in \{z, z^{-1}\}$, i.e. $[z'] = [z]$.

Surjective. For $x \in [-1, 1]$ set $\theta = \arccos x \in [0, \pi]$; then $c([e^{i\theta}]) = x$.

Continuous. The composite of c with the quotient map $U(1) \rightarrow U(1)/(z \sim z^{-1})$ is $z \mapsto \operatorname{Re} z$, which is continuous. By the characteristic property of the quotient topology, c is continuous.

Homeomorphism. $U(1)$ is compact, hence so is its continuous image $U(1)/(z \sim z^{-1})$; and $[-1, 1]$ is Hausdorff. A continuous bijection from a compact space to a Hausdorff space is a homeomorphism: it is a closed map, since closed subsets of a compact space are compact, their continuous images are compact, and compact subsets of a Hausdorff space are closed; a closed continuous bijection is a homeomorphism. $\square$

11.4. The refinement map, concretely. The relation family $\{r_{\text{hol}}\}$ mediates the spectral target in the sense of [Theorem 8.2](#): by (11.5) the spectrum is a function of the holonomy, so there is $\Phi_{\text{spec}}: T \times U(1) \rightarrow A_{Q_{\text{spec}}}$ with $a_{\text{spec}}(z) = \Phi_{\text{spec}}(p_F(z), r_{\text{hol}}(z))$. [Theorem 8.3](#) therefore supplies a unique surjection $\chi: E_{\text{hol}} \rightarrow E_{\text{spec}}$, and under the identifications (11.4) and [Theorem 11.1](#) it is exactly

$$\chi: U(1) \longrightarrow [-1, 1], \quad e^{i\theta} \longmapsto \cos \theta. \quad (11.9)$$

It is surjective, and it is *not* injective: $\chi(e^{i\theta}) = \chi(e^{-i\theta})$ while $e^{i\theta} \neq e^{-i\theta}$ whenever $\theta \not\equiv 0, \pi$. So this is a concrete realization of the strict inequality permitted by (8.6): here $\sim_A \subsetneq \sim_Q$, and the target completion is strictly coarser than the relation completion.

The physics of the strictness is transparent. The full oriented holonomy retains an *orientation*: it distinguishes transport of charge around the loop in one sense from transport in the other, i.e. z from z^{-1} . The charged-scalar spectrum does not require that information—by (11.5) it is invariant under $\theta \mapsto -\theta$ —and so the minimal completion *for that target* discards it. Each completion is minimal for its own nomination. Which is appropriate depends on which targets one has nominated, and that is a physical decision made before the mathematics begins, not one the mathematics can make.

Remark 11.2. The two fixed points of the inversion, $z = 1$ ($\theta = 0$) and $z = -1$ ($\theta = \pi$), are exactly the points where χ fails to be two-to-one; they map to the endpoints ± 1 of the interval. They are the connections whose holonomy is its own inverse, i.e. for which orientation carries no information.

12. THE EXACT SCOPE OF RECONSTRUCTION

Completion always reconstructs the joint readout/relation signature; that is what it is. Whether it reconstructs the source state is a separate question with a sharp answer.

Proposition 12.1 (Reconstruction criterion). *The completion map $e_A: S \rightarrow E_A$ is injective—equivalently, $\sim_A = \Delta_S$; equivalently, e_A is a bijection onto E_A —if and only if the family*

$$\{p\} \cup \{r_\alpha : \alpha \in A\} \quad (12.1)$$

separates the points of S : for all $s \neq s'$ in S , either $p(s) \neq p(s')$ or $r_\alpha(s) \neq r_\alpha(s')$ for some $\alpha \in A$.

Proof. By [Theorem 3.1](#) and (3.5), for all $s, s' \in S$,

$$\begin{aligned} e_A(s) = e_A(s') &\iff j_A(s) = j_A(s') \\ &\iff [p(s) = p(s') \text{ and } r_\alpha(s) = r_\alpha(s') \ \forall \alpha \in A] \\ &\iff s \sim_A s'. \end{aligned}$$

Hence e_A is injective iff $(s \sim_A s' \Rightarrow s = s')$, iff $\sim_A = \Delta_S$ (the reverse inclusion always holding by reflexivity), iff no pair of distinct points is identified by the whole family—which is the stated

separation property, in contrapositive form. Surjectivity of e_A is automatic (Theorem 3.2), so injectivity is equivalent to bijectivity. $\square$

Remark 12.2 (Two senses of “reconstruction”; only one is automatic). Distinguish carefully:

- (1) *Reconstruction of the joint signature.* The completed state determines the readout value and every nominated relation value: this is (3.4), and it holds *always*, by construction. It is the entire purpose of the object.
- (2) *Reconstruction of the reduced source state.* The completed state determines $s \in S$ itself. This holds *if and only if* the nominated family separates points, by Theorem 12.1, and in general it fails—necessarily so, since by Theorem 5.3 E_A is the coarsest sufficient extension and therefore retains nothing the nomination does not force it to retain.

Neither sense licenses a third: recovery of *raw* configuration data in Ω , or selection of a preferred representative of a physical equivalence class or of a readout fiber. Theorem 9.1 forbids the former and Theorem 3.7 the latter. The word “reconstruction” should not be used in this subject without saying which of (1) and (2) is meant.

Section 11 contains an instance of each case: the transport nomination $\{p_F, r_{\text{hol}}\}$ separates points of S_{flat} , so e_{hol} is bijective and completion recovers the whole reduced source; the spectral nomination does not separate z from z^{-1} , so the spectral completion is a proper quotient, and no amount of further analysis of the spectrum will recover the orientation, because the orientation is not there to be recovered.

13. DISCUSSION AND SCOPE

13.1. What has been established. Given a reduced source space S , a surjective readout $p: S \rightarrow T$, and a set-indexed family of eligible source relations $\{r_\alpha: S \rightarrow V_\alpha\}_{\alpha \in A}$, there exists a completion $E_A = \text{Im}(p, R_A)$ of the readout on which every nominated relation becomes well defined; it is canonically the quotient $S/\sim_A$; it is terminal among surjective relation-sufficient extensions and therefore unique up to unique isomorphism; it is the *coarsest* such extension, in the exact sense $\sim_e \subseteq \sim_A$ for every sufficient e ; its fibers over T are the realized relation signatures; it is functorial and cocyclic in the nomination; it is insensitive to faithful relabelling of relation values; it admits a canonical surjection onto the completion determined by any family of targets that the relations mediate; it restores no gauge redundancy; and in the topological setting it is terminal for the quotient topology, embedding into $T \times V_A$ exactly when the quotient and subspace topologies agree.

Together with [1] (obstruction) and [2] (target-relative necessity), this supplies the constructive step of the sequence: the object whose existence, uniqueness, and minimality were previously only desiderata.

13.2. The three standing confusions. *Minimality is not smallness.* Theorem 5.5: minimality is least retained distinction, an order-theoretic statement about equivalence relations, not a cardinality bound and not an optimality claim in any other metric.

Completion is not reconstruction. Theorem 12.2: the completed state always determines the nominated signature and determines the source state only under a separation hypothesis that must be checked, not assumed.

Coordinates are relative to the nomination. Theorem 6.4 and Theorem 7.3: the nominated relations are a presentation of the quotient, and irredundancy constrains that presentation without making it unique. Relative to a fixed nominated family, E_A does have canonical structure maps u_A and $(\rho_\alpha)_{\alpha \in A}$; what it lacks is a family-independent, intrinsically unique coordinate

presentation. Different faithful reparameterizations, and different irredundant families, may present the same quotient.

13.3. What has not been established. The theorem is model-relative, and the following are *not* consequences of it.

That the universe is a projection of a hidden source. Nothing here bears on that; S is an assumed structure, and the theorem tells one what follows from assuming it, not that it exists.

That the source space is empirically real. The construction is available whether S is a physical posit, a redundant description, or a bookkeeping device.

That every lost distinction ought to be restored. On the contrary: the entire point of [Theorem 8.3](#) and of [Section 11](#) is that a distinction is worth restoring only relative to a nominated target, and that nominating more than the target requires yields a completion that is finer than necessary—correct, but not minimal for that target.

That the full source can in general be reconstructed. [Theorem 12.1](#) makes this a hypothesis, and the spectral example shows it can fail.

That the nominated relation family is physically privileged. This is the load-bearing external input, and the mathematics cannot supply it. A relation is not physical because it varies on a readout fiber; if it were, every gauge artifact would be an observable. The nomination must be justified by the physics of the targets, and this paper’s theorems are conditional on it in exactly that one respect—which is, we think, where the conditionality belongs.

13.4. Outlook. The present construction is purely set-theoretic and, in [Section 10](#), topological. It says nothing about when E_A carries additional structure—smooth, symplectic, metric—or when the relation signatures arise as genuine geometric invariants rather than as formal coordinates. The companion manuscript in preparation on geometric realization takes up exactly that question, on restricted spaces of metrics, bundles, connections, and matter fields. We have deliberately imported none of that machinery here; the theorems above should stand or fall on their own, and they require nothing beyond images, quotients, and quotient topologies.

14. CONCLUSION

Source completion is neither arbitrary enrichment of the readout nor selection of a preferred state from each readout fiber. It is the canonical quotient of the reduced source that preserves the original readout while separating exactly those source states that the nominated family of invariant relations distinguishes—and no others. The joint-image construction proves existence; the terminal universal property, with its arrow running $E \rightarrow E_A$ from every sufficient extension onto the completion, proves minimality and uniqueness up to unique isomorphism; the indispensability criterion identifies which nominated coordinates do genuine work; the refinement tower organizes nominations by information content; the target-completion surjection separates “what the relations say” from “what the answers need”; the gauge boundary confirms that completion cannot undo reduction; and the topological theorem carries the whole construction into continuous settings without conflating the quotient and subspace topologies. What remains outside the theorem—whether a given source is real, and whether a given nomination is physical—remains outside it deliberately.

REFERENCES

- [1] 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).
- [2] 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).

- [3] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed., Graduate Texts in Mathematics 5, Springer-Verlag, New York, 1998.
- [4] J. R. Munkres, *Topology*, 2nd ed., Prentice Hall, Upper Saddle River, NJ, 2000.
- [5] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry*, Vol. I, Interscience Publishers, New York, 1963.
- [6] Y. Aharonov and D. Bohm, *Significance of electromagnetic potentials in the quantum theory*, Phys. Rev. **115** (1959), 485–491.
- [7] T. T. Wu and C. N. Yang, *Concept of nonintegrable phase factors and global formulation of gauge fields*, Phys. Rev. D **12** (1975), 3845–3857.

INDEPENDENT RESEARCHER