

SOURCE–READOUT NON-EQUIVALENCE: DESCENT AND EQUIVARIANT RECONSTRUCTION OBSTRUCTIONS

JEREMY RODGERS

ABSTRACT. Let Ω be a space of admissible realization states, let $\simeq_{\text{phys}}$ identify descriptions differing only by declared gauge or coordinate redundancy, and let $\Pi: \Omega \rightarrow T$ be a readout map invariant under $\simeq_{\text{phys}}$. After passing to the reduced source space $S = \Omega / \simeq_{\text{phys}}$, the readout induces a surjection $p: S \rightarrow T$. We show that p always presents T exactly as the quotient $S / \sim_p$ of S by equality of readouts, and that this exactness is fully compatible with the failure of p to be equivalent to S .

Two distinct obstructions are established. First, a physically invariant source relation $r: S \rightarrow V$ descends through the readout if and only if it is constant on every fiber of p ; a relation that varies within a fiber is therefore not a function of readout data. Second, if a group G acts on S and T with p equivariant, a G -equivariant section of p exists if and only if every readout stabilizer G_t fixes a point of its fiber F_t ; a fiber with no stabilizer-fixed point admits no symmetry-compatible deterministic representative. A finite marked-carrier occupation model and the Hopf fibration $h: S^3 \rightarrow S^2$ realize these obstructions in the discrete and continuous settings respectively, and the Hopf example additionally separates the equivariant obstruction from the independent topological obstruction furnished by a nonvanishing first Chern class.

The results separate three statements that are frequently conflated: exactness of a quotient presentation, existence of a selected representative, and equivalence with the realization source. The first may hold while the second and third fail.

1. INTRODUCTION

A recurring pattern in mathematical physics is that the quantities a theory treats as public data are obtained from a richer underlying description by a map that forgets structure. Occupation numbers forget which particle sits where; a density matrix forgets a purification; a coarse-grained field forgets microscopic detail; a gauge-invariant observable forgets a gauge frame. In such situations one wants to know precisely what has been lost, and in what sense, if any, it can be recovered.

The last of these examples is deliberately of a different character from the others, and the difference is the hinge on which everything below turns. A gauge frame is, by the declaration of the theory that employs it, not a feature of the physical situation at all, and a map that forgets it has forgotten a redundancy rather than a structure. Whether forgotten data is physically meaningful or merely redundant is not something the forgetting map can decide; it must be decided in advance, by the theory, and the obstruction framework developed here applies only after that decision has been made. The same piece of mathematics will be shown, in Theorem 7.4 and again in Theorem 9.6, to support opposite physical readings depending on which way the decision goes.

The question is easy to pose badly. Three different assertions are routinely run together.

- *Exact quotient presentation.* The readout represents, without error, the equivalence classes of source states that share a readout value.

Date: June 2026.

Website: everythingequation.com.

- *Representative selection.* From each readout value one may choose a source state producing it.
- *Source equivalence.* The readout determines the source state uniquely, so that source and readout carry the same information.

The first of these is essentially automatic once the codomain is restricted to the image of the readout. The second may hold or fail depending on the admissibility conditions imposed on the choice. The third is a strong statement and rarely holds when the readout is genuinely coarser than the source. Nothing about the first two implies the third, and yet the language of “exactness”, “faithful representation”, and “canonical choice” invites exactly that inference.

There is also a prior modeling step that is easy to skip. Not every difference between two descriptions of a physical situation is a physical difference. A theory declares certain transformations redundant: gauge transformations, coordinate changes, relabelings that do not affect any admissible measurement or any admissible law. Differences removed by such declarations are not candidates for lost realization structure, and a map that forgets them has forgotten nothing physical. Consequently, any claim of information loss must be made *after* redundancy has been quotiented out, not before. Failing to do this produces spurious “obstructions” that are artifacts of an over-fine description.

Contribution. This paper carries out the analysis in the order the logic demands. We begin with a realization space Ω , a declared physical equivalence $\simeq_{\text{phys}}$, and the reduced source space $S = \Omega / \simeq_{\text{phys}}$. Only then do we compare S with the readout target T , via the induced surjection $p: S \rightarrow T$.

Within this setting we establish the following.

- (1) The readout is an exact presentation of the quotient of S by readout equivalence: $T \cong S / \sim_p$ (Theorem 3.1). It is source-equivalent, in the set-theoretic sense, exactly when every reduced fiber is a singleton (Theorem 3.2).
- (2) A physically invariant source relation $r: S \rightarrow V$ factors through the readout if and only if it is constant on each fiber of p , and the factorization is then unique (Theorem 4.2). If two reduced source states share a readout but differ in r , then no function of the readout alone reproduces r on S (Theorem 4.3).
- (3) Relative to a family $\mathcal{R}$ of source relations that separates the points of S , completeness of the readout for $\mathcal{R}$ is equivalent to bijectivity of p (Theorem 4.5).
- (4) A section $\sigma: T \rightarrow S$ satisfies $p \circ \sigma = \text{id}_T$ by definition, but satisfies $\sigma \circ p = \text{id}_S$ if and only if p is bijective (Theorem 5.2). Representative selection is therefore not source reconstruction.
- (5) If a group G acts on S and T with p equivariant, then a G -equivariant section exists if and only if each readout stabilizer G_t fixes a point of its fiber (Theorem 6.3). A single fiber with $\text{Fix}_{G_t}(F_t) = \emptyset$ suffices to exclude every equivariant section (Theorem 6.6), and this happens whenever G_t acts transitively on a fiber with more than one point (Theorem 6.7).
- (6) A finite marked-carrier occupation model realizes the equivariant obstruction with $G = S_N$ (Section 7), while remaining explicit about the modeling boundary: the carrier marks are declared physically meaningful, and the obstruction disappears the moment they are declared redundant instead.
- (7) An invariant probability-valued assignment on each fiber does exist in the finite model (Theorem 8.2), which shows that the equivariant obstruction is an obstruction to *deterministic* selection and not to invariant ensembles. An invariant ensemble is not a recovery of the realized source state.

- (8) The Hopf fibration realizes the obstruction in the continuous setting, and does so twice over, by two logically independent mechanisms: a free and transitive fiber action, which excludes equivariant selection; and a nonzero first Chern class, which excludes every continuous global section (Section 9). Neither argument is needed for the other.
- (9) The combined source-readout non-equivalence theorem (Theorem 10.1) asserts exactly what its hypotheses support: under either obstruction, p is noninjective and the corresponding form of reconstruction fails, while T remains the exact quotient $S/\sim_p$.

The mathematics used is elementary in its ingredients — quotients, group actions, stabilizers, principal bundles, characteristic classes — and the individual facts are standard [1, 2, 5, 4, 3, 6, 7, 8, 9]. What is not standard, and what this paper is for, is the discipline of keeping the levels apart: physical equivalence before comparison; descent of relations distinct from reconstruction of states; selection distinct from inversion; and set-theoretic, equivariant, continuous, smooth, measurable, local, and covariant reconstruction requirements each treated on its own terms.

Scope. The theorems below are theorems about maps of sets, equivariant maps, and bundles. They do not, on their own, establish that any particular physical theory has a hidden source layer. Applying them to a physical theory requires that the theory independently supply: a reduced source domain, an explicit physical-equivalence relation, an explicit readout map, an explicit statement of which fiber distinctions are physically meaningful, and either a non-descending relation or an applicable reconstruction obstruction. Section 11 makes this conditional structure precise for the research program in which the present paper is the first of a sequence. Sections 12 and 13 discuss the resulting picture and what it does and does not license.

2. REALIZATION STATES, PHYSICAL EQUIVALENCE, AND READOUT

2.1. The realization space and its reduction. Let Ω be a set of *admissible realization states*. Depending on the application, an element of Ω may be a microscopic configuration, a field configuration, a geometric structure, a solution together with boundary data, or a complete realization package specified by the theory in question. Nothing below depends on which of these it is; what matters is that the theory has committed to Ω as its description of what there is.

A theory does not, however, treat every distinction available in its own description as a physical distinction. Let $\simeq_{\text{phys}}$ be the equivalence relation on Ω generated by all transformations that the theory declares physically redundant. Typically this includes gauge transformations and coordinate changes, in the cases where such transformations alter only the representation of a state and not the state itself. The declaration is part of the theory's content, not a mathematical consequence of it; the mathematics below takes it as given.

Definition 2.1 (Reduced source space). The *reduced source space* is the quotient

$$S := \Omega / \simeq_{\text{phys}}, \quad q: \Omega \longrightarrow S, \tag{1}$$

with q the canonical projection sending a realization state to its physical-equivalence class.

The passage from Ω to S is not a technical convenience; it is the step that makes the rest of the paper meaningful. A difference between two elements of Ω that is removed by $\simeq_{\text{phys}}$ is, by the theory's own declaration, not a difference in the world. If a readout map forgets such a difference, it has forgotten a redundancy, not a physical structure, and no obstruction worth the name has been exhibited. Every claim of lost source structure made below is therefore a claim about S , never about Ω .

2.2. The readout and its induced map. Let

$$\Pi: \Omega \longrightarrow T \quad (2)$$

be the *readout map*, whose values are the data the theory treats as public. Throughout we adopt the convention

$$T = \text{Im } \Pi, \quad (3)$$

so that Π is surjective onto its codomain. This is a normalization, not an assumption: readout values that no realization state produces play no role, and retaining them would only clutter statements about fibers with vacuous cases.

We assume the readout is *physically invariant*, that is, invariant under physical equivalence:

$$\omega \simeq_{\text{phys}} \omega' \implies \Pi(\omega) = \Pi(\omega'). \quad (4)$$

This is the minimal coherence requirement on a readout: if two descriptions describe the same physical situation, they must yield the same public data. A map violating (4) would let one read off gauge-dependent information, which the theory has declared meaningless.

By the universal property of the quotient (1), condition (4) yields a unique map

$$p: S \longrightarrow T, \quad \Pi = p \circ q, \quad (5)$$

and p is surjective because Π is. From this point onward, *all* source–readout comparisons are made between S and T by means of p . Gauge and coordinate redundancy have been removed before any question of non-equivalence is raised.

Definition 2.2 (Reduced fibers and readout equivalence). For $t \in T$, the *reduced source fiber* over t is

$$F_t := p^{-1}(t) \subseteq S. \quad (6)$$

Readout equivalence on S is the relation

$$s \sim_p s' \iff p(s) = p(s'). \quad (7)$$

By (3) every fiber F_t is nonempty, and the fibers partition S into precisely the $\sim_p$ -classes.

It is worth fixing the vocabulary that the rest of the paper will keep separate. *Equality* in Ω is identity of descriptions. *Physical equivalence* $\simeq_{\text{phys}}$ is the theory's declared redundancy. *Readout equivalence* $\sim_p$ is indistinguishability by public data. These are three different relations, and they are related by inclusion in one direction only: equality refines $\simeq_{\text{phys}}$, and $\simeq_{\text{phys}}$ (transported to S as equality) refines $\sim_p$. The content of this paper lies in the second refinement being strict, and in what that strictness does and does not entail.

3. QUOTIENT EXACTNESS AND SOURCE EQUIVALENCE

The first two results are elementary, and stating them explicitly is the point: they are what separates a correct claim about exactness from an incorrect claim about equivalence.

Proposition 3.1 (Exact quotient presentation). *The map*

$$\widehat{p}: S/\sim_p \longrightarrow T, \quad \widehat{p}([s]_p) = p(s), \quad (8)$$

is a well-defined bijection. Hence

$$T \cong S/\sim_p. \quad (9)$$

Proof. Well-definedness: if $[s]_p = [s']_p$ then $s \sim_p s'$, so $p(s) = p(s')$ by (7), and the prescription (8) does not depend on the chosen representative.

Surjectivity: given $t \in T$, the convention (3) supplies $s \in S$ with $p(s) = t$, whence $\widehat{p}([s]_p) = t$.

Injectivity: if $\widehat{p}([s]_p) = \widehat{p}([s']_p)$ then $p(s) = p(s')$, so $s \sim_p s'$ by (7), and therefore $[s]_p = [s']_p$. $\square$

Theorem 3.1 says exactly what exactness of a quotient presentation means, and no more: the readout labels the readout-equivalence classes of S without error and without omission. Every $\sim_p$ -class has exactly one label, and every label names exactly one class.

It is worth being precise about the scope of this, because it is easy to ask it to do more work than it can. Theorem 3.1 produces a *bijection of sets* and nothing further. What it licenses is transport: any structure that is *already well defined on the quotient* $S/\sim_p$ — that is, any structure whose defining data is constant on $\sim_p$ -classes — can be carried across $\hat{p}$ to T without loss. What it does not license is the assumption that some nominated structure on S is of that kind. Whether a given relation, law, or evolution is well defined on the quotient is precisely the descent question of Section 4, and it must be settled by Theorem 4.2 or its analogue, not by Theorem 3.1.

In particular, Theorem 3.1 does *not* establish that the readout level is lawful, internally consistent as a dynamical theory, or dynamically closed. It establishes none of these, and asserting them requires independent descent or closure results for the specific laws and evolutions at issue. The readout level may well be lawful and dynamically coherent — nothing here argues otherwise — but that would be a separate theorem with separate hypotheses. Whether projected dynamics descends, and under what conditions closure fails, are questions this paper deliberately leaves open and does not prejudge.

What (9) also does not say is that the classes it labels are singletons. Exactness is a statement about the relationship between T and $S/\sim_p$; source equivalence would be a statement about the relationship between T and S . Conflating them amounts to assuming that $S/\sim_p$ and S are the same object, which is precisely the question at issue.

Proposition 3.2 (Criterion for source equivalence). *The induced readout $p: S \rightarrow T$ is bijective if and only if every reduced source fiber F_t , $t \in T$, is a singleton.*

Proof. Surjectivity holds by the convention (3), so bijectivity is equivalent to injectivity. Now p is injective if and only if $p(s) = p(s')$ implies $s = s'$, that is, if and only if no fiber contains two distinct points. Since every fiber is nonempty by surjectivity, this is equivalent to every fiber being a singleton. $\square$

Combining the two propositions: a single fiber with more than one point already suffices to make exact quotient presentation and source equivalence come apart. Theorem 3.1 continues to hold verbatim, while Theorem 3.2 fails. What remains to be determined — and what occupies the rest of the paper — is *what* is lost across a nontrivial fiber, and under what admissibility conditions it can or cannot be restored.

We record the terminology to be used for the strongest of the three notions.

Definition 3.3 (Source equivalence). We say that the readout is *source-equivalent* to the reduced source, and write $T \simeq_{\text{source}} S$, when p is a bijection. When p is not a bijection we write $T \not\simeq_{\text{source}} S$.

The notation is deliberately distinguished from the isomorphism symbol in (9), which is a bijection of T with a quotient of S , and from $\simeq_{\text{phys}}$, which is a relation on Ω . It is also, at this stage, a purely set-theoretic notion; Sections 5, 6 and 9 will show that stronger admissibility conditions on reconstruction yield strictly stronger obstructions, and that a map may fail to be reconstructible in one of those senses while succeeding in another.

4. DESCENT OF SOURCE RELATIONS

Nontriviality of a fiber is a statement about the map p . To convert it into a statement about lost *physics*, one must exhibit a physically meaningful quantity that varies across a fiber. This section makes that precise.

4.1. Relations and descent. Let V be a set of values and let

$$r: S \longrightarrow V \quad (10)$$

be a *source relation*: a map assigning to each reduced source state a value that the theory regards as physically meaningful. Because the domain of r is S rather than Ω , the relation is automatically invariant under physical equivalence; this is the payoff of having performed the reduction first. Concrete instances that arise in practice include a global transport class, a boundary relation, a correlation functional, a spectral invariant, or a source property required as an input by some further construction. Nothing in what follows depends on which; r is an arbitrary map out of S .

Definition 4.1 (Descent). The relation r *descends through the readout* if there exists a map

$$\bar{r}: T \longrightarrow V \quad \text{with} \quad r = \bar{r} \circ p. \quad (11)$$

If r descends, then $r(s)$ is a function of the readout value $p(s)$ alone, and an agent with access only to public data can evaluate r . If r does not descend, no such evaluation is possible, however the agent proceeds. It is only after descent has been established that it is legitimate to call r an *observable* of the readout; a source relation is not observable at the readout level merely because it is physically meaningful at the source level.

Theorem 4.2 (Fiber-descent theorem). *Let $r: S \rightarrow V$. The following are equivalent.*

- (1) *There exists $\bar{r}: T \rightarrow V$ with $r = \bar{r} \circ p$.*
- (2) *r is constant on every fiber of p ; that is,*

$$p(s) = p(s') \implies r(s) = r(s') \quad (s, s' \in S). \quad (12)$$

When these hold, $\bar{r}$ is unique.

Proof. **1 $\Rightarrow$ 2.** Suppose $r = \bar{r} \circ p$ and $p(s) = p(s')$. Then

$$r(s) = \bar{r}(p(s)) = \bar{r}(p(s')) = r(s').$$

2 $\Rightarrow$ 1. Suppose r is constant on fibers. Fix $t \in T$. By surjectivity of p the fiber F_t is nonempty, so we may pick $s \in F_t$ and set

$$\bar{r}(t) := r(s). \quad (13)$$

This prescription is independent of the choice of $s \in F_t$: any two elements of F_t have the same image under p , namely t , so (12) forces them to have the same r -value. Hence $\bar{r}$ is a well-defined map $T \rightarrow V$. For arbitrary $s \in S$, applying (13) at $t = p(s)$ with the representative s itself gives $\bar{r}(p(s)) = r(s)$, that is, $r = \bar{r} \circ p$.

Uniqueness. If $\bar{r}_1 \circ p = \bar{r}_2 \circ p$ then, since p is surjective, every $t \in T$ is of the form $p(s)$, and $\bar{r}_1(t) = r(s) = \bar{r}_2(t)$. Hence $\bar{r}_1 = \bar{r}_2$. $\square$

Note that the proof of **2 $\Rightarrow$ 1** requires only a choice of one representative per fiber *for the purpose of evaluating a constant*, and the constancy hypothesis makes the choice immaterial. No choice principle is needed, and no representative survives into the conclusion. This is in sharp contrast to the selection problems of Sections 5 and 6, where the chosen representative *is* the output.

Corollary 4.3 (Descent obstruction). *Suppose there exist $s_0, s_1 \in S$ with*

$$p(s_0) = p(s_1), \quad r(s_0) \neq r(s_1). \quad (14)$$

Then r does not descend through p : there is no map $\bar{r}: T \rightarrow V$ with $r = \bar{r} \circ p$. Consequently no function of the readout alone reproduces r on all of S .

Proof. Condition (14) exhibits a fiber on which r is not constant, so 2 of Theorem 4.2 fails, and therefore 1 fails. $\square$

The content of Theorem 4.3 should not be mistaken for a restatement of noninjectivity of p . Noninjectivity says that the fibers are not all singletons. Theorem 4.3 says something further and more specific: that a *nominated* physically meaningful quantity, one the theory actually cares about, is not among the quantities the readout carries. The hypothesis (14) is a statement one must earn by producing the relation r ; it does not follow from coarseness of the readout by itself, since a noninjective readout may perfectly well transmit every relation the theory needs. This is the difference between observing that a map forgets something and identifying what it forgets.

4.2. Relative completeness. The preceding suggests grading the readout not absolutely but against a nominated family of relations.

Definition 4.4 (Relative source completeness; separation). Let $\mathcal{R}$ be a family of source relations on S (with values in possibly different value sets). The readout is *complete relative to $\mathcal{R}$* if every $r \in \mathcal{R}$ descends through p .

The family $\mathcal{R}$ *separates the points of S* if for all $s \neq s'$ in S there exists $r \in \mathcal{R}$ with $r(s) \neq r(s')$.

Theorem 4.5 (Separating-family criterion). *Let $\mathcal{R}$ separate the points of S . Then the readout is complete relative to $\mathcal{R}$ if and only if p is bijective.*

Proof. ($\Leftarrow$) If p is bijective, then for any $r: S \rightarrow V$ the map $\bar{r} := r \circ p^{-1}$ satisfies $\bar{r} \circ p = r$. Hence every relation whatsoever descends, in particular every member of $\mathcal{R}$.

($\Rightarrow$) Suppose every $r \in \mathcal{R}$ descends, and let $s, s' \in S$ satisfy $p(s) = p(s')$. By Theorem 4.2, each $r \in \mathcal{R}$ is constant on the fibers of p , so $r(s) = r(s')$ for every $r \in \mathcal{R}$. Since $\mathcal{R}$ separates the points of S , this forces $s = s'$. Hence p is injective, and it is surjective by (3). $\square$

Theorem 4.5 gives relative completeness an operational reading, but the reading is conditional and the condition must be checked. The theorem says that a readout is source-complete exactly when it carries a family of relations that separates source states. Whether that reading is *operational* depends entirely on whether the separating family $\mathcal{R}$ has been identified as physically meaningful. An arbitrary separating family always exists — take $\mathcal{R}$ to be all maps $S \rightarrow \{0, 1\}$, or all indicator functions of points — and its existence tells one nothing, since its members need not correspond to anything a theory recognizes or an experiment can access. The theorem earns its interpretation only when $\mathcal{R}$ is supplied by the physics.

5. REPRESENTATIVE SELECTION IS NOT RECONSTRUCTION

Noninjectivity of p excludes recovery of the realized source state by a map of readout data. It does not, by itself, exclude the weaker operation of choosing, for each readout value, some source state that would have produced it. The two are routinely conflated, and the distinction matters.

Definition 5.1 (Section). A *section* of p is a map

$$\sigma: T \longrightarrow S \quad \text{with} \quad p \circ \sigma = \text{id}_T. \quad (15)$$

Equivalently, $\sigma(t) \in F_t$ for every $t \in T$: a section picks one element from each fiber. When T is a set with no further structure and no admissibility conditions are imposed, a section exists whenever a choice function for the family $\{F_t\}_{t \in T}$ of nonempty sets exists; the axiom of choice for that family supplies one, and for finite or countable T weaker principles suffice. Mere existence in this sense therefore conveys no information about p at all. The substantive question is never whether *some* section exists, but whether one exists that is compatible with the structure the problem carries: equivariance, continuity, smoothness, measurability, locality, covariance. Sections 6 and 9 take up two such conditions.

What a section is *not*, in any case, is an inverse.

Proposition 5.2 (Selection is not reconstruction). *Let σ be a section of p . Then*

$$\sigma \circ p = \text{id}_S \quad (16)$$

holds if and only if p is bijective. In that case $\sigma = p^{-1}$.

Proof. Suppose (16) holds and $p(s) = p(s')$. Applying σ ,

$$s = \sigma(p(s)) = \sigma(p(s')) = s',$$

so p is injective; it is surjective by (3), hence bijective. Moreover (15) and (16) then identify σ as the two-sided inverse of p .

Conversely, if p is bijective then p^{-1} exists, satisfies $p \circ p^{-1} = \text{id}_T$, and is therefore a section; and any section σ satisfies, for each $s \in S$, that $\sigma(p(s))$ lies in $F_{p(s)} = \{s\}$, so $\sigma \circ p = \text{id}_S$. $\square$

The asymmetry of the two identities (15) and (16) is the whole point. The first says: whatever source state σ nominates, it does produce the observed readout. The second says: the source state σ nominates is the one that was actually there. The first is a consistency condition on a convention; the second is a claim of recovery. A section supplies a representative convention; only a two-sided inverse supplies source reconstruction, and by Theorem 5.2 a two-sided inverse exists precisely in the degenerate case in which nothing was lost.

Nor can the gap be closed by refining the convention. If $|F_t| > 1$, then for every section σ there exists $s \in F_t$ with $\sigma(p(s)) = \sigma(t) \neq s$: the section misidentifies at least one source state in that fiber, and no amount of cleverness in choosing σ changes this, because the readout data available to σ is identical for all elements of F_t .

6. EQUIVARIANT RECONSTRUCTION

In structured settings one typically wishes to restrict attention to sections compatible with the structure at hand, and the results below concern one such restriction. Suppose the theory carries a symmetry acting on both source and readout and respected by the readout map. A selection rule that fails to respect that symmetry introduces a preference among source states that the symmetry does not distinguish, and that preference has to come from somewhere.

The status of this restriction should be stated carefully, because the theorem below is conditional and its interpretation must be too. If G -equivariance is imposed as an admissibility requirement on reconstruction, then failure of the stabilizer-fixed-point condition excludes every selector in that class — which is a strong and unconditional mathematical conclusion *within* the class. It is not a claim that a non-equivariant selector is inherently unlawful or physically impossible. A non-equivariant selection rule may be perfectly legitimate once it is justified by

something further: symmetry-breaking data, boundary conditions, an environment, additional structure on S , or an openly stipulated convention.

The distinction must be drawn carefully, since it is easy to overstate. A non-equivariant section is, formally, a map $T \rightarrow S$ like any other: it takes a readout value as input and returns a source state, and in that sense it is certainly a function of the readout variable. What it is not is *determined* by the readout and the declared G -symmetric structure alone. Its particular choice among the many available maps $T \rightarrow S$ cannot be selected, derived, or justified from those ingredients; it encodes an additional convention or symmetry-breaking ingredient supplied from outside them, and different conventions yield different and equally admissible non-equivariant sections with no symmetric ground for preferring one. The obstruction is therefore an obstruction relative to a declared reconstruction class, and the class must be declared.

This section characterizes exactly when a G -equivariant selection rule exists.

6.1. Setup. Let a group G act on S and on T , and assume p is *equivariant*:

$$p(g \cdot s) = g \cdot p(s) \quad \text{for all } g \in G, s \in S. \quad (17)$$

Equivariance is the statement that the readout commutes with the symmetry: transforming the source and then reading out is the same as reading out and then transforming. Note that (17) implies

$$g \cdot F_t = F_{g \cdot t} \quad (g \in G, t \in T), \quad (18)$$

since $s \in F_t$ gives $p(g \cdot s) = g \cdot t$, and the reverse inclusion follows by applying g^{-1} .

Definition 6.1 (Stabilizer; fixed set of a fiber). For $t \in T$ let

$$G_t := \{g \in G : g \cdot t = t\} \quad (19)$$

be the stabilizer of t . By (18), G_t maps F_t to itself and hence acts on F_t . Set

$$\text{Fix}_{G_t}(F_t) := \{s \in F_t : g \cdot s = s \text{ for every } g \in G_t\}. \quad (20)$$

The set (20) has a direct reading. The subgroup G_t consists of exactly those symmetries that the readout cannot see at t : they move source states around, possibly, but they leave the readout value t where it is. An element of $\text{Fix}_{G_t}(F_t)$ is a source state over t that is unmoved even by these readout-invisible symmetries: a *symmetry-compatible candidate representative* over t . It is important not to overstate what such a point provides. The set (20) may contain many elements, and when it does, the symmetry itself does not adjudicate among them; a further choice is required, and that choice is not supplied by G . A nonempty fixed set therefore certifies that symmetry-compatible selection is *possible* over t ; it does not by itself single out a point, and it does not make any selected point canonical. When no such point exists, the readout-invisible symmetry group acts on F_t without fixed points, and every element of F_t is carried to a different element by some symmetry that the readout cannot detect. Selecting one of them is then necessarily an act of breaking the symmetry by fiat.

Definition 6.2 (Equivariant section). An *equivariant section* of p is a section $\sigma : T \rightarrow S$ satisfying

$$\sigma(g \cdot t) = g \cdot \sigma(t) \quad \text{for all } g \in G, t \in T. \quad (21)$$

6.2. The criterion.

Theorem 6.3 (Equivariant-section criterion). *Let G act on S and T with p equivariant. For the implication $3 \Rightarrow 1$ assume, in addition, a choice function for the family of G -orbits in T and a choice function for the associated family of nonempty fixed sets $\text{Fix}_{G_{t_i}}(F_{t_i})$ indexed by the selected orbit representatives (see Theorem 6.4). The following are equivalent.*

- (1) An equivariant section $\sigma: T \rightarrow S$ exists.
- (2) $\text{Fix}_{G_t}(F_t) \neq \emptyset$ for every $t \in T$.
- (3) Every G -orbit in T contains at least one representative t with $\text{Fix}_{G_t}(F_t) \neq \emptyset$.

Proof. $1 \Rightarrow 2$. Let σ be an equivariant section and let $t \in T$. Since $p(\sigma(t)) = t$, we have $\sigma(t) \in F_t$. For $g \in G_t$, equivariance (21) and $g \cdot t = t$ give

$$g \cdot \sigma(t) = \sigma(g \cdot t) = \sigma(t).$$

Hence $\sigma(t)$ is fixed by every element of G_t , so $\sigma(t) \in \text{Fix}_{G_t}(F_t)$, which is therefore nonempty.

$2 \Rightarrow 3$. Immediate, since every orbit is nonempty.

$3 \Rightarrow 2$. This is the conjugacy step. Let $t \in T$ be arbitrary; by hypothesis its orbit contains a representative t_0 with some $s_0 \in \text{Fix}_{G_{t_0}}(F_{t_0})$, and there is $g \in G$ with $t = g \cdot t_0$. Stabilizers along an orbit are conjugate:

$$G_{g \cdot t_0} = g G_{t_0} g^{-1}, \quad (22)$$

since $h \cdot (g \cdot t_0) = g \cdot t_0$ if and only if $(g^{-1}h) \cdot t_0 = t_0$. We claim $g \cdot s_0 \in \text{Fix}_{G_t}(F_t)$. First, $g \cdot s_0 \in F_t$ by (18). Second, let $h \in G_t$. By (22) we may write $h = g k g^{-1}$ with $k \in G_{t_0}$, and then

$$h \cdot (g \cdot s_0) = (g k g^{-1} g) \cdot s_0 = g \cdot (k \cdot s_0) = g \cdot s_0,$$

using $k \cdot s_0 = s_0$. Hence $\text{Fix}_{G_t}(F_t) \neq \emptyset$. Thus the fixed-point condition propagates from any orbit representative to the whole orbit, and 2 and 3 are equivalent conditions.

$3 \Rightarrow 1$. Choose a set $\{t_i\}_{i \in I}$ containing exactly one representative from each G -orbit in T , and for each i choose

$$s_i \in \text{Fix}_{G_{t_i}}(F_{t_i}),$$

which is possible by hypothesis. Every $t \in T$ lies in exactly one orbit, so $t = g \cdot t_i$ for a unique $i \in I$ and some $g \in G$. Define

$$\sigma(t) := g \cdot s_i. \quad (23)$$

Well-definedness. The index i is determined by t , but g need not be. Suppose $g \cdot t_i = g' \cdot t_i = t$. Then $(g'^{-1}g) \cdot t_i = t_i$, so $g'^{-1}g \in G_{t_i}$, and since s_i is fixed by G_{t_i} ,

$$g \cdot s_i = g' \cdot ((g'^{-1}g) \cdot s_i) = g' \cdot s_i.$$

Hence (23) does not depend on the choice of g , and σ is a well-defined map $T \rightarrow S$. It is precisely the fixed-point property of s_i that makes this work; an arbitrary choice of $s_i \in F_{t_i}$ would not.

Section property. With $t = g \cdot t_i$, equivariance (17) of p and $s_i \in F_{t_i}$ give

$$p(\sigma(t)) = p(g \cdot s_i) = g \cdot p(s_i) = g \cdot t_i = t.$$

Equivariance. Let $h \in G$ and $t = g \cdot t_i$. Then $h \cdot t = (hg) \cdot t_i$ lies in the same orbit, so by (23),

$$\sigma(h \cdot t) = (hg) \cdot s_i = h \cdot (g \cdot s_i) = h \cdot \sigma(t).$$

Thus σ is an equivariant section, establishing 1. $\square$

Remark 6.4 (Where choice enters). Two independent selections were made in the proof of $3 \Rightarrow 1$, and *both* require a choice assumption; neither one suffices alone.

The first is the selection of one orbit representative t_i from each G -orbit in T : a choice function for the family $\{\mathcal{O}\}_{\mathcal{O} \in T/G}$ of nonempty orbits. The second is the selection of one point s_i from each nonempty fixed set $\text{Fix}_{G_{t_i}}(F_{t_i})$: a choice function for the family $\{\text{Fix}_{G_{t_i}}(F_{t_i})\}_{i \in I}$ indexed by the representatives produced by the first selection. The second family is not the first, and hypothesizing choice for the orbits alone leaves the fixed points unselected. This is exactly the point recorded in Theorem 6.5: a nonempty $\text{Fix}_{G_{t_i}}(F_{t_i})$ may have more than one

element, and nothing in the hypotheses distinguishes one of them. Both choice functions are therefore needed, and it is their conjunction that the theorem assumes.

Neither assumption is needed when T/G is finite and each relevant fixed set is finite, since finite choice is a theorem; nor when the problem independently supplies canonical orbit representatives and canonical fixed points, in which case the selections are definable rather than chosen.

The obstruction direction is a different matter. The implication $1 \Rightarrow 2$ constructs nothing: it merely evaluates a given σ at a given t and observes that $\sigma(t)$ is fixed. It uses no choice principle whatever, and consequently Theorems 6.6 and 6.7, Theorem 7.3, Theorem 9.1, and the equivariant half of Theorem 10.1 — that is, every negative result in this paper — are choice-free.

Remark 6.5 (Existence, uniqueness, and canonicity are three things). Theorem 6.3 characterizes the *existence* of an equivariant section. It characterizes nothing else, and three statements should be kept apart.

Existence. An equivariant section exists if and only if every stabilizer fixes a point of its fiber (together with the choice assumptions of Theorem 6.4). This is the theorem.

Uniqueness. An equivariant section need not be unique. If some $\text{Fix}_{G_t}(F_t)$ contains two distinct points $s \neq s'$, then the construction (23) run with s and with s' produces two equivariant sections differing at t , and both are equally valid. Nonempty fixed sets certify possibility, not uniqueness.

Canonicity. A fortiori, an element of $\text{Fix}_{G_t}(F_t)$ is not thereby a canonical representative. It is a representative that the symmetry *permits*; the symmetry has simply declined to rule it out. Selecting among several permitted representatives requires a principle external to G , and any such principle is additional structure, not a consequence of the symmetry.

The theorem should therefore be read as delimiting a reconstruction class, not as producing a preferred reconstruction within it. Where the fixed set happens to be a singleton for every t , the equivariant section is unique and is canonical relative to G ; but that is an extra hypothesis, and it is not what Theorem 6.3 asserts.

6.3. Consequences.

Corollary 6.6 (No-selector theorem). *If there exists $t_0 \in T$ with*

$$\text{Fix}_{G_{t_0}}(F_{t_0}) = \emptyset, \tag{24}$$

then no G -equivariant section of p exists.

Proof. This is the contrapositive of $1 \Rightarrow 2$ in Theorem 6.3, which requires no choice principle. $\square$

Corollary 6.7 (Transitive nontrivial fiber). *Suppose that for some $t_0 \in T$ the stabilizer G_{t_0} acts transitively on F_{t_0} and $|F_{t_0}| > 1$. Then no G -equivariant section of p exists.*

Proof. Suppose some $s \in F_{t_0}$ were fixed by every element of G_{t_0} . Then the G_{t_0} -orbit of s inside F_{t_0} is the singleton $\{s\}$. But G_{t_0} acts transitively on F_{t_0} , so that orbit is all of F_{t_0} , giving $F_{t_0} = \{s\}$ and contradicting $|F_{t_0}| > 1$. Hence $\text{Fix}_{G_{t_0}}(F_{t_0}) = \emptyset$, and Theorem 6.6 applies. $\square$

Theorem 6.7 isolates the mechanism in its cleanest form. If the readout-invisible symmetries can move any point of a fiber to any other, then no point of that fiber is distinguished, and any selection must distinguish one. The resulting obstruction is not the observation that p is many-to-one, and it is important to see why. A many-to-one map may perfectly well admit sections that are continuous, measurable, smooth, or equivariant; being many-to-one is compatible with all of these. What Theorems 6.6 and 6.7 identify is a specific incompatibility between the act

of selecting a representative and the symmetry that the readout preserves. It is the symmetry, not the mere coarseness, that does the work.

Two further observations delimit the result. First, the obstruction is sensitive to the group: enlarging or shrinking G changes the stabilizers G_t and can create or destroy fixed points. A section may exist equivariantly for one symmetry and not for another. Second, the obstruction is to *deterministic* selection. Section 7 exhibits a case in which no equivariant deterministic section exists while an equivariant probability-valued assignment on fibers does.

7. A FINITE MARKED-CARRIER OCCUPATION MODEL

We now realize the equivariant obstruction in a completely explicit finite setting.

7.1. The model. Fix integers $N \geq 2$ and $K \geq 2$. Consider N individually marked carriers distributed among K cells. The source state space is the set of assignments

$$S = \{1, \dots, K\}^{\{1, \dots, N\}}, \quad (25)$$

so that a source state $s \in S$ is a function assigning to each marked carrier $i \in \{1, \dots, N\}$ the cell $s(i) \in \{1, \dots, K\}$ that it occupies. Then $|S| = K^N$.

The physical status of the marks is a modeling declaration, and here it is made explicitly: *the carrier identities are physically meaningful source distinctions*. Two assignments that differ by which marked carrier sits in which cell are different source states, not two descriptions of one source state. Accordingly, the physical equivalence $\simeq_{\text{phys}}$ on the underlying realization space is trivial as far as the marks are concerned, and (25) already *is* the reduced source space S in the sense of Theorem 2.1. Theorem 7.4 below examines what happens if this declaration is reversed.

The readout records only how many carriers occupy each cell:

$$p(s) = (n_1(s), \dots, n_K(s)), \quad n_j(s) := \#\{i : s(i) = j\}, \quad (26)$$

with values in

$$T = \left\{ (n_1, \dots, n_K) \in \mathbb{N}_0^K : \sum_{j=1}^K n_j = N \right\}. \quad (27)$$

Every such vector is realized by some assignment, so (3) holds with the T of (27).

7.2. Fibers.

Proposition 7.1 (Fiber cardinality). *For $n = (n_1, \dots, n_K) \in T$,*

$$|F_n| = \binom{N}{n_1, \dots, n_K} = \frac{N!}{n_1! \cdots n_K!}. \quad (28)$$

Proof. An element of F_n is precisely an ordered partition of the set $\{1, \dots, N\}$ of marks into blocks $B_j = s^{-1}(j)$ with $|B_j| = n_j$ for $j = 1, \dots, K$; conversely any such ordered partition determines a unique s with $p(s) = n$. Counting these partitions: choose B_1 in $\binom{N}{n_1}$ ways, then B_2 from the remaining $N - n_1$ marks in $\binom{N - n_1}{n_2}$ ways, and so on. The product telescopes to the multinomial coefficient (28). $\square$

7.3. The symmetry and the obstruction. Let $G = S_N$, the symmetric group on the carrier marks, acting on S by

$$(g \cdot s)(i) := s(g^{-1}i) \quad (g \in S_N, s \in S, i \in \{1, \dots, N\}). \quad (29)$$

This is a left action: $(g \cdot (h \cdot s))(i) = (h \cdot s)(g^{-1}i) = s(h^{-1}g^{-1}i) = ((gh) \cdot s)(i)$. Physically it permutes which marked carrier occupies which cell, while the marks themselves retain their identities.

Proposition 7.2 (Invariance and transitivity). *For every $g \in S_N$ and $s \in S$,*

$$p(g \cdot s) = p(s). \quad (30)$$

Consequently S_N acts trivially on T , the map p is equivariant for this pair of actions, and $G_n = S_N$ for every $n \in T$. Moreover S_N acts transitively on each fiber F_n .

Proof. Invariance: for each cell j ,

$$n_j(g \cdot s) = \#\{i : s(g^{-1}i) = j\} = \#\{i' : s(i') = j\} = n_j(s),$$

where the middle equality is the substitution $i' = g^{-1}i$, a bijection of $\{1, \dots, N\}$. This proves (30). If S_N is given the trivial action on T , then $g \cdot p(s) = p(s) = p(g \cdot s)$, which is exactly the equivariance condition (17); and every g fixes every n , so $G_n = S_N$.

Transitivity: let $s, s' \in F_n$. Then for each j the sets $s^{-1}(j)$ and $s'^{-1}(j)$ both have n_j elements, so there is a bijection $\beta_j : s'^{-1}(j) \rightarrow s^{-1}(j)$. The blocks $s'^{-1}(1), \dots, s'^{-1}(K)$ partition $\{1, \dots, N\}$, as do the blocks $s^{-1}(1), \dots, s^{-1}(K)$, so the β_j assemble into a permutation $g \in S_N$ with $g(s'^{-1}(j)) = s^{-1}(j)$ for all j . Then $s(gi') = j$ exactly when $i' \in s'^{-1}(j)$, i.e. $s \circ g = s'$, i.e. $s'(i) = s(gi)$. Writing $h = g^{-1}$ gives $(h \cdot s)(i) = s(h^{-1}i) = s(gi) = s'(i)$, so $h \cdot s = s'$. $\square$

The situation is thus the one isolated by Theorem 6.7 in its most extreme form: the symmetry group is entirely invisible to the readout, so every stabilizer is the whole group, and that whole group acts transitively on every fiber.

Example 7.3 (A fiber with no invariant representative). Take

$$n_0 = (1, N-1, 0, \dots, 0) \in T, \quad (31)$$

which is a legitimate element of (27) since $1 + (N-1) = N$ and $K \geq 2$. By Theorem 7.1,

$$|F_{n_0}| = \frac{N!}{1!(N-1)!0!\dots 0!} = N > 1, \quad (32)$$

using $N \geq 2$: the fiber consists of the N assignments in which exactly one marked carrier occupies cell 1 and the remaining $N-1$ occupy cell 2, indexed by which carrier is the solitary one.

By Theorem 7.2, $G_{n_0} = S_N$ acts transitively on F_{n_0} , and $|F_{n_0}| > 1$. Theorem 6.7 therefore gives

$$\text{Fix}_{S_N}(F_{n_0}) = \emptyset, \quad (33)$$

and no S_N -equivariant deterministic section $T \rightarrow S$ exists. That is: *no permutation-equivariant rule assigns to occupation data a specific marked-carrier assignment*. Any rule that names the carrier sitting alone in cell 1 must name a particular mark, and conjugating the rule by the transposition exchanging that mark with another produces a different answer from identical readout data.

Directly: if σ were equivariant, then $\sigma(n_0) \in F_{n_0}$ would be some assignment in which carrier i_* is alone in cell 1; taking g to be a transposition moving i_* , we get $g \cdot \sigma(n_0) \neq \sigma(n_0)$, while $\sigma(g \cdot n_0) = \sigma(n_0)$ because S_N acts trivially on T . This contradicts (21).

Note that ordinary sections abound here: T is finite, and one may simply declare that the solitary carrier is always mark 1. What fails is not selection but *symmetric* selection. The convention “mark 1” is a choice with no warrant in the data, and the symmetry of the problem exposes it as such.

Remark 7.4 (The modeling boundary). The obstruction in Theorem 7.3 depends entirely on the declaration made after (25), and it is worth stating exactly where the line falls.

Suppose instead that the carrier marks were declared *physically redundant* — that permuting which carrier is which is a relabeling and not a physical operation. Then the permutations (29) belong to $\simeq_{\text{phys}}$, and the reduction of Theorem 2.1 must be performed with respect to them. The reduced source space becomes

$$S' = \{1, \dots, K\}^{\{1, \dots, N\}} / S_N,$$

whose elements are exactly the occupation vectors: two assignments have the same S_N -orbit if and only if they induce the same $(n_1, \dots, n_K)$, by the transitivity established in Theorem 7.2. The induced readout $p': S' \rightarrow T$ is then a bijection, every reduced fiber is a singleton, and by Theorem 3.2 the readout is source-equivalent. The obstruction has vanished — and correctly so, because on this declaration there was never anything to lose.

The two models are mathematically adjacent and physically opposite. The theorems of Sections 4 and 6 detect lost source structure only after the theory has independently specified which distinctions are physical. They do not, and cannot, make that specification on the theory's behalf. In particular, in Theorem 7.3 the group S_N is *not* a gauge group: it acts on physically distinct states and its orbits are not points of the reduced source space. Treating it as gauge would be a different model, the one just described, with a different and trivial conclusion.

8. INVARIANT ENSEMBLES VERSUS DETERMINISTIC SELECTION

The obstruction of Theorem 7.3 is an obstruction to equivariant *deterministic* selection. Weakening the demand from a point of the fiber to a probability distribution over the fiber removes it, and it is important to see both that this is so and that it does not amount to reconstruction. Making this precise requires typing the weakened object correctly: it is not a section of p , because its values do not lie in S .

8.1. Fiber-supported probability kernels. The only application in this paper is the finite model of Section 7, and the construction is stated accordingly: everything below is discrete, and no measure-theoretic apparatus is required or invoked.

For a set X , let $\mathcal{P}_{\text{fin}}(X)$ denote the set of *finitely supported probability distributions* on X : functions $\nu: X \rightarrow [0, 1]$ that vanish outside a finite subset and satisfy $\sum_{x \in X} \nu(x) = 1$, the sum having only finitely many nonzero terms. For $\nu \in \mathcal{P}_{\text{fin}}(X)$ define the *algebraic support*

$$\text{supp}(\nu) := \{x \in X : \nu(x) > 0\}, \quad (34)$$

a finite set. For $x \in X$ write $\delta_x \in \mathcal{P}_{\text{fin}}(X)$ for the point distribution, $\delta_x(x) = 1$ and $\delta_x(y) = 0$ for $y \neq x$. For a map $f: X \rightarrow Y$ the *discrete pushforward* $f_*\nu \in \mathcal{P}_{\text{fin}}(Y)$ is

$$(f_*\nu)(y) := \sum_{x \in f^{-1}(y)} \nu(x), \quad (35)$$

a finite sum, since only the finitely many points of $\text{supp}(\nu)$ contribute. Note $f_*\delta_x = \delta_{f(x)}$.

Definition 8.1 (Probability-valued lift). A *probability-valued lift* of $p: S \rightarrow T$, equivalently a *fiber-supported probability kernel*, is a map

$$\mu: T \longrightarrow \mathcal{P}_{\text{fin}}(S), \quad t \longmapsto \mu_t, \quad (36)$$

such that each μ_t is supported on the fiber over t :

$$\text{supp}(\mu_t) \subseteq F_t, \quad \text{equivalently} \quad p_*\mu_t = \delta_t \quad (t \in T). \quad (37)$$

If a group G acts on S and T with p equivariant, the lift μ is *equivariant* when

$$g_*\mu_t = \mu_{g \cdot t} \quad \text{for all } g \in G, t \in T. \quad (38)$$

The two formulations of (37) agree, by (35) applied to p . Indeed $(p_*\mu_t)(t') = \sum_{s \in F_{t'}} \mu_t(s)$, so $p_*\mu_t = \delta_t$ holds precisely when this sum is 1 for $t' = t$ and 0 for every $t' \neq t$; since the fibers partition S and $\mu_t \geq 0$, that is precisely the condition that μ_t vanish off F_t , i.e. $\text{supp}(\mu_t) \subseteq F_t$.

The relationship to sections is exactly this: a section σ of p induces the probability-valued lift $t \mapsto \delta_{\sigma(t)}$, which satisfies (37) because $\sigma(t) \in F_t$, and which is equivariant in the sense of (38) precisely when σ is equivariant in the sense of (21), since $g_*\delta_x = \delta_{g \cdot x}$. The deterministic selectors thus embed among the probability-valued lifts as exactly the point-distribution-valued ones. A probability-valued lift is emphatically not a map $T \rightarrow S$, and calling it a section would misdescribe its type; the question below is whether the larger class contains an equivariant member when the point-distribution-valued subclass does not.

8.2. The finite model.

Proposition 8.2 (An equivariant probability-valued lift exists). *In the model of Section 7, for each $n \in T$ let $\mu_n \in \mathcal{P}_{\text{fin}}(S)$ be the uniform distribution on the finite fiber F_n :*

$$\mu_n(s) = \frac{1}{|F_n|} \quad (s \in F_n), \quad \mu_n(s) = 0 \quad (s \in S \setminus F_n). \quad (39)$$

This lies in $\mathcal{P}_{\text{fin}}(S)$ because F_n is finite, by Theorem 7.1. Then $n \mapsto \mu_n$ is a fiber-supported probability kernel in the sense of Theorem 8.1, each μ_n is invariant under the stabilizer $G_n = S_N$, and the kernel is S_N -equivariant. Hence an equivariant probability-valued lift of p exists even though, by (33), no equivariant deterministic section exists.

Proof. Fiber support. By construction μ_n vanishes off F_n , so $\text{supp}(\mu_n) \subseteq F_n$, and by the equivalence noted after Theorem 8.1 this is (37), i.e. $p_*\mu_n = \delta_n$.

Stabilizer invariance. By (18) and the triviality of the S_N -action on T , each $g \in S_N$ maps F_n bijectively onto $F_{g \cdot n} = F_n$. Writing g also for the induced bijection $s \mapsto g \cdot s$ of S , the discrete pushforward (35) along a bijection is $(g_*\mu_n)(s) = \mu_n(g^{-1} \cdot s)$, since $g^{-1}(s)$ is the single point $g^{-1} \cdot s$. Hence for $s \in F_n$ we have $g^{-1} \cdot s \in F_n$ and

$$(g_*\mu_n)(s) = \mu_n(g^{-1} \cdot s) = \frac{1}{|F_n|} = \mu_n(s),$$

while for $s \notin F_n$ both sides vanish, because $g^{-1} \cdot s \notin F_n$. Thus $g_*\mu_n = \mu_n$ for every $g \in S_N$: each μ_n is G_n -invariant.

Equivariance. Condition (38) requires $g_*\mu_n = \mu_{g \cdot n}$. The left side equals μ_n by the invariance just proved, and the right side equals μ_n because S_N acts trivially on T , so $g \cdot n = n$. The two sides agree.

The nonexistence of an equivariant deterministic section is (33), established in Theorem 7.3. $\square$

8.3. An ensemble is not a realized state. The contrast is precise, and it locates the obstruction correctly. Symmetry is capable of determining a distinguished *ensemble* over a fiber: here the uniform distribution is in fact the unique S_N -invariant element of $\mathcal{P}_{\text{fin}}(S)$ supported on F_n , since invariance under a transitive action forces equal weights on all points of the fiber. What symmetry cannot do is single out a *point*. The reason is visible in the proof: the object that survives the group action is precisely the object that has forgotten which element of the fiber it is concentrated on.

Consequently an equivariant probability-valued lift is not a reconstruction of the realized source state, and the existence of the former is no evidence for the availability of the latter. In any given realization one specific assignment $s \in F_n$ obtains — one specific marked carrier is the one sitting alone in cell 1 — and μ_n assigns that assignment probability $1/|F_n|$, exactly the same weight it assigns to every other element of the fiber. The kernel μ is a correct, symmetry-compatible description of the readout-conditional uncertainty about s ; it is not, and does not become, a determination of s . Nothing in μ_n distinguishes the realized state from the $|F_n| - 1$ states that were not realized.

The general moral is that weakening the target of reconstruction weakens the conclusion by exactly as much. A statistical reconstruction and an exact realization reconstruction are different mathematical tasks with different success conditions, and success at the first is silent about the second.

9. A CONTINUOUS OBSTRUCTION: THE HOPF FIBRATION

The occupation model is finite, its symmetry group is finite, and no topology is involved. We now exhibit the obstruction in a continuous setting, where two logically independent mechanisms are simultaneously available. It is essential to keep them apart: one is an instance of Theorem 6.7 and uses no topology, the other is topological and uses no equivariance.

9.1. The Hopf bundle. Regard $S^3 \subset \mathbb{C}^2$ as the unit sphere $\{(z_1, z_2) : |z_1|^2 + |z_2|^2 = 1\}$ and $S^2 \cong \mathbb{CP}^1$ as the space of complex lines in $\mathbb{C}^2$. The *Hopf map*

$$h: S^3 \longrightarrow S^2, \quad h(z_1, z_2) = [z_1 : z_2], \quad (40)$$

sends a unit vector to the complex line it spans. It is a smooth surjection, and it is the projection of a principal $U(1)$ -bundle for the action

$$\lambda \cdot (z_1, z_2) := (\lambda z_1, \lambda z_2), \quad \lambda \in U(1), \quad (41)$$

which preserves the unit sphere and does not change the spanned line. Two unit vectors span the same line if and only if they differ by such a phase, so the fibers of h are exactly the $U(1)$ -orbits, each a circle. The action (41) is free (if $\lambda z_i = z_i$ for $i = 1, 2$ with $(z_1, z_2) \neq 0$, then $\lambda = 1$) and transitive on each fiber. These facts, and the local triviality of h , are standard [3, 4, 9].

9.2. First obstruction: no equivariant section. Let $G = U(1)$ act on $S = S^3$ by (41) and *trivially* on $T = S^2$. The map h is equivariant for this pair of actions, since $h(\lambda \cdot z) = h(z) = \lambda \cdot h(z)$. Hence, exactly as in Theorem 7.2, every stabilizer is the whole group:

$$G_t = U(1) \quad \text{for every } t \in S^2, \quad (42)$$

and this group acts transitively on the circle fiber $F_t = h^{-1}(t)$, which has more than one point.

Proposition 9.1 (No equivariant selector for the Hopf map). *Fix $_{U(1)}(F_t) = \emptyset$ for every $t \in S^2$, and consequently h admits no $U(1)$ -equivariant section.*

Proof. The action of $G_t = U(1)$ on F_t is free, so the only element that fixes any point of F_t is $\lambda = 1$; since $U(1) \neq \{1\}$, no point of F_t is fixed by all of G_t . (Equivalently, the action is transitive on F_t and $|F_t| > 1$, so Theorem 6.7 applies directly.) Theorem 6.6 now excludes an equivariant section. $\square$

This argument is purely group-theoretic. It uses no topology of S^3 or S^2 beyond the identification of the fibers with $U(1)$ -orbits, and it would apply verbatim to the set-theoretic version of the same action. In particular, it does not depend on the Hopf bundle being nontrivial; the same conclusion holds for the trivial bundle $S^2 \times U(1) \rightarrow S^2$ with $U(1)$ acting by translation on the fiber and trivially on the base, which certainly does admit continuous global sections.

9.3. Second obstruction: no continuous section. The second obstruction is of an entirely different kind: it concerns continuity, not symmetry, and it *does* depend on the topology of the bundle.

Lemma 9.2 (A global section trivializes a principal bundle). *Let $P \rightarrow X$ be a principal $U(1)$ -bundle admitting a continuous global section $\sigma: X \rightarrow P$. Then P is isomorphic, as a principal $U(1)$ -bundle, to the trivial bundle $X \times U(1)$.*

Proof. Because the $U(1)$ -action on P is free and transitive on each fiber, the map

$$\Phi: X \times U(1) \longrightarrow P, \quad \Phi(x, \lambda) := \lambda \cdot \sigma(x),$$

is well defined and continuous, is fiber-preserving over X , and is $U(1)$ -equivariant for the translation action on the second factor. It is bijective on each fiber: surjective by transitivity, injective by freeness. Its inverse is continuous because local trivializations of P turn Φ into the map $(x, \lambda) \mapsto (x, \lambda \tau(x))$ for a continuous $U(1)$ -valued function τ obtained from σ in that trivialization, whose inverse $(x, \nu) \mapsto (x, \nu \tau(x)^{-1})$ is continuous. Hence Φ is an isomorphism of principal bundles. This is standard; see [3, 4, 5]. $\square$

Proposition 9.3 (No continuous global section of the Hopf map). *The Hopf bundle $h: S^3 \rightarrow S^2$ admits no continuous global section.*

Proof. Principal $U(1)$ -bundles over a paracompact base X are classified up to isomorphism by a first Chern class in $H^2(X; \mathbb{Z})$, the trivial bundle corresponding to 0 [4, 7, 6]. Fix a generator

$$u \in H^2(S^2; \mathbb{Z}) \cong \mathbb{Z}.$$

For the Hopf bundle,

$$c_1(h) = \pm u \neq 0, \tag{43}$$

that is, $c_1(h)$ is a generator of $H^2(S^2; \mathbb{Z})$. Equivalently, h is the unit-circle bundle of a tautological line bundle over $\mathbb{CP}^1$, whose Euler class generates $H^2(S^2; \mathbb{Z})$ [7, 8]. Since $c_1(h) \neq 0$, the Hopf bundle is not isomorphic to the trivial bundle. By Theorem 9.2, a continuous global section would force triviality. Hence no continuous global section exists. $\square$

Remark 9.4 (On the sign). The sign in (43) is convention-dependent and immaterial here. It turns on the choice of orientation of S^2 , the choice of generator u , and the convention by which a line bundle is associated to the principal bundle — in particular, whether one associates the tautological bundle $\mathcal{O}(-1)$ or its dual $\mathcal{O}(1)$, which have Chern classes differing by a sign. Fixing a set of conventions pins the value to $+u$ or $-u$, and different standard sources fix it differently. The argument of Theorem 9.3 uses only $c_1(h) \neq 0$, and is therefore insensitive to all of these choices. We accordingly refrain from asserting a signed value.

Remark 9.5 (The two obstructions are independent). Theorems 9.1 and 9.3 are not two proofs of one fact. They exclude different things, by different means, and neither implies the other.

That the equivariant obstruction does not entail the topological one is settled by an example. Let $P = S^2 \times U(1) \rightarrow S^2$ be the *trivial* principal $U(1)$ -bundle, with $U(1)$ acting by translation on the fiber factor and trivially on the base. Then $c_1(P) = 0$ and P admits continuous global sections in profusion — for instance $x \mapsto (x, 1)$. Nevertheless the hypotheses of Theorem 9.1 hold verbatim: every stabilizer is all of $U(1)$, which acts freely and transitively on each fiber, so $\text{Fix}_{U(1)}(F_x) = \emptyset$ and Theorem 6.6 excludes every $U(1)$ -equivariant section. A bundle can therefore have no equivariant section while having continuous sections, and Theorem 9.1 cannot be a proof of Theorem 9.3 in disguise.

The converse direction must be argued differently, and the same trivial-bundle example does not serve; it establishes one implication's failure, not both. What can be said is the following. The two propositions constrain *distinct admissibility classes*, and the class is not determined by the map: it is fixed by the additional data one specifies — a group action in one case, a topology in the other. Neither proof, taken by itself, establishes the other's conclusion for an independently specified group action or an independently specified topology, since neither proof mentions the other's data. Concretely, if one takes the group to be trivial, $G = \{1\}$, then every set-theoretic section of h is automatically G -equivariant, so the equivariant reconstruction class is unobstructed for that group; yet Theorem 9.3 continues to exclude every continuous section of h . The topological obstruction thus persists in a setting where the equivariant obstruction is vacuous, and it plainly does not follow from equivariance considerations, which in that setting say nothing at all.

Set-theoretic sections of h , meanwhile, exist whenever a choice function for the fibers does; the axiom of choice supplies one, and no obstruction of either kind touches them. The Hopf map is thus a single example on which three admissibility conditions — bare choice, $U(1)$ -equivariance, continuity — yield three different verdicts, and the verdicts are logically separate.

The lesson generalizes beyond the example. The conditions one imposes on a reconstruction — equivariance, continuity, smoothness, measurability, locality, covariance — are not decoration; they determine what obstructions are available and what the failure of reconstruction means. A bare set-theoretic choice of fiber representatives is possible whenever choice is available and is almost never physically meaningful on its own. Conversely, a proof that no continuous section exists says nothing about equivariance, and a proof that no equivariant section exists says nothing about continuity. Each requirement must be stated, and the corresponding obstruction proved, on its own.

9.4. The modeling boundary, again. The results just proved are theorems about the bundle h , and as theorems they stand without qualification. Their *physical* reading, however, depends on a declaration that the mathematics does not make, and the situation is exactly the one already met in Theorem 7.4.

Remark 9.6 (Two readings of the Hopf example). Applying Theorem 10.1 with $S = S^3$ and $T = S^2$ treats distinct points of a $U(1)$ fiber as *distinct elements of the reduced source*. Two readings are available, and they are not interchangeable.

Reading one: the phase is physical. If the theory has independently declared that the fiber distinctions are physically meaningful — that two unit vectors differing by a phase are different source states, not two descriptions of one — then S^3 is the reduced source space in the sense of Theorem 2.1, the fibers $F_t \cong U(1)$ are nonsingleton fibers of physically distinct states, and Theorems 9.1 and 9.3 are obstructions to reconstructing physically meaningful lost structure.

This is the reading under which the example instantiates Theorem 10.1, and it is legitimate only if the declaration has actually been made.

Reading two: the phase is gauge. If instead the $U(1)$ action (41) is declared to be gauge redundancy — as it is in the usual reading of the overall phase of a normalized state vector — then the action belongs to $\simeq_{\text{phys}}$, and Theorem 2.1 requires that it be quotiented out *before* any source-readout comparison is made. The reduced source space is then

$$S' = S^3/U(1) \cong S^2,$$

the induced readout $p': S' \rightarrow S^2$ is the identity map under this identification, every reduced fiber is a singleton, and by Theorem 3.2 the readout is source-equivalent. Nothing physical has been lost, and Theorem 10.1 has no instance here.

Under reading two the no-section results retain their mathematical content but change their meaning entirely. Theorem 9.3 then says that there is no global continuous *gauge choice* on S^2 : one cannot continuously pick a phase representative over the whole base, and any concrete phase convention must be introduced patchwise. That is a genuine and well-known fact about the bundle, and it is the reason gauge fields on S^2 require more than one chart. It is not a statement that some physically meaningful source structure has been lost, because on this reading there is no such structure in the fiber to lose.

The point to be drawn is therefore a negative one, and it should be stated plainly: *bundle nontriviality is not, by itself, evidence of a physically richer source*. A nonzero Chern class obstructs global continuous sections whatever the physical status of the fiber, and it obstructs them just as firmly when the fiber is pure gauge as when it is not. To convert the topology into a claim about lost realization structure, one must first supply the declaration that the fiber distinctions are physical — and that declaration is an input to the framework, never an output of it.

10. THE SOURCE-READOUT NON-EQUIVALENCE THEOREM

The threads can now be drawn together.

Theorem 10.1 (Source-readout non-equivalence). *Let Ω be a realization space with declared physical equivalence $\simeq_{\text{phys}}$, let $S = \Omega/\simeq_{\text{phys}}$ be the reduced source space, and let $p: S \rightarrow T$ be the induced surjective readout, as in Section 2. Suppose at least one of the following holds.*

(D) (Descent obstruction.) *There is a physically meaningful source relation $r: S \rightarrow V$ and a pair $s_0, s_1 \in S$ with*

$$p(s_0) = p(s_1), \quad r(s_0) \neq r(s_1).$$

(E) (Equivariant obstruction.) *A group G acts on S and on T , the map p is equivariant, and some $t \in T$ satisfies*

$$\text{Fix}_{G_t}(F_t) = \emptyset.$$

Then p is not injective, so $T \not\preceq_{\text{source}} S$ in the sense of Theorem 3.3. More precisely:

- (1) *under (D), the relation r does not descend to T , and hence T is not source-complete for r ;*
- (2) *under (E), no G -equivariant section $T \rightarrow S$ exists;*
- (3) *in either case, $T \cong S/\sim_p$ continues to hold exactly.*

Proof. Assume (D). Theorem 4.3 applied to s_0, s_1 shows that r does not descend, giving conclusion (1). Moreover $r(s_0) \neq r(s_1)$ forces $s_0 \neq s_1$, while $p(s_0) = p(s_1)$; hence p identifies two distinct points of S and is not injective.

Assume (E). Theorem 6.6 excludes a G -equivariant section, giving conclusion (2). For noninjectivity, observe that F_t cannot be a singleton: if $F_t = \{s\}$, then by (18) every $g \in G_t$ maps F_t into $F_{g \cdot t} = F_t = \{s\}$, forcing $g \cdot s = s$ for every $g \in G_t$ and hence $s \in \text{Fix}_{G_t}(F_t)$, contradicting the hypothesis that this set is empty. So $|F_t| \geq 2$, and p is not injective by Theorem 3.2.

Conclusion (3) is Theorem 3.1, whose proof uses nothing beyond surjectivity of p and the definition (7) of $\sim_p$, and is therefore unaffected by either hypothesis. $\square$

Theorem 10.1 claims exactly what its hypotheses support, and it is worth being explicit about the division of labor among them, because each part is doing work that none of the others can do.

The declaration of $\simeq_{\text{phys}}$ determines which source distinctions are genuine. Without it there is no fact of the matter about whether a fiber contains physically distinct states or merely redundant descriptions of one state, and Theorem 7.4 shows that the two possibilities give opposite verdicts on identical mathematics.

Hypothesis (D) determines which relations fail to descend. Coarseness of the readout alone does not identify anything as lost; a nominated relation r , varying within a fiber, does.

Hypothesis (E) determines whether a symmetry-compatible representative can be selected at all. Coarseness alone does not exclude a section, and neither does noninjectivity; symmetry — or, in the topological setting of Theorem 9.3, characteristic-class data — does. Note what (E) does and does not deliver in the negative and positive cases respectively: its failure excludes every selector in the equivariant class, while its success certifies only that some equivariant selector exists, not that one is uniquely determined or canonical (Theorem 6.5).

Conclusion (3) records what survives, and only that. The readout remains an exact, error-free presentation of the quotient $S/\sim_p$: a bijection of sets, permitting transport of any structure already well defined on $\sim_p$ -classes. Whether the readout level is additionally lawful, self-consistent, or dynamically closed is not settled by conclusion (3) and would require separate descent or closure arguments. What is settled is that exactness of the presentation is compatible with, and says nothing against, the non-equivalence established by (1) and (2).

Finally, a boundary that the theorem does not cross. Theorem 10.1 is a theorem about a given Ω , $\simeq_{\text{phys}}$, p , and (under (D) or (E)) a given r or G . It establishes non-equivalence *relative to* that identified lost relation or that reconstruction class. It does not establish, and cannot be used to establish, that any actual physical system possesses a source layer richer than its observed description; that would require producing the ingredients, not merely the theorem. Noninjectivity of a map is not by itself evidence of a physically real upstream source, and no abstract obstruction licenses an unrestricted ontological conclusion.

11. INTERPRETATION AND SCOPE

The research program under which the present results are developed — Shadow Theory — posits a reduced source-realization domain S and a readout domain T_1 , related by a physically invariant readout

$$p: S \longrightarrow T_1, \tag{44}$$

in exactly the sense of Section 2. The mathematics of the preceding sections is stated in standard language and requires no commitment to the program; what the program supplies is a candidate instantiation of the hypotheses.

Within that setting, Theorem 10.1 yields the following statements, each conditional on the hypothesis it invokes. If a realization-relevant relation varies within a T_1 -fiber, that relation is not a function of the T_1 -state, and no construction taking only T_1 -data as input can evaluate it. If, in addition, a fiber carries a readout-invisible symmetry with no invariant representative,

then no symmetry-compatible deterministic reconstruction selects the realized source state from T_1 -data. Both conclusions coexist with the exactness of the quotient presentation, so that the structure is

$$T_1 \cong S/\sim_p \quad \text{while} \quad T_1 \not\cong_{\text{source}} S, \quad (45)$$

the second relation holding *relative to* the lost relation or the reconstruction class established by the theorem. The qualification is not a hedge; it is the content. Non-equivalence is always non-equivalence with respect to something identified, and (45) means nothing more and nothing less than that.

The picture that results is coherent but modest, and its modesty should not be trimmed away. Theorem 3.1 establishes that T_1 is an exact presentation of $S/\sim_p$: as a description of $\sim_p$ -classes it is neither approximate nor defective. It does not establish that the T_1 -level description is lawful, or internally consistent as a dynamical theory, or dynamically closed. A T_1 -level description *may* be all of these while nevertheless being a readout of a richer realization carrier — there is no tension whatever between the two, and demonstrating exactness does not argue against a richer source — but establishing lawfulness or closure at the readout level requires descent results for the specific laws in question, and those are not proved here. This paper neither asserts nor denies that projected dynamics descends; it fixes the framework within which that question can be posed.

Establishing that a particular physical theory has this structure is a separate task from proving the theorems above, and it requires that the theory independently specify all of the following:

- (1) the reduced source domain S , and the realization space Ω from which it is obtained;
- (2) the physical-equivalence relation $\simeq_{\text{phys}}$, that is, an explicit declaration of which transformations are redundant;
- (3) the readout map, and hence p and the fibers F_t ;
- (4) which distinctions within a fiber are physically meaningful rather than gauge — the declaration whose reversal, in Theorem 7.4, dissolved the obstruction entirely;
- (5) either a non-descending physically meaningful relation r , or a reconstruction class (equivariant, continuous, smooth, measurable, local, covariant) together with an applicable obstruction in that class.

Absent these, Theorem 10.1 is a theorem with no instance. Supplied with them, it is a theorem with a specific and bounded conclusion. In neither case does it license the inference that reality is a projection, or that some particular observed physics has been shown to be one; those are claims about instantiation, and instantiation is not proved by abstraction.

12. DISCUSSION

12.1. The levels of obstruction. It may be useful to set the distinct notions side by side, since the burden of the paper is to keep them apart.

Exact quotient presentation (Theorem 3.1) always holds. It asserts a bijection between T and $S/\sim_p$, permits transport of structures already well defined on $\sim_p$ -classes, and asserts nothing about the size of the classes, nothing about which nominated structures descend, and nothing about lawfulness or dynamical closure at the readout level. It is compatible with total loss of source structure.

Source equivalence (Theorem 3.2, Theorem 3.3) holds exactly when the fibers are singletons. It is the strongest of the notions and is the one at issue.

Descent of a relation (Theorem 4.2) is a property of a nominated relation r , not of p alone. It holds exactly when r is fiber-constant. Its failure identifies a specific lost quantity. It is a statement about *functions on the source*, not about source states.

Existence of a section (Theorem 5.1) is nearly vacuous without admissibility conditions, and is never equivalent to inversion unless p is bijective (Theorem 5.2). It is a statement about *representative conventions*, not about recovery.

Existence of an equivariant section (Theorem 6.3) holds exactly when every stabilizer fixes a point of its fiber, the constructive direction requiring choice for both the orbit family and the family of fixed sets. Its failure is a genuine symmetry obstruction and is choice-free to prove. Its success is existence only: not uniqueness, and not canonicity (Theorem 6.5). The obstruction is moreover relative to the imposition of equivariance as an admissibility requirement. A non-equivariant selector is not thereby unlawful, and it is formally a map $T \rightarrow S$ like any other; what it is not is determined by the readout together with the declared G -symmetric structure alone, since its choice encodes an additional convention or symmetry-breaking ingredient.

Existence of a continuous section is a topological question, governed in the principal-bundle case by characteristic classes (Theorem 9.3), and is logically independent of the equivariant question (Theorem 9.5).

Between the second and the last of these lies a range of intermediate requirements — smoothness, measurability, locality, covariance — each of which defines its own notion of admissible reconstruction and its own obstruction theory. Nothing in this paper collapses them, and a claim of “no reconstruction” is meaningless until the class is named.

12.2. What the theorems do not say. Several inferences are tempting and unsupported, and it is worth stating them plainly as non-consequences.

Noninjectivity does not by itself exclude a section: many-to-one maps routinely admit continuous, smooth, measurable, and even equivariant sections. Only the specific hypotheses of Theorem 6.6 or Theorem 9.3 exclude them.

A section does not constitute an inverse. This is Theorem 5.2, and it is the point at which “we can always pick a representative” fails to become “we can always recover the state”.

An invariant probability distribution over a fiber does not constitute recovery of the realized source state. This is the burden of Section 8. Ensembles and realizations are different objects with different information content.

Gauge redundancy is not lost structure. If a fiber consists of gauge-equivalent descriptions, then it should have been collapsed by $\simeq_{\text{phys}}$ before the analysis began, and the obstruction it appears to exhibit is an artifact. Theorem 7.4 exhibits both readings of a single piece of mathematics.

Bundle nontriviality is not evidence of a physically richer source. A nonzero characteristic class obstructs continuous sections whether the fiber is physically meaningful or pure gauge, and Theorem 9.6 exhibits both readings of the same bundle. Topology cannot be substituted for the physical declaration.

Quotient exactness is not dynamical closure. Theorem 3.1 is a bijection of sets. That the readout level supports a closed dynamics, if true, is a separate theorem requiring its own descent hypotheses.

An abstract obstruction is not a physical ontology. The theorems establish what fails to descend and what cannot be selected, given hypotheses. They do not establish that the hypotheses hold of the world.

12.3. Position within the wider program. The results here supply a foundation on which further questions can be posed precisely, and it is worth naming them, though nothing in this paper depends on their answers and this paper is complete without them.

Theorem 4.2 makes it meaningful to ask which non-descending relations obstruct a specified mathematical or physical target, and what the resulting failure looks like. Given a family of such relations, one may ask for the coarsest extension of the readout on which they all do descend — a completion problem, whose input is exactly the fiber-constancy criterion established here. One may then ask whether such completions admit geometric or field-theoretic realizations, in what sense the resulting observable quotient supports a projected dynamics, and under what conditions the source is identifiable from readout data with additional structure. Each of these questions presupposes the distinctions drawn above: without them, “completion” and “identifiability” have no determinate meaning.

13. CONCLUSION

An exact quotient is not a source-equivalent representation. The readout $p: S \rightarrow T$ always presents T faithfully as $S/\sim_p$, and this fact carries no implication whatever about whether S can be recovered from T .

The difference between the two is captured by two independent tests, each with a clean criterion and each with a clean failure mode. A source relation is present at the readout level exactly when it is constant on readout fibers; if it varies within a fiber, no function of readout data reproduces it. A symmetry-compatible deterministic representative exists exactly when every readout stabilizer fixes a point of its fiber; if some stabilizer fixes no point, no equivariant selection rule exists, though invariant ensembles may still exist and set-theoretic sections exist under the relevant choice principle. The finite marked-carrier occupation model realizes the second obstruction discretely, and the Hopf fibration realizes it continuously while simultaneously exhibiting an independent topological obstruction of a different character.

The failure of either test yields a precise source–readout obstruction, relative to the relation or the reconstruction class that the test names. Together they give the first mathematical layer of the framework: the distinction between a realization carrier and the exact but potentially non-equivalent image through which that carrier is read.

REFERENCES

- [1] S. Lang, *Algebra*, 3rd revised ed., Graduate Texts in Mathematics 211, Springer, New York, 2002.
- [2] G. E. Bredon, *Introduction to Compact Transformation Groups*, Pure and Applied Mathematics 46, Academic Press, New York, 1972.
- [3] N. Steenrod, *The Topology of Fibre Bundles*, Princeton Mathematical Series 14, Princeton University Press, Princeton, 1951.
- [4] D. Husemöller, *Fibre Bundles*, 3rd ed., Graduate Texts in Mathematics 20, Springer, New York, 1994.
- [5] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry, Volume I*, Interscience Publishers, New York, 1963.
- [6] A. Hatcher, *Algebraic Topology*, Cambridge University Press, Cambridge, 2002.
- [7] J. W. Milnor and J. D. Stasheff, *Characteristic Classes*, Annals of Mathematics Studies 76, Princeton University Press, Princeton, 1974.
- [8] R. Bott and L. W. Tu, *Differential Forms in Algebraic Topology*, Graduate Texts in Mathematics 82, Springer, New York, 1982.
- [9] G. L. Naber, *Topology, Geometry, and Gauge Fields: Foundations*, 2nd ed., Texts in Applied Mathematics 25, Springer, New York, 2011.
- [10] J. Rodgers, *A Source-to-Readout Architecture for a Theory of Everything: Sector Physics as Coupled Projections from a Common Realization Carrier*, Zenodo, 2025. doi:10.5281/zenodo.21366204.

INDEPENDENT RESEARCHER