

TARGET-RELATIVE NECESSITY OF COMPLETION: WHEN READOUT LOSS OBSTRUCTS, AND WHAT A SUFFICIENT EXTENSION MUST RETAIN

JEREMY RODGERS

ABSTRACT. A surjective, noninjective readout map $p: S \rightarrow T$ from a reduced source space necessarily discards distinctions among reduced source states. It does not follow that the discarded distinctions matter. Whether they matter is a property of the question being asked, not of the readout alone. This paper makes that dependence precise. A *target* is specified by a correct-answer map $a_Q: S \rightarrow A_Q$, and we prove that Q admits an exact readout-only solution if and only if a_Q is constant on every fiber of p ; consequently a lost source relation obstructs Q precisely when it is *active* for Q , i.e. when the realized answer varies across some readout fiber. When exactness fails, the compatible-answer set $A_Q(t) = a_Q(p^{-1}(t))$ is shown to be the unique pointwise-least uniformly sound set-valued rule available without further information; it is a set-theoretic, not probabilistic, object. We then characterize what any repair must contain. Every surjective readout extension $e: S \rightarrow E$ through which a_Q factors must separate every pair of source states with different correct answers, and the joint target image $E_Q := \text{Im}(p, a_Q)$ is itself such an extension and is universally coarsest: every Q -sufficient extension e admits a unique surjection $\varphi_E: E \rightarrow E_Q$ with $j_Q = \varphi_E \circ e$. The framework is realized concretely for flat $\text{U}(1)$ connections on a circle. Curvature is a deliberately restricted readout that collapses the entire flat moduli space $S_{\text{flat}} \cong \text{U}(1)$ to a point, yet the spectrum of the charged covariant Laplacian $H_\theta = -D_x^2$ separates $\theta = 0$ from $\theta = \pi$; we prove that $\text{Spec}(H_\theta) = \text{Spec}(H_{\theta'})$, with multiplicities, if and only if $\theta' \equiv \pm\theta \pmod{2\pi}$, so that the coarsest completion sufficient for the spectral target is the inversion quotient coordinatized by $\cos\theta$, strictly coarser than full holonomy. Finally the same criterion is applied to finite-time projected dynamics: an autonomous readout propagator U_τ with $p \circ \Phi_\tau = U_\tau \circ p$ exists if and only if the future readout is constant on present readout fibers.

1. INTRODUCTION

1.1. The question left open by exactness. Physical descriptions are routinely obtained by passing from a rich source description to a coarser one. Gauge fixing, symmetry reduction, coarse graining, restriction to a chosen family of observables, and projection onto a subsystem all have the same formal shape: a surjection

$$p: S \longrightarrow T \tag{1}$$

from a space S of admissible states onto a space T of readout values. The companion paper to the present one, on source–readout non-equivalence [8], established two facts about such maps. First, an exact quotient readout need not be equivalent to the reduced source it was computed from: even when no approximation is committed and no information is thrown away by error, the map p may fail to be injective, and then T simply does not carry all the distinctions that S

Date: June 2026.

2020 Mathematics Subject Classification. Primary 58J50, 81Q70; Secondary 58J53, 81Q10, 53C29.

Key words and phrases. Quotient readout, information loss, target solvability, flat connections, holonomy, covariant Laplacian, spectral rigidity, projected dynamics.

Website: everythingequation.com.

carries. Second, a source-defined relation $r: S \rightarrow V$ descends to a relation on T precisely when r is constant on the fibers $F_t := p^{-1}(t)$.

That is a clean criterion, but by itself it is a statement about the relation r , and it is tempting to over-read it. The over-reading runs: since p is noninjective, information is lost; since information is lost, the readout description is defective; since the readout description is defective, it must be repaired by additional variables. Each step of this chain is either false or unsupported. Noninjectivity is a structural fact about p , and structural loss is cheap: almost every useful description loses something. The relevant question is never whether p loses distinctions, but whether it loses *the distinctions that the question at hand depends on*. A thermodynamic readout that forgets molecular positions is not thereby unfit to answer thermodynamic questions.

The present paper is devoted to making that dependence exact. We ask:

Given a nominated mathematical or physical question, when does the information discarded by the readout genuinely obstruct answering it, and, when it does, what must *any* successful repair retain?

Both halves of this question are answered here by positive mathematics. In particular the impossibility half is not obtained by assuming that no readout-only construction exists. It is *derived* from the observation that a deterministic construction taking only the readout as input is necessarily of the form $b \circ p$, hence constant on fibers, whereas the correct answer may vary along a fiber. Variation of the correct-answer map on a single fiber is the entire source of the obstruction, and the obstruction is therefore exactly as strong, and no stronger, than that variation.

1.2. Targets and correct-answer maps. The technical device that makes the theory target-relative is elementary and, we will argue, unavoidable. A *target* Q is not a piece of prose. It is a pair, consisting of a set A_Q of admissible answers together with a map

$$a_Q: S \longrightarrow A_Q \tag{2}$$

assigning to each reduced source state the answer that is in fact correct for that state. Nothing mysterious is intended by “correct”: if the target is “does the charged Laplacian have a zero mode?” then $A_Q = \{\text{yes}, \text{no}\}$ and a_Q records the true answer for each connection; if the target is “what is the spectrum, with multiplicities?” then A_Q is a set of multisets and a_Q records the actual spectrum. The values of a_Q may be Boolean, numerical, functional, geometric, spectral, or dynamical; the theory is indifferent.

What the device buys is precision about the object of the obstruction. “Information is lost” is a statement about p . “ Q cannot be solved from the readout” is a statement about the pair (p, a_Q) , and [Theorem 3.1](#) shows the second is equivalent to the failure of a_Q to be fiber-constant. Two different targets, evaluated against the same lossy readout, may receive opposite verdicts, and in [Sections 9 to 11](#) we exhibit exactly that: a curvature readout that is perfectly adequate for one class of questions about flat $U(1)$ connections and provably inadequate for another, together with a completion that is minimal for the second class and strictly insufficient for a third.

1.3. Contributions. The paper is organized around a single logical progression, and we summarize it here so that the dependency order is visible from the outset.

- (1) [Section 2](#) fixes the reduced source space, the readout, and the target, and insists that all statements concern reduced source states — states already quotiented by every declared gauge and coordinate redundancy — so that no apparent “loss” is an artifact of retaining redundant labels.

- (2) [Section 3](#) proves the exact solvability criterion ([Theorem 3.1](#)): Q is exactly solvable from the readout iff a_Q is constant on every fiber of p ; and derives the obstruction ([Theorem 3.2](#)) in the form actually needed, namely that a single fiber on which a_Q takes two values already defeats every deterministic readout-only construction. We show additionally ([Theorem 3.4](#)) that auxiliary quantities computed from the readout cannot repair this, because they too factor through p .
- (3) [Section 4](#) separates *structural* loss from *target-active* loss. A source relation r is active for Q on a fiber when the realized answers it induces there differ; it is inactive when they agree, *even if r itself varies*. Activity on one fiber obstructs ([Theorem 4.3](#)); global inactivity preserves exact solvability ([Theorem 4.4](#)).
- (4) [Section 5](#) identifies what is still available when exactness fails: the compatible-answer set $A_Q(t) = a_Q(F_t)$, proved to be uniformly sound and pointwise contained in every uniformly sound set-valued rule ([Theorem 5.4](#)). We stress that $\mathfrak{P}(A_Q)$ is the power set and that this is a set-valued, not probabilistic, statement.
- (5) [Sections 6](#) and [7](#) define readout extensions and target sufficiency. We then prove the necessary-separation theorem ([Theorem 7.1](#)): every Q -sufficient extension must separate each pair of source states whose correct answers differ.
- (6) [Section 8](#) constructs the joint target image $E_Q = \text{Im}(p, a_Q)$, proves it is a Q -sufficient extension, and proves the universal factorization ([Theorem 8.3](#)) in the correct direction: every Q -sufficient extension E maps onto E_Q , not conversely. E_Q is therefore the coarsest sufficient completion for the single target Q .
- (7) [Sections 9](#) to [11](#) realize all of this for flat $\text{U}(1)$ connections on S^1 with a curvature readout, a charged covariant Laplacian, and its spectrum, culminating in the spectral rigidity result $\text{Spec}(H_\theta) = \text{Spec}(H_{\theta'}) \iff \theta' \equiv \pm\theta$.
- (8) [Section 12](#) applies the same criterion to autonomous projected dynamics at finite time.

[Section 13](#) delimits scope. The sequel [\[9\]](#) generalizes from a single correct-answer map to families of realization relations; continuous-time and memory-kernel questions belong to [\[10\]](#). Neither construction is used here, and the results below are self-contained.

2. REDUCED SOURCE, READOUT, AND TARGET

2.1. The reduced source space. Throughout, S denotes a *reduced* source space. By this we mean the space obtained from whatever raw description one began with *after* quotienting by every redundancy that has been declared physically inert: local gauge transformations, coordinate reparametrizations, relabelings of identical constituents, and any further identification that the modeling commitments of the problem require. The point of insisting on this is not fastidiousness. If S still carried gauge copies, then p would appear to lose information whenever it merely forgot a gauge label, and every theorem below would be vacuous or misleading: one would be proving obstructions that could be dissolved by a change of gauge. By working on the reduced space we guarantee that two distinct points of S are genuinely distinct physical configurations, so that a fiber of p containing two points is a genuine collapse of physical distinctions and not a bookkeeping artifact.

Definition 2.1 (Readout). A *readout* is a surjective map $p: S \rightarrow T$ of sets. For $t \in T$ we write

$$F_t := p^{-1}(t) = \{ s \in S : p(s) = t \}, \quad (3)$$

and call F_t the *readout fiber* over t . Surjectivity of p means every fiber is nonempty, and the fibers partition S .

We write $\sim_p$ for the equivalence relation on S induced by p , so $s \sim_p s'$ iff $p(s) = p(s')$ iff s and s' lie in a common fiber. The readout is *lossy*, or noninjective, when some fiber contains more than one point. No topological, smooth, or measurable structure is required for the results of [Sections 3 to 8](#); those are statements about maps of sets, and their strength derives from that generality. Structure is imposed only in [Sections 9 to 11](#), where the source space is a moduli space of connections and the target values are spectra.

2.2. Targets.

Definition 2.2 (Target and correct-answer map). A *target* Q on S consists of

- (1) a set A_Q of *admissible answers*, and
- (2) a map $a_Q: S \rightarrow A_Q$, the *correct-answer map*, assigning to each reduced source state s the answer $a_Q(s)$ that is correct for s .

That a_Q is a well-defined map on all of S is a substantive standing assumption, and it is worth saying what it does and does not import. It asserts that the question has a determinate answer for each reduced source state — that once the physical configuration is fixed, so is the correct answer. It does *not* assert that the answer is computable, that it is known, or that it is accessible from any particular data. Indeed, the entire content of the paper is the analysis of when the answer, though determinate as a function on S , fails to be a function of $p(s)$. If a purported target does not have a determinate answer per reduced source state — if, say, the answer depends on further facts not encoded in S — then it is not a target on S in the sense of [Theorem 2.2](#), and the source space must be enlarged before the present analysis applies.

Definition 2.3 (Exact readout solvability). The target Q is *exactly solvable from the readout* p if there exists a map

$$\bar{a}_Q: T \longrightarrow A_Q \quad \text{such that} \quad a_Q = \bar{a}_Q \circ p. \quad (4)$$

We then say that a_Q *descends* along p , and call $\bar{a}_Q$ a *readout-only solution* of Q .

Remark 2.4 (Solving a target is weaker than reconstructing the source). [Theorem 2.3](#) asks only that the correct answer be recoverable from $t = p(s)$. It does not ask that s itself be recoverable. This is the asymmetry that makes the whole subject nontrivial. A readout can be violently noninjective — collapsing an infinite-dimensional fiber to a point — and still solve a target exactly, provided the target is blind to everything the fiber varies over. Conversely, an almost-injective readout that fails to separate a single critical pair of states can fail to solve a target that hinges on that pair. Injectivity of p is therefore sufficient for exact solvability of *every* target, but it is wildly stronger than necessary for any *given* target, and the gap between the two is where the useful mathematics lives.

3. EXACT TARGET SOLVABILITY AND ITS OBSTRUCTION

3.1. The criterion.

Theorem 3.1 (Exact target-solvability criterion). *Let $p: S \rightarrow T$ be a readout and let $Q = (A_Q, a_Q)$ be a target on S . The following are equivalent.*

- (1) Q is exactly solvable from the readout; that is, there is $\bar{a}_Q: T \rightarrow A_Q$ with $a_Q = \bar{a}_Q \circ p$.
- (2) a_Q is constant on every fiber of p : for all $t \in T$ and all $s, s' \in F_t$, $a_Q(s) = a_Q(s')$.
- (3) For all $s, s' \in S$, $p(s) = p(s') \implies a_Q(s) = a_Q(s')$.

Moreover, when these hold, $\bar{a}_Q$ is unique.

Proof. Conditions (2) and (3) are restatements of one another: s, s' lie in a common fiber precisely when $p(s) = p(s')$.

(1) $\Rightarrow$ (3). Suppose $a_Q = \bar{a}_Q \circ p$ and $p(s) = p(s')$. Applying $\bar{a}_Q$ to both sides of $p(s) = p(s')$ gives $a_Q(s) = \bar{a}_Q(p(s)) = \bar{a}_Q(p(s')) = a_Q(s')$.

(3) $\Rightarrow$ (1). We construct $\bar{a}_Q$. Let $t \in T$. Since p is surjective, $F_t \neq \emptyset$; choose $s \in F_t$ and set

$$\bar{a}_Q(t) := a_Q(s). \quad (5)$$

This is independent of the choice of s : if $s' \in F_t$ is another element, then $p(s) = t = p(s')$, so by (3) $a_Q(s) = a_Q(s')$. Hence $\bar{a}_Q: T \rightarrow A_Q$ is a well-defined map. (That some element may be chosen from each nonempty fiber is all that is used; no choice principle beyond the definition of a well-defined value is required, since the value assigned is independent of the representative.) Now for any $s \in S$, put $t = p(s)$; then $s \in F_t$, and by (5) together with representative-independence, $\bar{a}_Q(p(s)) = a_Q(s)$. Thus $a_Q = \bar{a}_Q \circ p$.

Uniqueness. Suppose $\bar{a}_Q$ and $\bar{a}'_Q$ both satisfy (4). Let $t \in T$ and choose $s \in F_t$, possible by surjectivity. Then $\bar{a}_Q(t) = \bar{a}_Q(p(s)) = a_Q(s) = \bar{a}'_Q(p(s)) = \bar{a}'_Q(t)$. As t was arbitrary, $\bar{a}_Q = \bar{a}'_Q$. $\square$

Theorem 3.1 is the fiber-descent theorem of [8] specialized to the correct-answer map, and we have given the argument in full so that the present paper stands on its own. What is new here is not the criterion but the object it is applied to: by taking the descending map to be the *answer* to a nominated question rather than an arbitrary source relation, the criterion becomes a statement about the adequacy of the readout *for that question*, and hence acquires the target-relativity that the rest of the paper develops.

3.2. The obstruction.

Corollary 3.2 (Exact target obstruction). *Suppose there exist $s_0, s_1 \in S$ with*

$$p(s_0) = p(s_1) \quad \text{and} \quad a_Q(s_0) \neq a_Q(s_1). \quad (6)$$

Then Q is not exactly solvable from the readout. Explicitly, for every map $b: T \rightarrow A_Q$ there exists $i \in \{0, 1\}$ with $b(p(s_i)) \neq a_Q(s_i)$: no deterministic readout-only construction is correct at both s_0 and s_1 .

Proof. Let $b: T \rightarrow A_Q$ be arbitrary and write $t := p(s_0) = p(s_1)$. The value $b(t) \in A_Q$ is a single element. If $b(t) = a_Q(s_0)$ then, since $a_Q(s_1) \neq a_Q(s_0)$, we get $b(p(s_1)) = b(t) = a_Q(s_0) \neq a_Q(s_1)$, so b errs at s_1 . If instead $b(t) \neq a_Q(s_0)$ then $b(p(s_0)) \neq a_Q(s_0)$ and b errs at s_0 . In either case b fails at one of the two states. Since b was arbitrary, no map $T \rightarrow A_Q$ satisfies $a_Q = b \circ p$, which is the negation of **Theorem 2.3**; this also follows from **Theorem 3.1** (3), whose hypothesis is contradicted by (6). $\square$

The structure of this argument deserves emphasis, because it is the pattern on which all the necessity results of this paper are built, and because it is easy to counterfeit. We did not assume that no readout-only construction exists and then declare completion necessary. We assumed only (6) — a fact about the pair (p, a_Q) , checkable in examples, and verified concretely in **Section 10** — and then *derived* the failure of every candidate b by a counting argument: $b(t)$ is one element of A_Q , while correctness at s_0 and s_1 demands two different elements. Impossibility follows from variation of the correct-answer map along a fiber, not from a stipulation of impossibility.

Proposition 3.3 (Every readout-only construction is fiber-constant). *Let X be any set and let $g: S \rightarrow X$ be a map that factors through the readout, i.e. $g = h \circ p$ for some $h: T \rightarrow X$.*

Then g is constant on every fiber of p . Conversely, every map $S \rightarrow X$ constant on all fibers of p factors through p , uniquely.

Proof. If $g = h \circ p$ and $s, s' \in F_t$, then $g(s) = h(t) = g(s')$. The converse and the uniqueness are [Theorem 3.1](#) applied with A_Q replaced by X and a_Q by g . $\square$

Proposition 3.4 (Auxiliary readout-computed variables cannot repair the obstruction). *Assume (6). Let W be any set and let $w: S \rightarrow W$ be an auxiliary quantity that is computed from the readout alone, meaning $w = \omega \circ p$ for some $\omega: T \rightarrow W$. Then for every map*

$$c: T \times W \longrightarrow A_Q,$$

the induced construction $s \mapsto c(p(s), w(s))$ fails to equal a_Q at s_0 or at s_1 .

Proof. Define $b: T \rightarrow A_Q$ by $b(t) := c(t, \omega(t))$. Then for every $s \in S$,

$$c(p(s), w(s)) = c(p(s), \omega(p(s))) = b(p(s)),$$

so the construction $s \mapsto c(p(s), w(s))$ is exactly $b \circ p$ and hence, by [Theorem 3.3](#), constant on fibers. By [Theorem 3.2](#) it errs at s_0 or at s_1 . $\square$

[Theorem 3.4](#) closes an escape route that is attempted surprisingly often. One cannot evade a fiber-variation obstruction by manufacturing additional quantities out of the readout — derived fields, filtered or smoothed versions of t , spectral transforms of t , statistics of t , or any other post-processing. Whatever is computed from t is a function of t ; composing functions of t leaves one with a function of t ; and a function of t is constant on F_t while a_Q is not. Genuine repair must inject information that is *not* a function of the readout. That is what [Section 6](#) formalizes.

4. ACTIVE AND INACTIVE INFORMATION LOSS

We now sharpen the picture. [Theorem 3.2](#) tells us that fiber variation of a_Q obstructs. But in practice one usually knows something more specific: one knows *which* source-level structure the readout has forgotten, and one wants to know whether *that particular* forgotten structure is what the target depends on. This is the distinction between information that is merely lost and information whose loss has consequences.

Definition 4.1 (Source relation and its target factorization). Let $r: S \rightarrow V$ be a map, called a *source relation*, thought of as recording some physically meaningful attribute of the reduced source state that the readout does not itself determine. We say that Q *factors through* (p, r) if there exists a map

$$\Phi_Q: T \times V \longrightarrow A_Q \quad \text{such that} \quad a_Q(s) = \Phi_Q(p(s), r(s)) \quad \text{for all } s \in S. \quad (7)$$

Equation (7) says that the readout value together with the relation value suffices to determine the correct answer. It is a modeling hypothesis about how much of the source the target actually consults; the interesting content lies in what one may then conclude.

Definition 4.2 (Activity and inactivity on a fiber). Assume (7), and let $t \in T$. The relation r is *active for Q on the fiber F_t* if there exist $s_0, s_1 \in F_t$ with

$$\Phi_Q(t, r(s_0)) \neq \Phi_Q(t, r(s_1)). \quad (8)$$

It is *inactive for Q on F_t* if no such pair exists, i.e. if the map $\Phi_Q(t, \cdot)$ is constant on the *realized set*

$$r(F_t) = \{ r(s) : s \in F_t \} \subseteq V. \quad (9)$$

We say r is *globally inactive for Q* if it is inactive on every fiber.

Two features of [Theorem 4.2](#) carry the weight of the section, and both are easy to get wrong.

First, inactivity is a statement about the *realized target values*, not about r . The relation r may vary wildly across F_t — $r(F_t)$ may be an enormous subset of V — and r may still be inactive for Q on F_t , provided $\Phi_Q(t, \cdot)$ happens to be constant on that subset. The readout has then forgotten a great deal, and none of it matters for Q . This is the normal situation in effective descriptions and is not a pathology.

Second, activity is evaluated on the realized set $r(F_t)$, not on all of V . That $\Phi_Q(t, \cdot)$ is nonconstant somewhere on V is irrelevant if the source states in F_t never realize the values at which it varies. The obstruction concerns configurations that actually occur in the fiber.

Theorem 4.3 (Active loss obstructs). *Assume (7). If r is active for Q on some fiber F_t , then Q is not exactly solvable from the readout p .*

Proof. Let $s_0, s_1 \in F_t$ realize (8). Since $s_0, s_1 \in F_t$ we have $p(s_0) = t = p(s_1)$. By (7),

$$a_Q(s_0) = \Phi_Q(p(s_0), r(s_0)) = \Phi_Q(t, r(s_0)) \neq \Phi_Q(t, r(s_1)) = \Phi_Q(p(s_1), r(s_1)) = a_Q(s_1).$$

Thus s_0, s_1 satisfy the hypothesis (6) of [Theorem 3.2](#), and Q is not exactly solvable from p . Indeed no map $b: T \rightarrow A_Q$ is correct at both. $\square$

Theorem 4.4 (Inactive loss does not obstruct). *Assume (7). If r is globally inactive for Q — that is, if for every $t \in T$ the map $\Phi_Q(t, \cdot)$ is constant on the realized set $r(F_t)$ — then Q is exactly solvable from the readout p . The unique readout-only solution is*

$$\bar{a}_Q(t) = \Phi_Q(t, v_t) \quad \text{for any } v_t \in r(F_t). \quad (10)$$

Proof. We verify condition (2) of [Theorem 3.1](#). Fix $t \in T$ and let $s, s' \in F_t$. Then $r(s), r(s') \in r(F_t)$, and by global inactivity $\Phi_Q(t, \cdot)$ is constant on $r(F_t)$; hence $\Phi_Q(t, r(s)) = \Phi_Q(t, r(s'))$. Using (7) and $p(s) = p(s') = t$,

$$a_Q(s) = \Phi_Q(p(s), r(s)) = \Phi_Q(t, r(s)) = \Phi_Q(t, r(s')) = a_Q(s').$$

So a_Q is constant on F_t , for every t . By [Theorem 3.1](#), Q is exactly solvable and the unique solution satisfies $\bar{a}_Q(t) = a_Q(s)$ for any $s \in F_t$; since $F_t \neq \emptyset$ we may pick such an s and set $v_t := r(s) \in r(F_t)$, giving (10). The right-hand side of (10) is independent of the choice of $v_t \in r(F_t)$, again by constancy of $\Phi_Q(t, \cdot)$ on $r(F_t)$. $\square$

Definition 4.5 (Structural versus consequential loss). Let $r: S \rightarrow V$ be a source relation that does not factor through p ; that is, r is not constant on some fiber. We call this *structural information loss*: the readout genuinely fails to determine r . If in addition r is active for a target Q on some fiber, we call the loss *consequential for Q* , or *target-active*.

[Theorems 4.3](#) and [4.4](#) should be read with their standing hypothesis kept in view. Both presuppose that the nominated variable r *mediates* the target, in the precise sense of (7): $a_Q(s) = \Phi_Q(p(s), r(s))$. Relative to such an r , the two theorems say that activity of r is exactly the condition under which the loss of r obstructs Q — active loss obstructs, and globally inactive loss does not. This is a statement about the chosen mediating variable, not about information loss in the abstract, and it is only as sharp as the choice of r .

The fully general criterion, carrying no such hypothesis, is the one already established in [Theorems 3.1](#) and [3.2](#): Q is obstructed from the readout precisely when a_Q is nonconstant on some fiber of p . Every statement about activity is a refinement of that criterion, obtained by localizing the fiber variation onto a particular mediating variable, and it adds diagnostic content — *which* forgotten structure is responsible — rather than logical strength.

With that scoping in place, the moral stands. Structural loss without activity is invisible to Q ; it is what makes coarse descriptions viable rather than defective. It is therefore *not* a

theorem of this framework — and we do not assert — that every noninjective readout requires completion, nor that every lost source distinction must be restored. Whether a given loss must be repaired is a question one can only answer after nominating a target, and the answer may legitimately be no. The examples of [Sections 9 to 11](#) include both verdicts for the same readout.

Remark 4.6. It is worth recording the logical relation between activity of r and descent of r . If r itself descends along p (is fiber-constant), then $r(F_t)$ is a singleton for each t and r is automatically inactive for every target that factors through (p, r) ; the loss is not merely inconsequential but nonexistent. The converse fails: r can vary on fibers — structural loss — while being inactive for Q . So activity implies structural loss, but structural loss does not imply activity, and it is precisely this asymmetry that forbids the inference from “noninjective” to “must be completed.”

5. MINIMAL SOUND UNCERTAINTY WITHOUT COMPLETION

Suppose the obstruction is real: a_Q genuinely varies on some fiber, so no readout-only construction returns the correct answer everywhere. It does not follow that the readout is useless for Q . It follows that the readout cannot deliver a *single* answer, and the honest question becomes: what is the most that the readout does license? The answer is a set-valued rule, and there is a canonical least one.

Definition 5.1 (Compatible-answer set). For $t \in T$, the *compatible-answer set* at t is the image of the fiber under the correct-answer map,

$$A_Q(t) := a_Q(F_t) = \{a_Q(s) : s \in S, p(s) = t\} \subseteq A_Q. \quad (11)$$

Since p is surjective, $F_t \neq \emptyset$ and hence $A_Q(t) \neq \emptyset$.

Definition 5.2 (Set-valued rules and uniform soundness). Let $\mathfrak{P}(A_Q)$ denote the *power set* of A_Q , i.e. the collection of all subsets of A_Q . A *set-valued rule* for Q is a map

$$C : T \longrightarrow \mathfrak{P}(A_Q). \quad (12)$$

The rule C is *uniformly sound* if it never excludes the truth:

$$a_Q(s) \in C(p(s)) \quad \text{for every } s \in S. \quad (13)$$

Remark 5.3 ($\mathfrak{P}$ is the power set, not a space of probability measures). We use the deliberately distinct symbol $\mathfrak{P}(A_Q)$ to forestall a conflation that would destroy the meaning of the results in this section. A set-valued rule assigns to each readout value a *subset* of admissible answers. It does not assign a probability distribution, a measure, a density, or a credence. Nothing in [Theorem 5.2](#) presupposes a σ -algebra on A_Q , a measure on S , or a prior on the fibers, and no such structure is introduced anywhere in this paper. See [Theorem 5.6](#) below for why this distinction is substantive and not merely notational.

Theorem 5.4 (The compatible-answer set is the least uniformly sound rule). *Define $C_{\min} : T \rightarrow \mathfrak{P}(A_Q)$ by $C_{\min}(t) := A_Q(t) = a_Q(F_t)$. Then:*

- (1) $C_{\min}$ is uniformly sound;
- (2) for every uniformly sound rule $C : T \rightarrow \mathfrak{P}(A_Q)$ and every $t \in T$,

$$C_{\min}(t) \subseteq C(t). \quad (14)$$

Consequently $C_{\min}$ is the unique pointwise-least uniformly sound set-valued rule for Q given the readout p .

Proof. (1) Let $s \in S$ and put $t := p(s)$. Then $s \in F_t$, so $a_Q(s) \in a_Q(F_t) = C_{\min}(t) = C_{\min}(p(s))$. This is exactly (13).

(2) Let C be uniformly sound and fix $t \in T$. Let $\alpha \in C_{\min}(t) = a_Q(F_t)$ be arbitrary. By definition of the image there exists $s \in F_t$ with $a_Q(s) = \alpha$. Since $s \in F_t$ we have $p(s) = t$, and uniform soundness (13) applied to this s gives

$$\alpha = a_Q(s) \in C(p(s)) = C(t).$$

As $\alpha \in C_{\min}(t)$ was arbitrary, $C_{\min}(t) \subseteq C(t)$.

Uniqueness of the least element: if C and C' are both uniformly sound and both pointwise least, then $C(t) \subseteq C'(t)$ and $C'(t) \subseteq C(t)$ for every t , so $C = C'$. Since $C_{\min}$ is uniformly sound by (1) and pointwise least by (2), it is the unique such rule. $\square$

Corollary 5.5 (Exactness as singleton-valuedness). *Q is exactly solvable from the readout p if and only if $A_Q(t)$ is a singleton for every $t \in T$. In that case the unique readout-only solution $\bar{a}_Q$ is characterized by $\{\bar{a}_Q(t)\} = A_Q(t)$.*

Proof. $A_Q(t) = a_Q(F_t)$ is a singleton for every t iff a_Q is constant on every fiber F_t (nonemptiness of F_t gives that the image is nonempty, so “singleton” is equivalent to “at most one element”), which by Theorem 3.1 is equivalent to exact solvability. If exact solvability holds, then for $s \in F_t$ we have $a_Q(F_t) = \{a_Q(s)\} = \{a_Q(p(s))\} = \{\bar{a}_Q(t)\}$. $\square$

Remark 5.6 (Set-valued uncertainty is not probability, and probability is not recovery). $A_Q(t)$ is the exact residual uncertainty that the readout partition induces about the answer to Q — and it exists prior to, and independently of, any probabilistic structure. Two consequences must be kept apart.

First, imposing a probability law on S , or on the fibers, changes the object under discussion but does not shrink $A_Q(t)$. A distribution on A_Q can *rank* the compatible answers, declaring some more likely than others. It cannot make two distinct correct answers into one answer, and it cannot identify which source state was in fact realized. We are careful about how to phrase the constraint this places on the distribution. We do not claim that an arbitrary probability law must assign positive mass to every element of $A_Q(t)$; a law is free to assign mass zero wherever it likes, and no full-support condition has been imposed here. What is true is conditional: *a probabilistic representation intended not to exclude any source-compatible outcome must have support containing $A_Q(t)$* . Under that stipulation — and it is a stipulation, not a theorem — the support still contains both $a_Q(s_0)$ and $a_Q(s_1)$ whenever $a_Q(s_0) \neq a_Q(s_1)$ and $p(s_0) = p(s_1)$, so the rule has not told us which of s_0, s_1 obtained, and Theorem 3.2 still forbids a correct deterministic answer.

Second, and consequently, a probabilistic prediction answers a *different target*. Replacing Q by the target “what is the distribution of answers, given t and given a specified law on the fibers?” is a legitimate move, but it is a change of question, and it is one that requires additional input — namely the law on the fibers, which is a piece of source-level information not contained in p . It is not a solution to the original Q , and it should not be reported as one. We return to this point in the dynamical setting in Theorem 12.11.

6. READOUT EXTENSIONS AND TARGET SUFFICIENCY

If the readout does not determine the answer, one may enlarge the readout. The following definition says what an enlargement is, and the crucial point is that it need not be an enlargement all the way back to the source.

Definition 6.1 (Readout extension). A *readout extension* (or *enrichment*) of $p: S \rightarrow T$ is a surjective map

$$e: S \longrightarrow E \quad \text{together with a map} \quad u: E \longrightarrow T \quad \text{such that} \quad p = u \circ e. \quad (15)$$

Lemma 6.2 (Basic properties of extensions). *Let (e, E, u) be a readout extension of p . Then:*

- (1) *u is uniquely determined by e and p ;*
- (2) *u is surjective;*
- (3) *the partition of S into fibers of e refines the partition into fibers of p ; equivalently $e(s) = e(s') \implies p(s) = p(s')$.*

Proof. (1) If u, u' both satisfy $p = u \circ e = u' \circ e$, let $x \in E$. By surjectivity of e choose s with $e(s) = x$; then $u(x) = u(e(s)) = p(s) = u'(e(s)) = u'(x)$.

(2) Let $t \in T$. By surjectivity of p choose s with $p(s) = t$; then $u(e(s)) = p(s) = t$, so $t \in \text{Im}(u)$.

(3) If $e(s) = e(s')$ then $p(s) = u(e(s)) = u(e(s')) = p(s')$. □

Theorem 6.2 (3) is the content of the word “extension”: E is at least as informative as T , because anything determined by E determines the T -value. Equivalently, the equivalence relation $\sim_e$ induced by e is contained in $\sim_p$. But E need not be S , and $\sim_e$ need not be equality: E sits somewhere between the readout and the source, an intermediate quotient of S , or, in partition language, a refinement of the readout partition that need not be the discrete partition. The two extreme cases are $E = T$ (with $e = p$, $u = \text{id}$), which adds nothing, and $E = S$ (with $e = \text{id}$, $u = p$), which retains everything. The interesting extensions are strictly between, and the point of the next two sections is that for a given target there is a canonical one, generally far closer to T than to S .

Definition 6.3 (Sufficiency for a target). A readout extension $e: S \rightarrow E$ is *sufficient for Q* , or *Q -sufficient*, if there exists a map

$$b: E \longrightarrow A_Q \quad \text{such that} \quad a_Q = b \circ e. \quad (16)$$

We also call a Q -sufficient readout extension a *completion of p for Q* .

Remark 6.4 (What “completion” means here). The word “completion” is used in [Theorem 6.3](#) in a sense that must be separated from several established and entirely different uses. It is *not* metric completion, not topological or uniform completion, not Cauchy completion, not the completion of a local ring, and not the adjunction of arbitrary hidden variables posited to restore determinism. A completion in the present sense is an *enrichment of the readout that suffices for a nominated target*: a surjection e through which both p and a_Q factor. It is target-relative by construction; an extension may complete p for one target while remaining insufficient for another. And it is constrained: [Theorem 7.1](#) shows that a completion cannot be arbitrary, and [Theorem 8.3](#) shows that there is a least one.

Lemma 6.5 (Sufficiency criterion). *A readout extension $e: S \rightarrow E$ is Q -sufficient if and only if a_Q is constant on every fiber of e ; and then b in (16) is unique.*

Proof. Apply [Theorem 3.1](#) with p replaced by the surjection e and T by E . □

Note the immediate sanity check: $e = \text{id}_S$ (with $E = S$, $u = p$, $b = a_Q$) is always Q -sufficient, so completions exist trivially. The theorems that follow are therefore not existence theorems. They are theorems about *how little* a completion can get away with, and *what it must contain no matter what*.

7. WHAT EVERY SUFFICIENT EXTENSION MUST RETAIN

Theorem 7.1 (Necessary separation). *Let $e: S \rightarrow E$ be any Q -sufficient readout extension. Then e separates every pair of source states with different correct answers:*

$$a_Q(s_0) \neq a_Q(s_1) \implies e(s_0) \neq e(s_1), \quad s_0, s_1 \in S. \quad (17)$$

Proof. We prove the contrapositive. Suppose $e(s_0) = e(s_1) =: x$. By Q -sufficiency there is $b: E \rightarrow A_Q$ with $a_Q = b \circ e$. Then

$$a_Q(s_0) = b(e(s_0)) = b(x) = b(e(s_1)) = a_Q(s_1),$$

since $b(x)$ is a single element of A_Q . Hence $e(s_0) = e(s_1) \implies a_Q(s_0) = a_Q(s_1)$, which is (17). $\square$

Corollary 7.2 (The critical case: fiber-mates with different answers). *Suppose $s_0, s_1 \in S$ satisfy $p(s_0) = p(s_1)$ and $a_Q(s_0) \neq a_Q(s_1)$ — that is, they form an obstructing pair in the sense of (6). Then every Q -sufficient readout extension e satisfies $e(s_0) \neq e(s_1)$, while $u(e(s_0)) = u(e(s_1))$. In particular e is strictly finer than p : the extension is not merely allowed to distinguish s_0 from s_1 , it is compelled to, and it must do so while remaining compatible with the readout value they share.*

Proof. The separation $e(s_0) \neq e(s_1)$ is Theorem 7.1. The identity $u(e(s_i)) = p(s_i)$ from (15), together with $p(s_0) = p(s_1)$, gives $u(e(s_0)) = u(e(s_1))$. Since $e(s_0) \neq e(s_1)$ but p does not separate s_0, s_1 , the partition $\sim_e$ is a strict refinement of $\sim_p$. $\square$

It is useful to restate Theorem 7.1 in partition language, where its content is transparent. Let $\sim_e, \sim_p$, and $\sim_{a_Q}$ denote the equivalence relations on S induced respectively by e , by p , and by a_Q (so $s \sim_{a_Q} s'$ iff $a_Q(s) = a_Q(s')$). Theorem 6.2 (3) says $\sim_e \subseteq \sim_p$, and Theorem 7.1 says $\sim_e \subseteq \sim_{a_Q}$. Thus:

A readout extension is sufficient for Q precisely when its induced partition refines both the readout partition and the answer partition.

Indeed the converse of Theorem 7.1 holds as well and is worth recording, since together they characterize sufficiency exactly.

Proposition 7.3 (Separation characterizes sufficiency). *A readout extension $e: S \rightarrow E$ is Q -sufficient if and only if (17) holds.*

Proof. Necessity is Theorem 7.1. For sufficiency, assume (17). Its contrapositive reads $e(s) = e(s') \implies a_Q(s) = a_Q(s')$, i.e. a_Q is constant on every fiber of e . By Theorem 6.5, e is Q -sufficient. $\square$

We emphasize that Theorem 7.1 does *not* say that e must separate every pair of distinct source states, i.e. that e must be injective, i.e. that the source must be reconstructed. It says only that e must separate pairs *whose answers differ*. Pairs with equal answers may be freely collapsed. That is the loophole that Section 8 exploits to construct a completion that is often dramatically smaller than S , and it is why a target-relative theory of completion is not a disguised demand for full reconstruction.

8. THE JOINT TARGET IMAGE AND UNIVERSAL MINIMALITY

Theorem 7.1 gives a lower bound on the distinctions a completion must carry. We now show that the bound is attained within the category of readout extensions defined above: there is a completion carrying exactly those distinctions and no others, and it receives a canonical map from every other completion. Attainment here is a statement about the abstract construction

$E_Q = \text{Im}(p, a_Q)$; whether a given physical situation supplies an operational realization of that completion is a separate matter, taken up after [Theorem 8.3](#).

Definition 8.1 (Joint target image). Define

$$j_Q: S \longrightarrow T \times A_Q, \quad j_Q(s) := (p(s), a_Q(s)), \quad (18)$$

and let

$$E_Q := \text{Im}(j_Q) = \{ (p(s), a_Q(s)) : s \in S \} \subseteq T \times A_Q. \quad (19)$$

We regard j_Q as a map onto its image, $j_Q: S \rightarrow E_Q$, and write $\pi_T: E_Q \rightarrow T$ and $\pi_A: E_Q \rightarrow A_Q$ for the restrictions to E_Q of the two coordinate projections of $T \times A_Q$.

Proposition 8.2 (E_Q is a Q -sufficient readout extension). *The map $j_Q: S \rightarrow E_Q$ is a surjective readout extension of p , with structure map $u = \pi_T$, and it is sufficient for Q , with $b = \pi_A$. Explicitly,*

$$p = \pi_T \circ j_Q, \quad a_Q = \pi_A \circ j_Q. \quad (20)$$

Moreover the fibers of j_Q are exactly the sets $F_t \cap a_Q^{-1}(\alpha)$ for $(t, \alpha) \in E_Q$.

Proof. $j_Q: S \rightarrow E_Q$ is surjective by construction, since E_Q was defined as the image of j_Q . No nonemptiness hypothesis is needed: if $S = \emptyset$ then $E_Q = \emptyset$ as well, and the identities below hold vacuously.

For (20), let $s \in S$. Then $j_Q(s) = (p(s), a_Q(s)) \in E_Q$, and applying the coordinate projections,

$$\pi_T(j_Q(s)) = \pi_T(p(s), a_Q(s)) = p(s), \quad \pi_A(j_Q(s)) = \pi_A(p(s), a_Q(s)) = a_Q(s).$$

The first identity says $p = \pi_T \circ j_Q$, so (j_Q, E_Q, π_T) is a readout extension in the sense of [Theorem 6.1](#). The second says $a_Q = \pi_A \circ j_Q$, so j_Q is Q -sufficient in the sense of [Theorem 6.3](#), with $b = \pi_A$.

For the fibers: $j_Q(s) = (t, \alpha)$ iff $p(s) = t$ and $a_Q(s) = \alpha$, i.e. iff $s \in F_t \cap a_Q^{-1}(\alpha)$. $\square$

We have thus verified directly, and not by appeal to any universal property, that E_Q does the job: it determines the readout value and it determines the correct answer. The next theorem says it does nothing more.

Theorem 8.3 (Universal factorization: E_Q is the coarsest Q -sufficient extension). *Let $e: S \rightarrow E$ be any surjective readout extension of p that is sufficient for Q , with structure map $u: E \rightarrow T$ satisfying $p = u \circ e$ and solution map $b: E \rightarrow A_Q$ satisfying $a_Q = b \circ e$. Then there is a unique map*

$$\varphi_E: E \longrightarrow E_Q \quad \text{such that} \quad j_Q = \varphi_E \circ e, \quad (21)$$

namely

$$\varphi_E(x) = (u(x), b(x)), \quad x \in E, \quad (22)$$

and φ_E is surjective. Moreover φ_E intertwines the structure maps and the solution maps:

$$\pi_T \circ \varphi_E = u, \quad \pi_A \circ \varphi_E = b. \quad (23)$$

Proof. We proceed in five steps.

Step 1: u and b are uniquely determined by e . This is [Theorem 6.2 \(1\)](#) for u , and [Theorem 6.5](#) for b ; in both cases uniqueness follows from surjectivity of e . We record the argument for b once more, since it is used below: if b, b' both satisfy $a_Q = b \circ e = b' \circ e$ and $x \in E$, choose s with $e(s) = x$; then $b(x) = b(e(s)) = a_Q(s) = b'(e(s)) = b'(x)$. Hence the formula (22) is unambiguous.

Step 2: φ_E is well defined as a map into E_Q . Define $\psi: E \rightarrow T \times A_Q$ by $\psi(x) := (u(x), b(x))$; this is manifestly a well-defined map into the product. We must check that its values lie in the subset E_Q . Let $x \in E$. By surjectivity of e there is $s \in S$ with $e(s) = x$. Then

$$\psi(x) = (u(e(s)), b(e(s))) = (p(s), a_Q(s)) = j_Q(s) \in E_Q, \quad (24)$$

using $p = u \circ e$ and $a_Q = b \circ e$. Hence ψ takes values in E_Q , and we may define $\varphi_E: E \rightarrow E_Q$ by $\varphi_E := \psi$, i.e. by (22). (Note that (24) did not require the choice of s to be canonical: $\psi(x)$ was already defined without reference to s , and s was used only to certify membership in E_Q . So no well-definedness issue arises from the choice of preimage.)

Step 3: the factorization $j_Q = \varphi_E \circ e$. Let $s \in S$. Then by (22) and the two factorizations through e ,

$$\varphi_E(e(s)) = (u(e(s)), b(e(s))) = (p(s), a_Q(s)) = j_Q(s).$$

Since s was arbitrary, $j_Q = \varphi_E \circ e$.

Step 4: φ_E is surjective. Let $y \in E_Q$. By (19) there is $s \in S$ with $j_Q(s) = y$. Put $x := e(s) \in E$. By Step 3, $\varphi_E(x) = \varphi_E(e(s)) = j_Q(s) = y$. Hence φ_E is onto E_Q .

Step 5: uniqueness of φ_E . Suppose $\varphi, \varphi': E \rightarrow E_Q$ both satisfy $j_Q = \varphi \circ e = \varphi' \circ e$. Let $x \in E$. By surjectivity of e choose $s \in S$ with $e(s) = x$. Then

$$\varphi(x) = \varphi(e(s)) = j_Q(s) = \varphi'(e(s)) = \varphi'(x).$$

As x was arbitrary, $\varphi = \varphi'$. This is precisely the point at which surjectivity of e — part of Theorem 6.1 — is essential: without it, φ_E would be unconstrained on $E \setminus \text{Im}(e)$.

Finally, (23) is immediate from (22): for $x \in E$, $\pi_T(\varphi_E(x)) = \pi_T(u(x), b(x)) = u(x)$ and $\pi_A(\varphi_E(x)) = \pi_A(u(x), b(x)) = b(x)$. $\square$

Corollary 8.4 (Information ordering). *Every Q -sufficient readout extension $e: S \rightarrow E$ carries at least the distinctions encoded by the pair (p, a_Q) : for all $s_0, s_1 \in S$,*

$$j_Q(s_0) \neq j_Q(s_1) \implies e(s_0) \neq e(s_1). \quad (25)$$

Equivalently, $\sim_e \subseteq \sim_{j_Q}$: the partition induced by e refines the partition induced by j_Q . Thus E_Q is the coarsest readout extension sufficient for Q .

Proof. By Theorem 8.3, $j_Q = \varphi_E \circ e$. If $e(s_0) = e(s_1)$ then $j_Q(s_0) = \varphi_E(e(s_0)) = \varphi_E(e(s_1)) = j_Q(s_1)$, which is the contrapositive of (25). Hence $\sim_e \subseteq \sim_{j_Q}$. Since j_Q is itself Q -sufficient (Theorem 8.2), it is a Q -sufficient extension whose partition is refined by that of every other, i.e. it is coarsest among Q -sufficient extensions. $\square$

Several points about Theorem 8.3 and Theorem 8.4 require care, and we state them explicitly.

The direction of the universal map. The canonical map runs

$$E \longrightarrow E_Q,$$

from an arbitrary sufficient extension *to* the joint target image — not from E_Q to E . This direction is forced and is not a convention. A surjection $E \rightarrow E_Q$ expresses that E is at least as fine as E_Q , since it says the E -value determines the E_Q -value. There is in general *no* map $E_Q \rightarrow E$ over S : E may retain distinctions that E_Q has discarded — distinctions irrelevant to Q but possibly relevant to other targets — and no map out of E_Q could recover them. Reversing the arrow would assert the opposite of what is true, namely that E_Q is the finest rather than the coarsest sufficient extension.

E_Q does not reconstruct the source. The joint target image retains exactly two things: the readout value $p(s)$, and the correct answer $a_Q(s)$. It discards everything else, and in particular it identifies any two source states that agree on both. This is usually a massive collapse. In

[Section 11](#) the target is the spectrum of a charged Laplacian, and E_Q turns out to be a one-dimensional interval’s worth of data — the inversion quotient of holonomy — even though the source moduli space $U(1)$ is strictly finer and the raw connections finer still. Nothing in [Theorem 8.3](#) says or implies that the source is recoverable; on the contrary, [Theorem 8.4](#) says the source is generally *not* recoverable from E_Q , because id_S is one of the sufficient extensions e and the canonical map goes $S \rightarrow E_Q$, not back.

One target at a time. [Theorem 8.3](#) is a statement about a single correct-answer map a_Q . If one nominates a family of targets, the coarsest extension sufficient for all of them simultaneously must, by [Theorem 7.1](#) applied to each member, separate every pair of source states on which *some* member’s answer differs. It is therefore at least as fine as the completion required by each individual member, and it may be strictly finer. It need not be: should the remaining members contribute no distinctions beyond those already forced by one of them, the family completion coincides with that member’s individual completion. Which case obtains depends on the family, and the construction in general requires machinery we do not develop here; that is the subject of [\[9\]](#). Within the present paper, “minimal” always means “minimal for the nominated Q ,” and we will see in [Section 11](#) that the minimal completion for the spectral target is strictly insufficient for a transport target on the same source space.

What the universal property does and does not supply. [Theorem 8.3](#) establishes that E_Q is the coarsest sufficient extension in the category of set-theoretic source quotients and readout extensions defined in this paper: among all Q -sufficient extensions, it is the one carrying the fewest distinctions, and every other maps canonically onto it. Because the construction of $E_Q = \text{Im}(p, a_Q)$ refers explicitly to the correct-answer map a_Q , the theorem is an information-theoretic characterization of the least distinctions any sufficient extension must carry; it identifies the floor abstractly, from the data (p, a_Q) . It does not, however, furnish an operational procedure for extracting a_Q from the original readout p — indeed [Theorems 3.2](#) and [3.4](#) show that no such extraction from p alone exists when Q is obstructed — and it does not by itself exhibit a physically measurable, experimentally accessible, or independently specified realization of E_Q . A concrete physical repair is a separate task: it must independently provide an extension $e: S \rightarrow E$ and an answer map $b: E \rightarrow A_Q$ through which a_Q factors, and must justify that e and b are available in the situation at hand. The universal property then locates any such repair relative to the floor — it certifies that e carries at least the distinctions of E_Q — but the provision and justification of e and b lie outside what the abstract quotient construction delivers. The flat-connection example of [Sections 9](#) and [11](#) supplies exactly such an independently specified extension, the charged spectrum, and it is that concrete construction, not [Theorem 8.3](#) alone, that makes the repair physical.

9. FLAT $U(1)$ CONNECTIONS ON THE CIRCLE

We now instantiate the framework in a setting where every ingredient — the reduction by gauge, the deliberately restricted readout, the fiber, the target, and the minimal completion — is explicit and computable.

9.1. The reduced moduli space. Let S^1 be a circle of circumference $L > 0$, with coordinate $x \in [0, L]$ and endpoints identified, $x \sim x + L$. Consider the trivial principal $U(1)$ -bundle over S^1 and its connections, described in the standard way by real one-forms A on S^1 ; the associated covariant derivative on a unit-charge complex scalar is $\partial_x - iA_x$, and the curvature is the two-form $F = dA$.

We state at once the feature that makes this example the sharpest possible illustration of the theory, rather than treating it as an obstacle to be maneuvered around. *Every two-form*

on a one-dimensional manifold vanishes identically, since $\Lambda^2 T_x^* S^1 = 0$ for dimensional reasons. Hence $F = dA = 0$ for *every* connection A on S^1 , with no flatness hypothesis required and no special choice of representative involved. The curvature readout on this space is therefore identically zero as a matter of dimension.

This is not a degeneracy to be apologized for; it is exactly the point. Curvature is a *local* invariant, and on the circle it is vacuous. Holonomy is a *global* invariant, and on the circle it is not: connections with identically vanishing field strength are nonetheless distinguished, gauge-invariantly, by the phase a charge accumulates on traversing the loop. The circle is thus the cleanest available laboratory for a global-holonomy-versus-local- curvature separation, and that separation is precisely what [Sections 10 and 11](#) will show a physical target can detect.

Consider now the one-parameter family of connection one-forms

$$A_\theta = \frac{\theta}{L} dx, \quad \theta \in \mathbb{R}. \quad (26)$$

Their field strength vanishes,

$$F_\theta = dA_\theta = d\left(\frac{\theta}{L}\right) \wedge dx = 0, \quad (27)$$

which holds for the two independent reasons just discussed: the coefficient θ/L is constant, and, more fundamentally, any two-form on S^1 is zero regardless. Every connection on the trivial bundle over S^1 is gauge-equivalent to exactly one such A_θ modulo the identifications described below; this is the standard parametrization of the moduli space of $U(1)$ connections on a circle [\[3, 4, 2\]](#). Since all connections here are flat, we retain the customary name *moduli space of flat connections* for it.

9.2. Gauge convention and the reduction. Gauge transformations act by

$$A \mapsto A + d\lambda, \quad (28)$$

where $g = e^{i\lambda}: S^1 \rightarrow U(1)$ is a smooth, single-valued map. Single-valuedness of g , not of λ , is the physical requirement, and it is exactly this that produces the identification we need. Writing $\lambda(x) = 2\pi kx/L$ with $k \in \mathbb{Z}$, the map

$$g_k(x) = \exp\left(\frac{2\pi i k x}{L}\right) \quad (29)$$

is smooth and single-valued on S^1 , since $g_k(x+L) = g_k(x)e^{2\pi i k} = g_k(x)$, even though λ itself increases by $2\pi k$ around the circle. Such g_k are the *large gauge transformations*; they are not connected to the identity, and they act on [\(26\)](#) by

$$A_\theta \mapsto A_\theta + d\lambda = \frac{\theta}{L} dx + \frac{2\pi k}{L} dx = A_{\theta+2\pi k}. \quad (30)$$

Consequently the parameters θ and $\theta + 2\pi k$ describe *the same physical configuration*, and the honest parameter space of flat connections is the quotient

$$\theta \in \mathbb{R}/2\pi\mathbb{Z}. \quad (31)$$

Small gauge transformations, with λ single-valued, change A_θ by an exact form $d\lambda$ with $\oint d\lambda = 0$; these do not alter $\theta \bmod 2\pi$ and, after fixing the constant gauge [\(26\)](#), do not move us within the family at all.

We may now define the reduced source space of this example in accordance with the standing convention of [Section 2](#):

$$S_{\text{flat}} := \{\text{flat } U(1) \text{ connections on } S^1\} / \{\text{gauge}\} \cong \mathbb{R}/2\pi\mathbb{Z} \cong U(1). \quad (32)$$

It is essential to the correctness of everything below that the identification $\theta \sim \theta + 2\pi k$ has *already been performed*. Distinct points of S_{flat} — distinct values of $e^{i\theta}$ — are gauge-inequivalent physical configurations, not gauge copies of one another. Any obstruction we exhibit between two such points is therefore an obstruction between genuinely different physical states, and cannot be dissolved by a change of gauge.

9.3. Holonomy. The gauge-invariant content of A_θ is its holonomy around the circle,

$$h(\theta) = \exp\left(i \oint_{S^1} A_\theta\right) = \exp\left(i \int_0^L \frac{\theta}{L} dx\right) = e^{i\theta} \in \text{U}(1). \quad (33)$$

Under a large gauge transformation (30), $h \mapsto e^{i(\theta+2\pi k)} = e^{i\theta}$: the holonomy is invariant, as it must be. Under a small gauge transformation the integral $\oint d\lambda$ vanishes and h is again unchanged. So h is a well-defined function on S_{flat} , and by (32) it is in fact a bijection $S_{\text{flat}} \rightarrow \text{U}(1)$. We will accordingly use $e^{i\theta}$, or $\theta \bmod 2\pi$, as the coordinate on the reduced source space.

That the holonomy is physically consequential — that it is not a gauge artifact but an observable feature of the configuration — is the content of the Aharonov–Bohm effect [1, 11], and it is worth stating the geometry of that effect accurately.

In the standard physical realization, the apparatus consists of a solenoid or flux tube together with the region exterior to it. The charged particle is confined to the *accessible region*, which excludes the flux-carrying interior and is therefore multiply connected; on that accessible region, and only there, the connection is flat and the field strength vanishes. The excluded region does carry field: the flux Φ threading the solenoid is nonzero, and it is exactly that flux which fixes the phase. It would be wrong to say that the electromagnetic field vanishes everywhere in the apparatus — it does not — and we do not say so. What is true, and what matters, is that the interference phase measured by the particle is determined by the gauge-invariant holonomy of the connection around a loop in the accessible region, a loop along which the particle never encounters any field.

The circle S^1 of this section models such an accessible loop. Its states $e^{i\theta}$ are distinguished by holonomy alone, with no local field anywhere along them, which is the physical warrant for treating $S_{\text{flat}} \cong \text{U}(1)$ as a space of genuinely distinct physical states rather than as gauge copies of a single trivial one.

9.4. The curvature readout. Define the readout to be the field strength:

$$p_F: S_{\text{flat}} \longrightarrow \{0\}, \quad p_F(\theta) = F_\theta = 0. \quad (34)$$

By (27) this is well defined and, trivially, surjective onto the one-point set $\{0\}$: the curvature vanishes on every connection, for the dimensional reason recorded above, so $\{0\}$ is the full set of curvature values realized. Its single fiber is everything:

$$F_0 = p_F^{-1}(0) = S_{\text{flat}} \cong \text{U}(1). \quad (35)$$

This is about as lossy as a readout can be. It is also, for exactly that reason, an ideal test case: any target whose correct answer is nonconstant on $\text{U}(1)$ is obstructed by p_F , and any target whose correct answer is constant on $\text{U}(1)$ is not.

Remark 9.1 (The curvature readout is a deliberate restriction, not a claim about physics). We are not asserting that curvature exhausts what can be observed of a $\text{U}(1)$ gauge field. It plainly does not: in the Aharonov–Bohm interference experiment the measured interference phase is determined by the gauge-invariant holonomy, and an experiment sensitive to holonomy would supply a strictly finer readout than p_F — indeed the finest possible one on S_{flat} , since h is injective there. The map p_F is the readout *under examination*. It is nominated, not

discovered, and it is nominated precisely because its extreme degeneracy makes the target-relative machinery of Sections 3, 4, 7 and 8 maximally visible. A reader who objects that “real physics sees the holonomy” is agreeing with the analysis: enlarging the readout to include holonomy is exactly the move that Section 6 formalizes, and Section 11 shows that for the spectral target one need not even go that far.

Example 9.2 (A target unobstructed by the curvature readout). Take $A_Q = \{\text{yes}, \text{no}\}$ and $a_Q(\theta) = \text{yes}$ iff the connection is flat. Then a_Q is constant on S_{flat} , hence constant on the fiber F_0 , hence by Theorem 3.1 exactly solvable from p_F : the readout-only solution is $\bar{a}_Q(0) = \text{yes}$. The curvature readout, degenerate as it is, answers this target perfectly. The next two sections exhibit a target on the same source space, with the same readout, for which it fails.

10. THE CHARGED COVARIANT LAPLACIAN AND THE CURVATURE OBSTRUCTION

10.1. Operator, domain, and spectrum. Fix $\theta \in \mathbb{R}$ representing a point of S_{flat} , and consider a complex scalar field of unit charge on S^1 minimally coupled to A_θ . The gauge-covariant derivative is

$$D_x := \partial_x - \frac{i\theta}{L}, \quad (36)$$

and the covariant Laplacian, or gauged kinetic operator, is

$$H_\theta := -D_x^2 = -\left(\partial_x - \frac{i\theta}{L}\right)^2. \quad (37)$$

We specify the operator precisely. Let

$$\mathcal{H} := L^2(S^1) = L^2([0, L], dx) \quad \text{with periodic identification } x \sim x + L, \quad (38)$$

with inner product $\langle \psi, \phi \rangle = \int_0^L \overline{\psi(x)} \phi(x) dx$, and take the operator domain to be the periodic Sobolev space

$$\mathcal{D}(H_\theta) := H_{\text{per}}^2(S^1) = \left\{ \psi \in L^2(S^1) : \sum_{n \in \mathbb{Z}} (1 + n^2)^2 |\hat{\psi}(n)|^2 < \infty \right\}, \quad (39)$$

where $\hat{\psi}(n)$ are the Fourier coefficients of ψ with respect to the orthonormal basis introduced below. Equivalently, $\mathcal{D}(H_\theta)$ consists of $\psi \in L^2$ whose distributional derivatives up to order two lie in L^2 and which, together with their first derivatives, satisfy the periodic boundary conditions $\psi(0) = \psi(L)$, $\psi'(0) = \psi'(L)$; we write $H_{\text{per}}^2(S^1)$ for this space. With this domain, H_θ is self-adjoint, and the cleanest justification is Fourier diagonalization rather than a perturbative argument. Let $\mathcal{F}: L^2(S^1) \rightarrow \ell^2(\mathbb{Z})$, $(\mathcal{F}\psi)(n) = \hat{\psi}(n) = \langle \psi_n, \psi \rangle$, be the unitary Fourier transform associated with the orthonormal basis $\{\psi_n\}_{n \in \mathbb{Z}}$ of (41) below. A direct computation, carried out in Theorem 10.1, gives $H_\theta \psi_n = \lambda_n(\theta) \psi_n$ with $\lambda_n(\theta) = ((2\pi n - \theta)/L)^2$. Hence $\mathcal{F}$ conjugates H_θ into the *multiplication operator* by the real sequence $(\lambda_n(\theta))_{n \in \mathbb{Z}}$ on $\ell^2(\mathbb{Z})$:

$$(\mathcal{F}H_\theta\psi)(n) = \lambda_n(\theta) \hat{\psi}(n), \quad \mathcal{F}(\mathcal{D}(H_\theta)) = \left\{ \hat{\psi} \in \ell^2(\mathbb{Z}) : \sum_{n \in \mathbb{Z}} \lambda_n(\theta)^2 |\hat{\psi}(n)|^2 < \infty \right\}. \quad (40)$$

Multiplication by a real sequence, taken on its maximal domain, is self-adjoint; and since $\lambda_n(\theta) \asymp (2\pi n/L)^2$ as $|n| \rightarrow \infty$, the maximal domain in (40) coincides with $\mathcal{F}(H_{\text{per}}^2(S^1))$, so $\mathcal{D}(H_\theta) = H_{\text{per}}^2(S^1)$ as claimed. Because the multiplier $\lambda_n(\theta) \rightarrow +\infty$, the operator has compact resolvent and purely discrete spectrum [6, 7]. No relative-boundedness estimate is needed, and none is claimed: the first-order term $\frac{2i\theta}{L}\partial_x$ is *not* bounded on $L^2(S^1)$, and the diagonalization above avoids any appeal to such a statement.

10.2. The spectrum.

Proposition 10.1 (Eigenvalues and multiplicities). *The functions*

$$\psi_n(x) = \frac{1}{\sqrt{L}} e^{2\pi i n x / L}, \quad n \in \mathbb{Z}, \quad (41)$$

form an orthonormal basis of $\mathcal{H}$ consisting of eigenvectors of H_θ , and

$$D_x \psi_n = \frac{i(2\pi n - \theta)}{L} \psi_n, \quad H_\theta \psi_n = \lambda_n(\theta) \psi_n \quad \text{with} \quad \lambda_n(\theta) = \left(\frac{2\pi n - \theta}{L} \right)^2. \quad (42)$$

Consequently the spectrum of H_θ is purely discrete and, counted with multiplicity, is the multiset

$$\text{Spec}(H_\theta) = \left\{ \left\{ \lambda_n(\theta) = \left(\frac{2\pi n - \theta}{L} \right)^2 : n \in \mathbb{Z} \right\} \right\}, \quad (43)$$

where $\{\{\cdot\}\}$ denotes a multiset: the eigenvalue λ occurs in $\text{Spec}(H_\theta)$ with multiplicity

$$m_\theta(\lambda) = \# \{ n \in \mathbb{Z} : \lambda_n(\theta) = \lambda \}. \quad (44)$$

Proof. Orthonormality and completeness of $\{\psi_n\}$ in $L^2(S^1)$ is the classical Fourier basis theorem. Each ψ_n is smooth and periodic, hence lies in $H^2(S^1) = \mathcal{D}(H_\theta)$. Compute

$$D_x \psi_n = \partial_x \psi_n - \frac{i\theta}{L} \psi_n = \frac{2\pi i n}{L} \psi_n - \frac{i\theta}{L} \psi_n = \frac{i(2\pi n - \theta)}{L} \psi_n,$$

which is the first identity in (42). Applying D_x once more,

$$D_x^2 \psi_n = \frac{i(2\pi n - \theta)}{L} D_x \psi_n = \left(\frac{i(2\pi n - \theta)}{L} \right)^2 \psi_n = - \left(\frac{2\pi n - \theta}{L} \right)^2 \psi_n,$$

so $H_\theta \psi_n = -D_x^2 \psi_n = \left(\frac{2\pi n - \theta}{L} \right)^2 \psi_n$, which is the second identity. Since $\{\psi_n\}_{n \in \mathbb{Z}}$ is a complete orthonormal system of eigenvectors with real eigenvalues $\lambda_n(\theta) \rightarrow +\infty$ as $|n| \rightarrow \infty$, the Fourier transform $\mathcal{F}$ carries H_θ with domain (39) onto multiplication by $(\lambda_n(\theta))_{n \in \mathbb{Z}}$ on its maximal domain, as in (40). That operator is self-adjoint with compact resolvent, hence so is H_θ , and its spectrum, counted with multiplicity, is the multiset (43) with multiplicities (44).

Finally we check that (43) is well defined on the reduced source space, i.e. invariant under $\theta \mapsto \theta + 2\pi k$: the substitution $n \mapsto n + k$ gives $\lambda_{n+k}(\theta + 2\pi k) = \left(\frac{2\pi(n+k) - \theta - 2\pi k}{L} \right)^2 = \lambda_n(\theta)$, and $n \mapsto n + k$ is a bijection of $\mathbb{Z}$, so the multiset (43) is unchanged, with all multiplicities preserved. Thus $\text{Spec}(H_\theta)$ depends only on $\theta \bmod 2\pi$, i.e. only on the point of S_{flat} . $\square$

Remark 10.2 (Multiplicities are part of the answer). We record (43) as a multiset, not as a subset of $\mathbb{R}$, because the multiplicity data is physically real and is not recoverable from the underlying set. For generic θ , distinct n give distinct $\lambda_n(\theta)$ and every eigenvalue is simple; but at $\theta = 0$ the eigenvalues λ_n and λ_{-n} coincide for all n , so every nonzero eigenvalue has multiplicity two while $\lambda_0 = 0$ is simple; and at $\theta = \pi$ one finds $\lambda_n(\pi) = \left(\frac{2\pi n - \pi}{L} \right)^2 = \left(\frac{\pi(2n-1)}{L} \right)^2$, so $\lambda_n(\pi) = \lambda_{1-n}(\pi)$ and every eigenvalue has multiplicity two. Discarding multiplicities would erase this structure and, in particular, would change the target: “the set of eigenvalues” and “the spectrum with multiplicities” are different correct-answer maps and may have different completions. We take the latter throughout.

10.3. The spectral target and the curvature obstruction.

Definition 10.3 (Spectral target). Let A_{spec} be the set of multisets of nonnegative reals arising as spectra of the operators H_θ , and define the correct-answer map

$$a_{\text{spec}}: S_{\text{flat}} \longrightarrow A_{\text{spec}}, \quad a_{\text{spec}}(\theta) := \text{Spec}(H_\theta) \quad (45)$$

with $\text{Spec}(H_\theta)$ as in (43). By the last paragraph of the proof of [Theorem 10.1](#), a_{spec} is well defined on S_{flat} , i.e. on $\theta \bmod 2\pi$. We denote this target by Q_{spec} .

Theorem 10.4 (The spectral target is obstructed by the curvature readout). *Q_{spec} is not exactly solvable from the curvature readout p_F of (34). Explicitly, the reduced source states $\theta = 0$ and $\theta = \pi$ satisfy*

$$p_F(0) = p_F(\pi) = 0, \quad \text{but} \quad \text{Spec}(H_0) \neq \text{Spec}(H_\pi). \quad (46)$$

Indeed $0 \in \text{Spec}(H_0)$ while $0 \notin \text{Spec}(H_\pi)$, and moreover $\min \text{Spec}(H_0) = 0 < \pi^2/L^2 = \min \text{Spec}(H_\pi)$.

Proof. First, $\theta = 0$ and $\theta = \pi$ are distinct points of $S_{\text{flat}} = \mathbb{R}/2\pi\mathbb{Z}$, since $\pi \not\equiv 0 \pmod{2\pi}$; equivalently their holonomies $h(0) = 1$ and $h(\pi) = -1$ are distinct elements of $\text{U}(1)$. They are therefore genuinely gauge-inequivalent physical configurations, not two labels for one state.

Both are flat: by (27), $F_0 = F_\pi = 0$, so $p_F(0) = p_F(\pi) = 0$ and the two states lie in the same (indeed the only) fiber of p_F .

Now compute. By (42) with $\theta = 0$,

$$\lambda_n(0) = \left(\frac{2\pi n}{L} \right)^2,$$

and taking $n = 0$ gives $\lambda_0(0) = 0$, so $0 \in \text{Spec}(H_0)$; the corresponding eigenvector is the constant function $\psi_0 = L^{-1/2}$, and indeed $D_x \psi_0 = \partial_x \psi_0 = 0$ when $\theta = 0$.

By (42) with $\theta = \pi$,

$$\lambda_n(\pi) = \left(\frac{2\pi n - \pi}{L} \right)^2 = \frac{\pi^2}{L^2} (2n - 1)^2.$$

Since $2n - 1$ is odd for every $n \in \mathbb{Z}$, we have $(2n - 1)^2 \geq 1$, hence

$$\lambda_n(\pi) \geq \frac{\pi^2}{L^2} > 0 \quad \text{for every } n \in \mathbb{Z}, \quad (47)$$

with equality precisely when $2n - 1 = \pm 1$, i.e. $n \in \{0, 1\}$. Therefore $0 \notin \text{Spec}(H_\pi)$ and $\min \text{Spec}(H_\pi) = \pi^2/L^2$, attained with multiplicity two.

Since $0 \in \text{Spec}(H_0)$ and $0 \notin \text{Spec}(H_\pi)$, the two multisets differ, so $a_{\text{spec}}(0) \neq a_{\text{spec}}(\pi)$ while $p_F(0) = p_F(\pi)$. This is exactly the hypothesis (6) of [Theorem 3.2](#), whence Q_{spec} is not exactly solvable from p_F : no map $b: \{0\} \rightarrow A_{\text{spec}}$ — that is, no single spectrum assigned to the curvature value 0 — is correct for both $\theta = 0$ and $\theta = \pi$. $\square$

Corollary 10.5 (Holonomy is target-active for the spectrum). *Let $r = h: S_{\text{flat}} \rightarrow \text{U}(1)$, $h(\theta) = e^{i\theta}$, be the holonomy relation, and let $\Phi_{\text{spec}}(0, e^{i\theta}) := \text{Spec}(H_\theta)$, which is well defined by [Theorem 10.1](#), so that $a_{\text{spec}} = \Phi_{\text{spec}} \circ (p_F, h)$ in the sense of (7). Then h is active for Q_{spec} on the single fiber F_0 of p_F .*

Proof. Take $s_0 = \theta = 0$ and $s_1 = \theta = \pi$ in [Theorem 4.2](#). Both lie in F_0 by (35), and by [Theorem 10.4](#), $\Phi_{\text{spec}}(0, h(0)) = \text{Spec}(H_0) \neq \text{Spec}(H_\pi) = \Phi_{\text{spec}}(0, h(\pi))$, which is (8). $\square$

The contrast with [Theorem 9.2](#) is the whole point. The readout p_F is the same in both cases, and the structural loss — the collapse of all of $U(1)$ to a point — is the same. For the flatness target the loss is inactive and the readout is fully adequate. For the spectral target the same loss is active, and by [Theorem 4.3](#) it obstructs. Whether completion is necessary is a fact about the pair, never about the readout alone.

Remark 10.6 (Physical reading). The physical content of [Theorem 10.4](#) is the Aharonov–Bohm phenomenon in spectral form [1, 11]. Along the circle the field strength vanishes — identically, and for dimensional reasons — so there is no local electromagnetic field anywhere on the particle’s accessible configuration space for it to feel. Nevertheless the charged particle’s energy levels depend on θ : the covariant derivative (36) twists the periodicity condition felt by the wavefunction, shifting the Fourier momenta from $2\pi n/L$ to $(2\pi n - \theta)/L$. At $\theta = \pi$ the shift is maximal, the momentum lattice is displaced by half a spacing, no mode has zero covariant momentum, and the ground-state energy is lifted from 0 to π^2/L^2 . The obstruction is not an artifact of the formalism; it is a measurable spectral shift produced by information that curvature does not record.

11. THE MINIMAL COMPLETION FOR THE SPECTRAL TARGET

Curvature is insufficient for Q_{spec} . Holonomy is manifestly sufficient, since by [Theorem 10.1](#) the spectrum is a function of $\theta \bmod 2\pi$. The question this section answers is the sharp one: *how much less than holonomy is enough?* The answer is: strictly less. The spectral target sees holonomy only up to inversion.

11.1. Holonomy is sufficient.

Proposition 11.1 (Holonomy completes the curvature readout for the spectrum). *The identity map $e = \text{id}: S_{\text{flat}} \rightarrow S_{\text{flat}}$, equivalently the holonomy map $h: S_{\text{flat}} \rightarrow U(1)$, $h(\theta) = e^{i\theta}$, is a readout extension of p_F that is sufficient for Q_{spec} .*

Proof. h is a bijection $S_{\text{flat}} \rightarrow U(1)$ by (32) and (33), hence surjective, and $p_F = u \circ h$ with $u: U(1) \rightarrow \{0\}$ the constant map, so h is a readout extension in the sense of [Theorem 6.1](#). Define $b: U(1) \rightarrow A_{\text{spec}}$ by $b(e^{i\theta}) := \text{Spec}(H_\theta)$; this is well defined because, as shown at the end of the proof of [Theorem 10.1](#), $\text{Spec}(H_\theta)$ depends only on $\theta \bmod 2\pi$, i.e. only on $e^{i\theta}$. Then $a_{\text{spec}} = b \circ h$, so h is Q_{spec} -sufficient by [Theorem 6.3](#). $\square$

11.2. Spectral rigidity: the exact equivalence.

Theorem 11.2 (Isospectrality iff $\theta' \equiv \pm\theta$). *For $\theta, \theta' \in \mathbb{R}$,*

$$\text{Spec}(H_\theta) = \text{Spec}(H_{\theta'}) \iff \theta' \equiv \theta \text{ or } \theta' \equiv -\theta \pmod{2\pi}, \quad (48)$$

where the spectra are compared as multisets, with multiplicities.

Proof. ($\Leftarrow$) *Reverse direction.* Suppose first $\theta' \equiv \theta \pmod{2\pi}$. Then $\text{Spec}(H_{\theta'}) = \text{Spec}(H_\theta)$ by the invariance established at the end of the proof of [Theorem 10.1](#).

Suppose next $\theta' \equiv -\theta \pmod{2\pi}$. By the same invariance it suffices to treat $\theta' = -\theta$. Then for every $n \in \mathbb{Z}$,

$$\lambda_{-n}(-\theta) = \left(\frac{2\pi(-n) - (-\theta)}{L} \right)^2 = \left(\frac{-(2\pi n - \theta)}{L} \right)^2 = \left(\frac{2\pi n - \theta}{L} \right)^2 = \lambda_n(\theta). \quad (49)$$

The map $n \mapsto -n$ is a bijection of $\mathbb{Z}$ onto itself. Hence (49) exhibits a bijection between the indexing set of $\text{Spec}(H_{-\theta})$ and that of $\text{Spec}(H_\theta)$ carrying each eigenvalue to an equal eigenvalue.

A multiset is determined by its indexed family up to bijective reindexing, so the two multisets are equal, *with all multiplicities*:

$$m_{-\theta}(\lambda) = \#\{n : \lambda_n(-\theta) = \lambda\} = \#\{n : \lambda_{-n}(-\theta) = \lambda\} = \#\{n : \lambda_n(\theta) = \lambda\} = m_\theta(\lambda).$$

So $\text{Spec}(H_{-\theta}) = \text{Spec}(H_\theta)$. Concretely, the reindexing $n \mapsto -n$ is implemented on the Hilbert space by complex conjugation, $\psi_n \mapsto \overline{\psi_n} = \psi_{-n}$, which is an antiunitary intertwiner between H_θ and $H_{-\theta}$; but the elementary reindexing argument above is already complete and requires nothing further.

($\Rightarrow$) *Forward direction.* Suppose $\text{Spec}(H_\theta) = \text{Spec}(H_{\theta'})$. We recover θ up to sign and 2π from the spectrum, using the lowest eigenvalue.

Define, for any $\theta \in \mathbb{R}$, the *principal representative*

$$\alpha(\theta) \in [0, \pi] \quad \text{the unique number in } [0, \pi] \text{ with } \theta \equiv \alpha(\theta) \text{ or } \theta \equiv -\alpha(\theta) \pmod{2\pi}. \quad (50)$$

Existence and uniqueness of $\alpha(\theta)$: reduce $\theta \bmod 2\pi$ to a representative $\tilde{\theta} \in [0, 2\pi)$. If $\tilde{\theta} \in [0, \pi]$ set $\alpha := \tilde{\theta}$; if $\tilde{\theta} \in (\pi, 2\pi)$ set $\alpha := 2\pi - \tilde{\theta} \in (0, \pi)$, and then $\theta \equiv \tilde{\theta} = -\alpha \pmod{2\pi}$. In both cases $\alpha \in [0, \pi]$ and $\theta \equiv \pm\alpha$. For uniqueness, suppose $\alpha, \beta \in [0, \pi]$ both work; then $\alpha \equiv \pm\beta \pmod{2\pi}$. If $\alpha \equiv \beta$ then $\alpha - \beta \in 2\pi\mathbb{Z}$ and $|\alpha - \beta| \leq \pi < 2\pi$, forcing $\alpha = \beta$. If $\alpha \equiv -\beta$ then $\alpha + \beta \in 2\pi\mathbb{Z}$ with $0 \leq \alpha + \beta \leq 2\pi$, so $\alpha + \beta \in \{0, 2\pi\}$, giving either $\alpha = \beta = 0$ or $\alpha = \beta = \pi$. In all cases $\alpha = \beta$.

We claim

$$\lambda_{\min}(\theta) := \min \text{Spec}(H_\theta) = \frac{\alpha(\theta)^2}{L^2}. \quad (51)$$

To see this, note $\lambda_n(\theta) = \left(\frac{2\pi n - \theta}{L}\right)^2$, so

$$\min_{n \in \mathbb{Z}} \lambda_n(\theta) = \frac{1}{L^2} \min_{n \in \mathbb{Z}} (2\pi n - \theta)^2 = \frac{1}{L^2} \left(\text{dist}(\theta, 2\pi\mathbb{Z}) \right)^2,$$

the minimum being attained since $(2\pi n - \theta)^2 \rightarrow \infty$ as $|n| \rightarrow \infty$. It remains to identify $\text{dist}(\theta, 2\pi\mathbb{Z}) = \alpha(\theta)$. Write $\theta \equiv \pm\alpha$ with $\alpha = \alpha(\theta) \in [0, \pi]$; since $2\pi\mathbb{Z}$ is invariant under negation and under translation by 2π , we get $\text{dist}(\theta, 2\pi\mathbb{Z}) = \text{dist}(\pm\alpha, 2\pi\mathbb{Z}) = \text{dist}(\alpha, 2\pi\mathbb{Z})$. For $\alpha \in [0, \pi]$ the nearest point of $2\pi\mathbb{Z}$ is 0, since the next candidates are $\pm 2\pi$ at distance $2\pi - \alpha \geq \pi \geq \alpha$ and $2\pi + \alpha > \alpha$. Hence $\text{dist}(\alpha, 2\pi\mathbb{Z}) = \alpha$, proving (51).

Now assume $\text{Spec}(H_\theta) = \text{Spec}(H_{\theta'})$. Equal multisets have equal minima, so by (51)

$$\frac{\alpha(\theta)^2}{L^2} = \lambda_{\min}(\theta) = \lambda_{\min}(\theta') = \frac{\alpha(\theta')^2}{L^2},$$

whence $\alpha(\theta)^2 = \alpha(\theta')^2$. Since $\alpha(\theta), \alpha(\theta') \in [0, \pi]$ are both nonnegative, $\alpha(\theta) = \alpha(\theta') =: \alpha$. By (50), $\theta \equiv \pm\alpha$ and $\theta' \equiv \pm\alpha \pmod{2\pi}$. Combining the four sign cases: if $\theta \equiv \alpha$ and $\theta' \equiv \alpha$ then $\theta' \equiv \theta$; if $\theta \equiv \alpha$ and $\theta' \equiv -\alpha$ then $\theta' \equiv -\theta$; if $\theta \equiv -\alpha$ and $\theta' \equiv \alpha$ then $\theta' \equiv -\theta$; if $\theta \equiv -\alpha$ and $\theta' \equiv -\alpha$ then $\theta' \equiv \theta$. In every case $\theta' \equiv \pm\theta \pmod{2\pi}$, as required. $\square$

Corollary 11.3 (The spectral target sees holonomy only modulo inversion). *Define the equivalence $\sim_{\text{spec}}$ on $S_{\text{flat}} \cong \text{U}(1)$ by*

$$e^{i\theta} \sim_{\text{spec}} e^{i\theta'} \iff \theta' \equiv \pm\theta \pmod{2\pi} \iff e^{i\theta'} \in \{e^{i\theta}, e^{-i\theta}\}. \quad (52)$$

Then $\sim_{\text{spec}}$ is exactly the answer equivalence of the spectral target: $a_{\text{spec}}(\theta) = a_{\text{spec}}(\theta')$ iff $\theta \sim_{\text{spec}} \theta'$.

Proof. Immediate from Theorem 11.2 and Theorem 10.3. $\square$

11.3. The coarsest sufficient completion and its coordinate.

Theorem 11.4 ($\cos \theta$ coordinatizes the minimal spectral completion). *Let*

$$c: S_{\text{flat}} \longrightarrow [-1, 1], \quad c(\theta) := \cos \theta. \quad (53)$$

Then:

- (1) c is a well-defined surjection onto $[-1, 1]$;
- (2) $c(\theta) = c(\theta')$ if and only if $\theta \sim_{\text{spec}} \theta'$; that is, c realizes exactly the spectral equivalence classes;
- (3) c is a readout extension of p_F that is sufficient for Q_{spec} ;
- (4) c is a coarsest such extension: the induced map $\varphi_{[-1, 1]}: [-1, 1] \rightarrow E_{Q_{\text{spec}}}$ of [Theorem 8.3](#) is a bijection. Consequently every Q_{spec} -sufficient readout extension e must separate source states with distinct values of $\cos \theta$.

Proof. (1) $\cos$ is 2π -periodic, so c descends to $\mathbb{R}/2\pi\mathbb{Z}$; its image is $[-1, 1]$ since $\cos$ restricted to $[0, \pi]$ is already onto $[-1, 1]$.

(2) The function $\cos$ is strictly decreasing on $[0, \pi]$, hence injective there. Given θ, θ' , write $\alpha = \alpha(\theta), \alpha' = \alpha(\theta') \in [0, \pi]$ as in (50); since $\cos$ is even and 2π -periodic, $\cos \theta = \cos \alpha$ and $\cos \theta' = \cos \alpha'$. Therefore

$$c(\theta) = c(\theta') \iff \cos \alpha = \cos \alpha' \iff \alpha = \alpha' \iff \theta' \equiv \pm \theta \pmod{2\pi} \iff \theta \sim_{\text{spec}} \theta',$$

where the second equivalence uses injectivity of $\cos$ on $[0, \pi]$, and the third is the argument at the end of the proof of [Theorem 11.2](#) (equality of principal representatives is equivalent to $\theta' \equiv \pm \theta$).

(3) c is surjective by (1), and $p_F = u \circ c$ with $u: [-1, 1] \rightarrow \{0\}$ constant, so c is a readout extension. For sufficiency, define $b_{\text{spec}}: [-1, 1] \rightarrow A_{\text{spec}}$ by

$$b_{\text{spec}}(\kappa) := \text{Spec}(H_{\arccos \kappa}), \quad \arccos \kappa \in [0, \pi]. \quad (54)$$

If $c(\theta) = \kappa$ then by (2) $\theta \sim_{\text{spec}} \arccos \kappa$ (both have principal representative $\arccos \kappa$), so by [Theorem 11.3](#) $\text{Spec}(H_\theta) = \text{Spec}(H_{\arccos \kappa}) = b_{\text{spec}}(c(\theta))$. Hence $a_{\text{spec}} = b_{\text{spec}} \circ c$, and c is Q_{spec} -sufficient.

(4) Write $\varphi := \varphi_{[-1, 1]}$ and $E_{\text{spec}} := E_{Q_{\text{spec}}}$ for brevity. By [Theorem 8.3](#), applied to the sufficient extension c , there is a unique surjection $\varphi: [-1, 1] \rightarrow E_{\text{spec}}$ with $j_{Q_{\text{spec}}} = \varphi \circ c$, given by

$$\varphi(\kappa) = (u(\kappa), b_{\text{spec}}(\kappa)) = (0, b_{\text{spec}}(\kappa)), \quad \kappa \in [-1, 1]. \quad (55)$$

It remains to prove injectivity. Suppose $\varphi(\kappa) = \varphi(\kappa')$. Comparing second coordinates in (55) gives $b_{\text{spec}}(\kappa) = b_{\text{spec}}(\kappa')$, which by (54) says

$$\text{Spec}(H_{\arccos \kappa}) = \text{Spec}(H_{\arccos \kappa'}). \quad (56)$$

By [Theorem 11.2](#), (56) forces $\arccos \kappa' \equiv \pm \arccos \kappa \pmod{2\pi}$. Both $\arccos \kappa$ and $\arccos \kappa'$ lie in $[0, \pi]$, so each is its own principal representative in the sense of (50); by the uniqueness established there, they are equal. Applying $\cos$ yields $\kappa = \kappa'$.

Hence φ is a bijection, and c carries precisely the distinctions of E_{spec} — no more and no fewer.

The final assertion follows from [Theorem 8.4](#): if e is Q_{spec} -sufficient and $c(\theta_0) \neq c(\theta_1)$, then by (2) $\theta_0 \not\sim_{\text{spec}} \theta_1$, so by [Theorem 11.3](#) $a_{\text{spec}}(\theta_0) \neq a_{\text{spec}}(\theta_1)$, and [Theorem 7.1](#) forces $e(\theta_0) \neq e(\theta_1)$. $\square$

Remark 11.5 (What is and is not claimed about $\cos \theta$). [Theorem 11.4](#) says that $\cos \theta$ coordinatizes the coarsest completion sufficient for Q_{spec} : its level sets are exactly the spectral

equivalence classes, and it separates exactly what must be separated. It does not say that $\cos \theta$ is the unique possible numerical coordinate on $E_{Q_{\text{spec}}}$. Any injective function of $\cos \theta$ — for instance $\frac{1}{2}(1 - \cos \theta)$, or $\alpha(\theta)$ itself, or the ground-state energy $\lambda_{\min} = \alpha(\theta)^2/L^2$ — has the same level sets and coordinatizes the same quotient equally well; the completion is the quotient, and the coordinate is a labeling of it. Nor does [Theorem 11.4](#) say that $\cos \theta$ is minimal for *every* physical target on S_{flat} . It is minimal for the nominated spectral target, and for that target only.

Example 11.6 (A target for which $\cos \theta$ is insufficient: oriented transport). Consider the target Q_{Hol} with $A_{Q_{\text{Hol}}} = \text{U}(1)$ and correct-answer map $a_{\text{Hol}}(\theta) = e^{i\theta}$: “what phase does a unit charge acquire on being transported once around the circle in the positive direction?” Take $\theta = \pi/2$ and $\theta' = -\pi/2$. Then

$$c(\pi/2) = \cos(\pi/2) = 0 = \cos(-\pi/2) = c(-\pi/2),$$

so c does not separate these two states; but $a_{\text{Hol}}(\pi/2) = e^{i\pi/2} = i \neq -i = e^{-i\pi/2} = a_{\text{Hol}}(-\pi/2)$. By [Theorem 7.1](#), c is not Q_{Hol} -sufficient. By [Theorem 11.2](#) the two states are nonetheless isospectral, $\text{Spec}(H_{\pi/2}) = \text{Spec}(H_{-\pi/2})$, consistent with c being sufficient for Q_{spec} .

The full holonomy map h is Q_{Hol} -sufficient (take $b = \text{id}_{\text{U}(1)}$), and by [Theorem 8.4](#) it is coarsest for that target, since $j_{Q_{\text{Hol}}} = (p_F, a_{\text{Hol}})$ has the same level sets as h . So on one and the same reduced source space, with one and the same readout, the minimal completion for the transport target is the full holonomy circle, while the minimal completion for the spectral target is its strictly coarser inversion quotient $\text{U}(1)/(z \sim \bar{z}) \cong [-1, 1]$. The scalar spectrum is blind to orientation; oriented transport is not. This is target-relativity of completion in its sharpest form, and it is not a defect of either description: each retains exactly what its question requires.

12. FINITE-TIME DYNAMICAL DESCENT

Our final application of the criterion concerns the question that motivates most practical coarse-graining: can the projected description evolve on its own?

12.1. Setting. Let $\Phi_\tau: S \rightarrow S$ denote the source evolution through elapsed time τ , assumed to be a well-defined map for the fixed τ under consideration. No group law, continuity, or reversibility is needed for what follows; τ is fixed throughout this section and we prove a finite-time statement. Let $p: S \rightarrow T$ be the readout as before.

Definition 12.1 (Future-readout target). The *future-readout target* at time τ is the target Q_τ with answer set $A_{Q_\tau} = T$ and correct-answer map

$$a_\tau := p \circ \Phi_\tau: S \longrightarrow T, \quad a_\tau(s) = p(\Phi_\tau(s)). \quad (57)$$

In words: “given the source state now, what will the readout be after time τ ?”

Definition 12.2 (Autonomous readout propagator). An *autonomous readout propagator* at time τ is a map $U_\tau: T \rightarrow T$ such that the projection intertwines the dynamics:

$$p \circ \Phi_\tau = U_\tau \circ p. \quad (58)$$

Existence of U_τ is precisely the statement that the readout description is dynamically closed at time τ : the future readout value can be computed from the present readout value alone, without consulting the source.

12.2. The closure criterion.

Theorem 12.3 (Existence criterion for the autonomous propagator). *An autonomous readout propagator U_τ satisfying (58) exists if and only if*

$$p(s) = p(s') \implies p(\Phi_\tau(s)) = p(\Phi_\tau(s')) \quad \text{for all } s, s' \in S. \quad (59)$$

When it exists, U_τ is unique.

Proof. Equation (58) states exactly that $a_\tau = p \circ \Phi_\tau$ factors through p , i.e. that the target Q_τ of Theorem 12.1 is exactly solvable from the readout with $\bar{a}_{Q_\tau} = U_\tau$. By Theorem 3.1 applied to Q_τ (with $A_{Q_\tau} = T$ and $a_{Q_\tau} = a_\tau$), such a factorization exists if and only if a_τ is constant on every fiber of p , which is condition (3) of that theorem and reads, after substituting $a_\tau = p \circ \Phi_\tau$, exactly as (59). Uniqueness of U_τ is the uniqueness clause of Theorem 3.1. $\square$

Corollary 12.4 (Dynamical obstruction). *Suppose there exist $s_0, s_1 \in S$ with*

$$p(s_0) = p(s_1) \quad \text{but} \quad p(\Phi_\tau(s_0)) \neq p(\Phi_\tau(s_1)). \quad (60)$$

Then no map $U_\tau: T \rightarrow T$ satisfies (58). Concretely: two source states with the same present readout but different future readouts cannot both be reproduced by any deterministic evolution law acting on the present readout state alone.

Proof. Apply Theorem 3.2 to the target Q_τ : (60) is precisely (6) for a_τ . For any candidate U_τ , the value $U_\tau(t)$ at the common readout $t = p(s_0) = p(s_1)$ is a single element of T and hence cannot equal both $p(\Phi_\tau(s_0))$ and $p(\Phi_\tau(s_1))$. $\square$

12.3. Prediction at time τ versus autonomous evolution. The general theory of Sections 7 and 8 applies to Q_τ verbatim, and it is worth stating exactly what it delivers — and, just as importantly, what it does not.

By Theorem 7.1, any readout extension $e: S \rightarrow E$ sufficient for Q_τ must separate the obstructing pair s_0, s_1 of (60). By Theorems 8.2 and 8.3, the coarsest such extension is the joint target image

$$E_{Q_\tau} = \text{Im}(p, p \circ \Phi_\tau) \subseteq T \times T, \quad (61)$$

the present readout paired with the correct future readout at time τ .

Proposition 12.5 (What Q_τ -sufficiency provides). *Let $e: S \rightarrow E$ be a readout extension sufficient for Q_τ . Then there is a unique map*

$$b_\tau: E \longrightarrow T \quad \text{with} \quad p \circ \Phi_\tau = b_\tau \circ e, \quad (62)$$

so that the E -value at the present instant determines the readout value at time τ .

Proof. This is Theorem 6.3 and Theorem 6.5 applied to Q_τ , whose answer set is $A_{Q_\tau} = T$ and whose correct-answer map is $a_\tau = p \circ \Phi_\tau$. $\square$

Theorem 12.5 is a statement about *prediction of a T -valued quantity at the one specified time τ* . It says the enriched present state determines the future readout. It does *not* say that the enriched state evolves — that is, it produces a map $E \rightarrow T$, not a map $E \rightarrow E$, and there is no sense in which (62) can be iterated. Autonomous evolution on E is a strictly stronger requirement, and it is governed by its own descent condition.

Definition 12.6 (Autonomous evolution on an extension). *Let $e: S \rightarrow E$ be a readout extension. An autonomous evolution on E at time τ is a map $V_\tau: E \rightarrow E$ satisfying*

$$e \circ \Phi_\tau = V_\tau \circ e. \quad (63)$$

Proposition 12.7 (Criterion for autonomous evolution on E). *A map V_τ satisfying (63) exists if and only if*

$$e(s) = e(s') \implies e(\Phi_\tau(s)) = e(\Phi_\tau(s')) \quad \text{for all } s, s' \in S, \quad (64)$$

and it is then unique.

Proof. Apply Theorem 3.1 to the surjection e and the map $e \circ \Phi_\tau: S \rightarrow E$: a factorization $e \circ \Phi_\tau = V_\tau \circ e$ exists iff $e \circ \Phi_\tau$ is constant on the fibers of e , which is (64), and V_τ is then unique by the uniqueness clause of that theorem. $\square$

Conditions (62) and (63) are different requirements, and the implication between them runs in exactly one direction.

Proposition 12.8 (Autonomous evolution implies Q_τ -sufficiency). *Let $e: S \rightarrow E$ be a readout extension, with structure map $u: E \rightarrow T$ satisfying $p = u \circ e$. If e admits an autonomous evolution $V_\tau: E \rightarrow E$ at time τ , in the sense of (63), then e is sufficient for Q_τ , and one may take*

$$b_\tau = u \circ V_\tau. \quad (65)$$

Proof. Using $p = u \circ e$ and then (63),

$$p \circ \Phi_\tau = u \circ e \circ \Phi_\tau = u \circ V_\tau \circ e = (u \circ V_\tau) \circ e,$$

which is (62) with $b_\tau := u \circ V_\tau: E \rightarrow T$. Hence $a_\tau = p \circ \Phi_\tau$ factors through e , and e is Q_τ -sufficient by Theorem 6.3. $\square$

The converse fails in general. Sufficiency for Q_τ constrains e to determine the future *readout* — a T -valued quantity — whereas (64) constrains e to determine the future *E -value*, which is in general a finer quantity, since e refines p (Theorem 6.2 (3)). Knowing where the present E -state sends the time- τ readout says nothing about where it sends the time- τ E -state. The following finite example shows that the gap is real and not merely unproved.

Example 12.9 (A Q_τ -sufficient extension admitting no autonomous evolution). Let

$$S = \{a, b, c, d\}, \quad T = \{*\},$$

with $p: S \rightarrow T$ the constant readout, and let

$$E = \{x, y, z\}, \quad e(a) = e(b) = x, \quad e(c) = y, \quad e(d) = z.$$

Then e is surjective and $p = u \circ e$ for the constant map $u: E \rightarrow T$, so e is a readout extension. Define a source evolution $\Phi_\tau: S \rightarrow S$ by

$$\Phi_\tau(a) = c, \quad \Phi_\tau(b) = d,$$

with $\Phi_\tau(c)$ and $\Phi_\tau(d)$ chosen arbitrarily — say $\Phi_\tau(c) = c$ and $\Phi_\tau(d) = d$.

Since p is constant, the composite $p \circ \Phi_\tau$ is constant, hence constant on every fiber of e ; by Theorem 6.5 it factors through e , so e is Q_τ -sufficient. (Concretely, $b_\tau: E \rightarrow T$ is the constant map to $*$.)

Nevertheless e admits no autonomous evolution. We have $e(a) = e(b) = x$, so a and b lie in a common fiber of e , while

$$e(\Phi_\tau(a)) = e(c) = y \neq z = e(d) = e(\Phi_\tau(b)).$$

Thus $e \circ \Phi_\tau$ is *not* constant on the fiber $e^{-1}(x)$, so (64) fails, and by Theorem 12.7 no map $V_\tau: E \rightarrow E$ satisfies $e \circ \Phi_\tau = V_\tau \circ e$. Directly: such a V_τ would have to send the single element x to both y and z .

So Q_τ -sufficiency does not imply autonomous E -evolution, and autonomy is strictly stronger in general, with Q_τ -sufficiency the weaker consequence of it established in [Theorem 12.8](#).

Remark 12.10 (What E_{Q_τ} is, and what it is not). Accordingly, the joint target image (61) should be described in exactly one way: it is the universally coarsest completion of the readout *for predicting the T -valued readout at the fixed time τ* . By [Theorem 8.3](#) every Q_τ -sufficient extension maps onto it, and by [Theorem 8.2](#) it is itself sufficient for that prediction.

Its Q_τ -sufficiency does not by itself make it an autonomous state space, supply an iterable dynamics, or furnish a semigroup completion. Nothing proved here supplies a map $E_{Q_\tau} \rightarrow E_{Q_\tau}$ satisfying (63), and the results of this section do not assert one. By [Theorem 12.8](#), autonomy would be a strictly stronger property than the Q_τ -sufficiency that [Theorem 8.3](#) establishes for E_{Q_τ} , and [Theorem 12.9](#) shows that the two can come apart; the stronger property is not established here.

This leaves the question open rather than settled against it. In a particular system, E_{Q_τ} may or may not happen to admit an autonomous V_τ ; what does not hold is that any such evolution *follows* from the universal construction of [Theorem 8.3](#). Whether a given extension admits an autonomous V_τ , whether such extensions can be chosen compatibly across a range of times τ , and whether the resulting family composes as a semigroup are all questions beyond [Theorems 12.3](#) and [12.7](#), which are finite-time factorization statements at a single fixed τ . Those questions belong to [\[10\]](#) and are not addressed here.

12.4. What the obstruction does and does not license. [Theorem 12.4](#) rules out one thing precisely: a deterministic autonomous law on T alone. It is silent about what to do instead, and it is worth being explicit about the available modeling responses, none of which is an automatic consequence of the theorem.

- *State refinement.* One may replace T by a Q_τ -sufficient extension E , on which, by [Theorem 12.5](#), a deterministic *prediction* $b_\tau: E \rightarrow T$ of the time- τ readout does exist. Obtaining a deterministic *evolution* on E is a further demand, requiring (64), and is not delivered by Q_τ -sufficiency. Whether an extension meeting either condition is available in practice, and whether its variables are measurable, are separate questions the theorem does not settle. What the theorem does settle is a lower bound on what E must distinguish.
- *Memory.* One may allow the evolution law to depend on the past history of the readout rather than on its instantaneous value. This changes the shape of the modeling problem — the state is now a trajectory segment, not a point of T — and whether the history supplies the missing distinctions is not settled by [Theorem 12.3](#), which concerns the instantaneous readout. Projection-operator methods in nonequilibrium statistical mechanics [\[12, 5\]](#) are the classical setting in which such history-dependent reduced equations arise; the construction of exact memory equations for the present framework belongs to [\[10\]](#) and is not developed here.
- *Probabilistic modeling.* One may posit a probability law and predict a distribution over future readouts. See [Theorem 12.11](#).
- *Domain restriction.* One may restrict S to a subset on which (59) does hold, obtaining an exactly closed law on the restricted problem. This is legitimate and often useful, but it is a change of the source domain, and the theorem then applies to the restricted problem rather than the original one.

Remark 12.11 (Stochasticity is not exact trajectory recovery). It is sometimes suggested that the obstruction of [Theorem 12.4](#) is dissolved by making the reduced law stochastic. It is not, and the reason is the one already given in [Theorem 5.6](#). Suppose $p(s_0) = p(s_1) = t$

but the futures differ. A stochastic law on T assigns to t a distribution over T . If that law is intended not to exclude any source-compatible future — and a law that did exclude one would be predicting the impossibility of something that can in fact occur — then its support must contain both $p(\Phi_\tau(s_0))$ and $p(\Phi_\tau(s_1))$. It therefore does not tell us which of the two *actually occurs*, and if the system is in fact in state s_0 , the stochastic law has not predicted the realized trajectory; it has predicted a spread that contains it. This is a strictly weaker claim, and it answers a different target — “what is the distribution of the future readout, given the present readout *and a specified law on the fibers?*” — than the one we posed.

Note also the italicized clause. Even the distributional target requires input the readout does not contain: a probability law on the fibers F_t , which is a piece of source-level information. Absent such a law, the most that can be said from the readout alone is the set-valued statement of [Theorem 5.4](#), namely that the future readout lies in $A_{Q_\tau}(t) = p(\Phi_\tau(F_t))$, and that this set is the least sound prediction available. Stochastic modeling is a legitimate and powerful response to [Theorem 12.4](#); it is not a refutation of it, and it should not be presented as recovering the realized state.

13. DISCUSSION AND SCOPE

13.1. What has been established. The results above are elementary in their machinery and, we would argue, correspondingly robust. Their content can be compressed into four statements.

Obstruction is target-relative. A noninjective readout is not, as such, defective. It fails a target exactly when the correct answer to that target varies within one of its fibers ([Theorems 3.1](#) and [3.2](#)) — that is the general criterion, and it is stated without auxiliary hypotheses. When a mediating variable r is nominated, so that $a_Q = \Phi_Q(p, r)$ as in (7), the same criterion localizes: the loss of r obstructs Q exactly when r is active for Q on some fiber ([Theorems 4.3](#) and [4.4](#)). The same readout may pass one target and fail another ([Theorem 9.2](#) versus [Theorem 10.4](#)). Consequently the inference from “information was lost” to “the description must be completed” is invalid without a nominated target, and we have been careful never to make it.

Impossibility is derived, not assumed. Every impossibility claim in this paper reduces to the same counting observation: a deterministic construction from the readout has the form $b \circ p$ and therefore assigns one answer per fiber, while the correct-answer map may assign two. No hypothesis of the form “no replacement exists” is anywhere invoked, and [Theorem 3.4](#) closes the natural attempt to evade the conclusion by manufacturing derived readout variables.

Without completion, there is still a best answer, and it is a set. The compatible-answer set $A_Q(t) = a_Q(F_t)$ is the pointwise-least uniformly sound set-valued rule ([Theorem 5.4](#)), and exactness is equivalent to its being everywhere a singleton ([Theorem 5.5](#)). This is a set-theoretic fact prior to any probabilistic modeling; a probability law over $A_Q(t)$ ranks its elements but does not merge them and does not identify which source state obtained ([Theorems 5.6](#) and [12.11](#)).

With completion, there is a least sufficient one. Every Q -sufficient extension must separate states with different answers ([Theorem 7.1](#)), and the joint target image $E_Q = \text{Im}(p, a_Q)$ both is sufficient ([Theorem 8.2](#)) and receives a unique surjection from every sufficient extension ([Theorem 8.3](#)). It is thus the coarsest completion for the target, and it retains the readout value and the answer distinction and nothing else. In particular it does not reconstruct the source: [Theorem 8.4](#) runs $E \rightarrow E_Q$, and the identity extension id_S is one of the E ’s.

13.2. The example as a test of the theory. The flat $U(1)$ example is not decoration. It exhibits, in a system where every object is computable in closed form, all four statements at once. The reduced source space $S_{\text{flat}} \cong U(1)$ is a genuine moduli space, obtained only after quotienting by large gauge transformations ([30](#)), so its points are physically inequivalent.

The curvature readout (34) collapses it entirely. Against the flatness target the collapse is inert (Theorem 9.2); against the spectral target it is fatal (Theorem 10.4), and the fatality is a measurable one — the ground-state energy jumps from 0 to π^2/L^2 between $\theta = 0$ and $\theta = \pi$, with the eigenvalue multiplicities changing as recorded in Theorem 10.2. Yet the completion demanded by the spectral target is strictly less than the source: Theorem 11.2 shows the spectrum determines θ only up to sign, so Theorem 11.4 identifies the coarsest sufficient completion as the inversion quotient, coordinatized by $\cos \theta$. And this coarsest completion is genuinely insufficient for a different, entirely physical target on the same space (Theorem 11.6): oriented transport requires the whole holonomy circle, because it can tell $e^{i\pi/2}$ from $e^{-i\pi/2}$ and the scalar spectrum cannot.

13.3. Scope, and what is not claimed. We close by delimiting the results, since each of the following overreaches is easy and none is supported.

We have not shown that every lossy description requires repair; Theorem 4.4 shows the opposite is routinely the case. We have not shown that the source is recoverable from any completion; Theorem 8.4 explicitly denies it. We have not shown that curvature is a poor physical observable, only that it is a poor readout *for the spectral question*, which is a claim about a pair (Theorem 9.1). We have not shown that $\cos \theta$ is canonical among numerical coordinates or minimal for all targets (Theorems 11.5 and 11.6). We have not shown that stochastic reduced dynamics is illegitimate, only that it does not recover the realized trajectory and does not answer the target we posed (Theorem 12.11). And we have not made, and do not make, any claim to have established that physical reality is a projection of a hidden source; the mathematics here constrains what a given projection can and cannot deliver, given a nominated target, and that is all it does.

Within the sequence of which this paper is a part, the role of the present results is to supply the discrimination that the descent theorem of [8] alone does not provide. Paper 1 identified when a source relation descends; this paper identifies when the failure to descend matters, and what a repair must contain. Two natural generalizations are deliberately deferred. Passing from a single correct-answer map to a family of realization relations, and constructing the canonical minimal completion for such a family, requires machinery not developed here and is the subject of [9]; the results above concern a single target throughout, and their minimality claims should be read accordingly. Passing from the finite-time statement of Theorem 12.3 to continuous-time or semigroup dynamics, and to the exact memory equations that govern projected evolution when (59) fails, likewise lies outside the present scope [10]; Section 12 proves exactly one thing about dynamics — the existence criterion for the finite-time propagator U_τ — and no more.

14. CONCLUSION

The information a readout discards is not, by itself, an obstruction. It becomes one against a nominated target, and it does so under an exactly identifiable condition: the correct answer must vary across a single readout fiber. When that condition fails, the coarse description is complete for the question at hand, however much it has forgotten (Theorems 4.4 and 9.2). When it holds, no deterministic construction from the readout — and no auxiliary quantity computed from the readout — can be correct throughout (Theorems 3.2 and 3.4), and the honest residual is the compatible-answer set, which is the least sound set-valued rule available and is not a probability (Theorem 5.4).

Repair, when required, is then constrained from below, and the lower bound is attained abstractly within the paper’s category of extensions. Any sufficient enrichment must separate every pair of states whose correct answers differ (Theorem 7.1), and hence must distinguish at

least the joint (p, a_Q) -classes (Theorem 8.4). The joint target image $E_Q = \text{Im}(p, a_Q)$ attains this lower bound exactly as a set-theoretic construction: it is sufficient for Q (Theorem 8.2), it retains no information beyond the original readout value and the target answer, and every sufficient enrichment maps onto it (Theorem 8.3). This characterizes the least distinctions any repair must carry; it does not by itself construct a physically measurable repair, which must independently supply the enriched extension and the answer map that factors a_Q . The flat-connection example below does supply one concretely.

This is a statement about the floor, not the ceiling. Other Q -sufficient extensions may — and in practice usually do — carry additional information beyond the (p, a_Q) -classes: the identity extension id_S retains the entire source, and any extension adequate for several targets at once will generally be strictly finer than E_Q . Nothing forbids this, and Theorem 8.3 does not assert otherwise; it asserts only that whatever else such an extension carries, it maps canonically onto E_Q , which is the least it could have carried and still answered Q .

The flat $U(1)$ connections on a circle make all of this concrete and, we think, persuasive. A curvature readout that erases the entire flat moduli space is perfectly adequate for some questions and provably inadequate for the spectrum of the charged covariant Laplacian, where H_0 has a zero mode and H_π has none. The completion the spectrum demands is not the source and is not even the full holonomy: the spectrum determines the holonomy only up to inversion, and the coarsest sufficient completion is $U(1)/(z \sim \bar{z})$, coordinatized by $\cos \theta$ — strictly coarser than the holonomy that an oriented transport question would require. Two targets, one source, one readout, two different minimal completions. That is the shape of the phenomenon, and the same criterion, applied to the future readout, decides when a projected dynamics can close on itself.

REFERENCES

- [1] Y. Aharonov and D. Bohm, *Significance of electromagnetic potentials in the quantum theory*, Physical Review **115** (1959), 485–491.
- [2] J. Baez and J. P. Muniain, *Gauge Fields, Knots and Gravity*, World Scientific, Singapore, 1994.
- [3] S. Kobayashi and K. Nomizu, *Foundations of Differential Geometry, Volume I*, Interscience Publishers, New York, 1963.
- [4] M. Nakahara, *Geometry, Topology and Physics*, 2nd ed., Institute of Physics Publishing, Bristol, 2003.
- [5] H. Grabert, *Projection Operator Techniques in Nonequilibrium Statistical Mechanics*, Springer, Berlin, 1982.
- [6] M. Reed and B. Simon, *Methods of Modern Mathematical Physics I: Functional Analysis*, revised and enlarged ed., Academic Press, New York, 1980.
- [7] M. Reed and B. Simon, *Methods of Modern Mathematical Physics IV: Analysis of Operators*, Academic Press, New York, 1978.
- [8] J. Rodgers, *Source–Readout Non-Equivalence: Descent and Equivariant Reconstruction Obstructions*, Zenodo, 2026. doi:10.5281/zenodo.21369967.
- [9] J. Rodgers, *Canonical minimal completion for families of realization relations*, forthcoming manuscript, 2026.
- [10] J. Rodgers, *Projected dynamics, semigroups, and exact memory equations*, forthcoming manuscript, 2026.
- [11] T. T. Wu and C. N. Yang, *Concept of nonintegrable phase factors and global formulation of gauge fields*, Physical Review D **12** (1975), 3845–3857.
- [12] R. Zwanzig, *Ensemble method in the theory of irreversibility*, Journal of Chemical Physics **33** (1960), 1338–1341.

INDEPENDENT RESEARCHER