

NON-SOURCE PROJECTION AND INTERNAL IDENTIFIABILITY

JEREMY RODGERS

ABSTRACT. Let S be a source-realization space obtained after removal of declared gauge and coordinate redundancy, and let $p: S \rightarrow T_1$ be an exact surjective readout onto an observable tier. We establish equivalent criteria for source completeness of p ; we combine fiber variation with target dependence and canonical minimal completion; and we prove that any model containing an *essential non-gauge distinction* inside a p -fiber — two physically inequivalent source states with the same readout but different correct answers to a nominated target — is, relative to that model and target, a non-source projection.

The dynamical consequences divide sharply. If a deterministic source evolution preserves readout fibers, it descends uniquely to an autonomous law on T_1 even when p is noninjective; autonomy of the projected law therefore does not imply source completeness. If the evolution breaks fibers, no deterministic law on the instantaneous T_1 state reproduces every source trajectory. What this establishes is a negative: the instantaneous readout is insufficient. It does not by itself supply a replacement — a completion variable, an unresolved-state distribution, or a memory representation each requires an additional construction, and memory in particular need not encode the missing coordinate at all.

We then bound the inference available to an observer confined to the readout. A family of probes identifies the source exactly when its joint map separates source states; the maximally identifiable object is otherwise the quotient by the intersection of the probe equivalence relations. Under the explicit hypothesis that every admissible internal probe factors through p , no probe value, derived statistic, or deterministic selector distinguishes members of a common fiber. The statistical counterpart replaces probe values by complete outcome laws, and the finite-state stochastic counterpart is a strong lumpability theorem giving the exact condition under which a projected Markov source process remains Markov.

The synthesis is a boundary theorem. The mathematics determines when a readout is a non-source quotient, what a completion must restore, when projected laws close, and what internal observation can identify. It does not by itself supply a source domain, a physical equivalence, or a witness; those must be instantiated and tested by a physical model, and we mark that boundary explicitly throughout.

1. INTRODUCTION

1.1. The problem. A recurrent structure in physical modelling is the following. A theory posits a space of realizations, declares some of its differences to be pure redundancy (gauge, coordinate choice, relabelling), and then describes what is actually read off — measured, recorded, or lived — by a map into a smaller space of observable states. The observable description may be perfectly self-consistent: its states may evolve by autonomous laws, its probes may be mutually coherent, and its predictions may be exact. None of that settles the question of whether the observable space *is* the physical state space or merely a faithful quotient of a larger one.

Date: June 2026.

2020 Mathematics Subject Classification. Primary 37N20, 62B05; Secondary 60J27, 47A05, 81P16.

Key words and phrases. quotient map, exact readout, identifiability, lumpability, projected dynamics, minimal completion.

Website: everythingequation.com.

This paper isolates the exact mathematical content of that question. We work with a surjection

$$p: S \longrightarrow T_1,$$

where S is a source space *already reduced* by the physical equivalence the model declares, and T_1 is the readout. Because the reduction has been performed first, a nontrivial fiber $p^{-1}(t)$ is not a residue of gauge freedom: its distinct points are physically inequivalent by the model's own standard. The question is then whether such a fiber *matters*, and what an observer restricted to T_1 could ever learn about it.

1.2. Contribution. We prove a connected sequence of results, each stating a different location at which the source/readout distinction becomes visible or provably invisible.

- (1) **Source completeness** (§3). Under the hypothesis that a nominated family of invariant relations separates points of S , five conditions coincide: injectivity of p ; singleton fibers; descent of every relation through p ; bijectivity of the canonical completion projection $u_A: E_A \rightarrow T_1$; and existence of a two-sided inverse of p (Theorem 3.4). The point-separating hypothesis is what makes the equivalence a theorem rather than a definition, and we mark precisely where it and surjectivity are used.
- (2) **Non-source projection** (§4). If some fiber carries a distinction that changes the correct answer to a target Q , then p is noninjective, no deterministic function of the readout answers Q correctly on all of S , the minimal target completion E_Q properly refines T_1 , and every exact extension sufficient for Q maps uniquely and surjectively onto E_Q (Theorem 4.5, Theorem 4.6). We prove the universal property directly, so that the manuscript is self-contained. Relation, bounded-linear, and equivariant-selector forms follow.
- (3) **Projected-law dichotomy** (§5). Deterministic source evolution either preserves readout fibers, in which case it descends uniquely to an autonomous readout law, or it does not, in which case no deterministic law on the instantaneous readout state reproduces every source trajectory (Theorem 5.2). We emphasize the negative corollary: coherent autonomous projected laws establish dynamical preservation of readout classes, *not* that those classes are singletons.
- (4) **Identifiability, and two distinct ceilings** (§§6–8). A probe family identifies the source iff its joint map separates points; otherwise the maximally identifiable object is $S/\sim_I \cong \text{Im } P_I$, and adding probes refines monotonically. Two ceilings then follow from two *separate* hypotheses, and we are careful not to derive either from the other. The *deterministic* ceiling assumes that every admissible probe factors through p , i.e. $P_i = \bar{P}_i \circ p$; it rules out deterministic selectors and any post-processing of those probe values (Theorem 7.2). The *statistical* ceiling assumes instead that the complete outcome law factors through p , i.e. $\mu_s = \mathbf{M}_{p(s)}$; only this yields the stronger conclusion that no test, at any sample size, with any randomization, can discriminate (Corollary 8.4).
- (5) **The exact relations among the obstructions** (§8.1). Kinematic incompleteness, dynamical nonclosure, and statistical nonidentifiability are *not* logically independent. Fiber-breaking implies noninjectivity but not conversely; statistical nonidentifiability is unrelated to p unless an outcome-law factorization is imposed, under which noninjectivity implies it within a fiber (Proposition 8.6).
- (6) **Exact Markovian projection** (§9). For a finite-state continuous-time chain, the projected process is Markov for *every* initial source distribution iff the block-sum (strong lumpability) condition holds; we prove necessity from instantaneous block rates and sufficiency at the process level, not merely for one-time marginals (Theorem 9.3).

- (7) **The instantiation boundary (§10).** We state exactly what a physical model must supply, and exactly what the theorems do and do not license.

1.3. Relation to established work. Every ingredient used below is classical, and we claim none of them. The factorization of a map through the quotient by its kernel pair, and the universal property of images, are elementary category- and set-theoretic facts [11]. Descent of functions through a quotient, and invariance of a subspace under a semigroup, are standard [5]. The Hilbert-space statement $\mathcal{H}_{\text{src}}/\ker K \cong \text{Ran } K$ for bounded K with closed range is the closed-range theorem [14]. The projection formalism producing initial-state and memory terms for a non-invariant projector is due to Nakajima, Zwanzig and Mori [13, 15, 12]. The identification of observability with injectivity of the state-to-output map, which is the form our linear obstruction takes in Corollary 4.9, is Kalman’s [7]. Sufficiency, statistical identifiability, and the equivalence of equal laws with equal risk for all decision rules are Blackwell’s and Le Cam’s [2, 9, 10]. Strong lumpability of Markov chains is due to Kemeny and Snell [8], with the continuous-time form in Ball and Yeo [1] and Buchholz [3]. The obstruction to an equivariant section by empty fixed-point sets is standard equivariant topology [4].

What is new here is the *integration*: the assembly of these ingredients into a single source–readout architecture in which (i) the criterion for non-source projection is stated on a *physically reduced* domain, so that a nontrivial fiber cannot be dismissed as gauge; (ii) the target obstruction is integrated with a canonical minimal completion that is terminal among sufficient extensions; (iii) the *exact logical relations* among kinematic incompleteness, dynamical non-closure, and statistical nonidentifiability are established — these are *not* independent, and Proposition 8.6 states precisely how they are ordered; (iv) the internal-probe factorization hypothesis is isolated, as an explicit hypothesis, as the exact source of the identifiability ceiling; and (v) the boundary between a model-internal theorem and a claim about the world is drawn explicitly and defended.

We claim no novelty for any component criterion. Quotient descent, invariant-subspace conditions for projected closure, statistical identifiability, and strong lumpability are established mathematics, due to the authors cited above; the descent lemma and the lumpability criterion in particular are classical and are reproved here only so that the manuscript is self-contained.

Companion papers of the same programme [16, 17, 18, 19, 20] are cited for context only. *No proof below depends on any of them*: where the source material appealed to a companion result, we supply the proof here.

1.4. Scope. We state the limitation once, in full, and repeat it where it bites.

Every theorem below is conditional on a supplied quadruple — a source domain, a physical-equivalence relation, a readout, and (as appropriate) a target, a dynamics, or a family of outcome laws. The mathematics does not generate these objects. In particular, Theorem 4.5 does not prove that any actual physical system possesses a non-source projection; it proves that *a model exhibiting an essential non-gauge fiber distinction is one*, relative to that model and that target. Noninjectivity of an arbitrarily impoverished descriptor is not a witness; neither is a map merely named a projection. The conversion of the theorem into a claim about nature requires physical instantiation of its hypotheses and empirical discrimination of its alternatives, and we do not perform either. Section 10 is devoted to this boundary.

2. THE SOURCE-TO-READOUT ARCHITECTURE

2.1. Reduction and readout.

Definition 2.1 (Reduced source). Let Ω be a raw realization space carrying an equivalence relation $\simeq_{\text{phys}}$, the *declared physical equivalence* of the model: $\omega \simeq_{\text{phys}} \omega'$ means that the model

regards ω and ω' as descriptions of the same physical situation, differing only by gauge choice, coordinate choice, or relabelling. Define the *reduced source*

$$S := \Omega / \simeq_{\text{phys}}, \quad q: \Omega \twoheadrightarrow S \quad (1)$$

with q the canonical surjection.

The order of operations in (1) is the crux of the entire paper and we pause on it. If one did *not* quotient first, then a noninjective readout would be uninformative: its fibers could consist entirely of gauge copies of a single physical state, and their nontriviality would be an artifact of description. By constructing S first, we ensure that

$$s \neq s' \text{ in } S \implies s, s' \text{ are not related by any redundancy the model declares.} \quad (2)$$

Consequently a nontrivial p -fiber in S is a genuine multiplicity of physically inequivalent states *by the model's own standard*. This does not make it physically real — the model may be wrong — but it removes the cheapest objection.

Definition 2.2 (Exact readout). Let $\Pi_{T_1}: \Omega \rightarrow T_1$ be a surjection that is invariant under $\simeq_{\text{phys}}$, i.e. $\omega \simeq_{\text{phys}} \omega' \Rightarrow \Pi_{T_1}(\omega) = \Pi_{T_1}(\omega')$. We call T_1 the *readout* or *Tier 1*.

Lemma 2.3 (Induced readout). *There is a unique map $p: S \rightarrow T_1$ with*

$$\Pi_{T_1} = p \circ q, \quad (3)$$

and p is surjective. Moreover, writing $s \sim_p s' \iff p(s) = p(s')$, the map p induces a bijection

$$S / \sim_p \xrightarrow{\cong} T_1. \quad (4)$$

Proof. Existence and uniqueness of p are the universal property of the quotient. Define

$$p(q(\omega)) := \Pi_{T_1}(\omega).$$

This is well defined precisely by the $\simeq_{\text{phys}}$ -invariance of Π_{T_1} , and it is forced by (3), which gives uniqueness. Surjectivity of p follows from surjectivity of Π_{T_1} , since $\Pi_{T_1} = p \circ q$ and hence $\text{Im } p \supseteq \text{Im } \Pi_{T_1} = T_1$. Finally, $\sim_p$ is by construction the kernel pair of p , so the induced map $[s] \mapsto p(s)$ on $S / \sim_p$ is well defined and injective; it is surjective because p is. This is (4). $\square$

We record the terminology that the rest of the paper turns on.

Definition 2.4 (Exactness). We call p an *exact* readout to mean precisely (4): T_1 is, up to canonical bijection, the quotient of S by the relation “has the same readout”. That is, “exact” means exact *as the quotient by the readout's own kernel relation $\sim_p$* , and nothing more: it does not mean that p is injective, that the readout is source-complete, or that information is preserved at the source level. Exactness is a statement about *no information being lost within a fiber's label*; it is emphatically *not* a statement that fibers are singletons. The whole content of this paper lies in that gap.

2.2. Relations and the minimal completion.

Definition 2.5 (Invariant relations). Let $\mathcal{R} = \{r_\alpha: S \rightarrow V_\alpha\}_{\alpha \in A}$ be a family of maps, the *physically invariant source relations*. (Invariance is already built in: they are defined on S , hence constant on $\simeq_{\text{phys}}$ -classes.) Write

$$R_A := (r_\alpha)_{\alpha \in A}: S \longrightarrow \prod_{\alpha \in A} V_\alpha.$$

Definition 2.6 (Minimal completion). Define

$$E_A := \text{Im}(p, R_A) = \left\{ (p(s), R_A(s)) : s \in S \right\} \subseteq T_1 \times \prod_{\alpha} V_{\alpha}, \quad (5)$$

with the two canonical maps

$$e_A: S \rightarrow E_A, \quad e_A(s) = (p(s), R_A(s)), \quad u_A: E_A \rightarrow T_1, \quad u_A(t, v) = t. \quad (6)$$

The coordinate maps on the completion are $\rho_{\alpha}: E_A \rightarrow V_{\alpha}$, $\rho_{\alpha}(t, v) = v_{\alpha}$.

Lemma 2.7 (Elementary properties of the completion). *e_A is surjective; u_A is surjective; $p = u_A \circ e_A$; and $r_{\alpha} = \rho_{\alpha} \circ e_A$ for every α .*

Proof. e_A is surjective onto E_A by the definition (5) of E_A as an image. Then $u_A(e_A(s)) = p(s)$ gives $p = u_A \circ e_A$, and since p is surjective (Lemma 2.3) so is u_A . The identity $\rho_{\alpha}(e_A(s)) = r_{\alpha}(s)$ is immediate. $\square$

Thus E_A sits between S and T_1 : it is the readout *augmented by exactly the values of the nominated relations*, and nothing more. Its universal character is proved in §4 in the form we actually need (Theorem 4.6).

2.3. Specializations. The abstract triple $(S, p, \mathcal{R})$ has three standard realizations, which we record to fix the reader's picture. They are *specializations of one architecture*, not competing definitions.

Example 2.8 (Geometric orbit space and geometric realization). Let $\mathcal{C}_{\text{geo}}$ be a space of field configurations and $\simeq_{\text{geo}}$ the action of the gauge and diffeomorphism group. Then $S_{\text{geo}} = \mathcal{C}_{\text{geo}} / \simeq_{\text{geo}}$ is the orbit space, and the readout is a map $p_{\text{geo}}: S_{\text{geo}} \rightarrow T_1$ recording whatever subset of invariants it resolves. The relations are now geometric: write

$$J_A = (J_{\alpha})_{\alpha \in A}, \quad J_{\alpha}: S_{\text{geo}} \rightarrow V_{\alpha},$$

for a family of gauge- and diffeomorphism-invariant functionals: integrated, spectral, or relational curvature invariants; holonomy conjugacy classes; and Wilson-loop observables based on invariantly specified loops. Specializing Definition 2.6 gives the *geometric realization* of the completion,

$$E_A^{\text{geo}} := \text{Im}(p_{\text{geo}}, J_A) = \left\{ (p_{\text{geo}}(s), J_A(s)) : s \in S_{\text{geo}} \right\}, \quad (7)$$

the set of completed signatures actually *attained* by configurations.

The word “realization” carries the entire content of (7), and it is a restriction, not a licence. E_A^{geo} is by construction an *image*: it contains only those pairs (readout value, invariant values) that arise from an *admissible, globally compatible* field configuration in $\mathcal{C}_{\text{geo}}$. It is emphatically *not* the full product $T_1 \times \prod_{\alpha} V_{\alpha}$. An abstract signature — an arbitrary assignment of values to the invariants, consistent with no constraint one has bothered to impose — need not be realized by any global field at all: field equations, Bianchi identities, integrality and holonomy constraints, and global topological obstructions all cut the attainable set down. Determining *which* completed signatures are realized by compatible global geometric fields is thus a substantive geometric question about $\mathcal{C}_{\text{geo}}$, to be settled by geometry, and none of the theorems below settles it or presumes an answer.

We make no smoothness, Hausdorffness, or slice claims about S_{geo} : nothing below requires them, and orbit spaces are frequently badly behaved.

Example 2.9 (Bounded linear observation). Let $K: \mathcal{H}_{\text{src}} \rightarrow \mathcal{H}_{\text{obs}}$ be a bounded linear map with closed range. Taking $S = \mathcal{H}_{\text{src}}$ (with $\simeq_{\text{phys}}$ trivial) and p the corestriction of K to its range, the closed-range theorem gives

$$T_1 \cong \mathcal{H}_{\text{src}} / \ker K \cong \text{Ran } K. \quad (8)$$

Fibers of p are the cosets $\Psi + \ker K$. Here the “essential distinction” of §4 is a vector $n \in \ker K$, $n \neq 0$, that nevertheless changes some quantity of interest. Closedness of the range is used only for (8) to be an isomorphism of Banach spaces; the set-theoretic quotient statement (4) needs no such hypothesis.

Example 2.10 (Coarse-grained state). Let S be a finite set of microstates modulo relabelling and p a partition into macrostates. This is the setting of §9.

3. EQUIVALENT FORMS OF SOURCE COMPLETENESS

We now fix the vocabulary that the literature routinely conflates.

Definition 3.1 (Three completeness notions). Let $p: S \rightarrow T_1$ be an exact readout.

- (1) p is *kinematically source-complete* if p is injective (hence bijective, by Lemma 2.3).
- (2) p is *complete relative to a family $\mathcal{R}$* if every r_α *descends*, i.e. there is $\bar{r}_\alpha: T_1 \rightarrow V_\alpha$ with $r_\alpha = \bar{r}_\alpha \circ p$. Equivalently (Lemma 3.2), each r_α is constant on p -fibers. *Relative to a target Q* , completeness means a_Q descends.
- (3) p is *dynamically sufficient* (or *closed*) for a source evolution Φ_t if Φ_t preserves p -fibers, so that an autonomous law exists on T_1 (§5).

These are genuinely distinct, and it is worth saying why before proving anything, since the entire architecture would collapse into triviality if they were not. (1) implies (2) and (3) trivially. Neither converse holds: Example 5.5 exhibits a dynamically sufficient, hence perfectly autonomous, readout that is kinematically incomplete; Example 4.13 exhibits a readout complete relative to one relation family and incomplete relative to a larger one. *Relative* completeness is always completeness *with respect to a nominated family*, and this relativity is not a weakness of the framework but its content: it is what allows the question “does the lost distinction matter?” to have a nontrivial answer.

Lemma 3.2 (Descent lemma). *Let $p: S \rightarrow T_1$ be surjective and let $r: S \rightarrow V$ be any map. The following are equivalent:*

- (1) *there exists $\bar{r}: T_1 \rightarrow V$ with $r = \bar{r} \circ p$;*
- (2) *r is constant on every fiber of p : $p(s) = p(s') \Rightarrow r(s) = r(s')$;*
- (3) *$\sim_p \subseteq \sim_r$, where $s \sim_r s' \iff r(s) = r(s')$.*

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

Proof. (1) $\Rightarrow$ (2): if $p(s) = p(s')$ then $r(s) = \bar{r}(p(s)) = \bar{r}(p(s')) = r(s')$. (2) $\Leftrightarrow$ (3) is a restatement. (2) $\Rightarrow$ (1): given $t \in T_1$, surjectivity of p supplies some s with $p(s) = t$; set $\bar{r}(t) := r(s)$. Condition (2) says this does not depend on the choice of s , so $\bar{r}$ is well defined, and $r = \bar{r} \circ p$ by construction. Uniqueness: if $\bar{r} \circ p = \bar{r}' \circ p$ then $\bar{r} = \bar{r}'$ on $\text{Im } p = T_1$, using surjectivity again. $\square$

Note that surjectivity of p is used twice: once for existence of $\bar{r}$ on all of T_1 , once for uniqueness. This is the first of the two hypotheses whose location the following theorem is designed to expose.

Assumption 3.3 (Separating family). Throughout this section, $\mathcal{R} = \{r_\alpha\}_{\alpha \in A}$ separates the points of S :

$$s \neq s' \implies r_\alpha(s) \neq r_\alpha(s') \text{ for some } \alpha \in A. \quad (9)$$

Equivalently, R_A is injective.

Assumption 3.3 is what prevents the next theorem from being a tautology. Without it, “every r_α descends” is a statement about a possibly impoverished family and carries no information about p : the empty family descends vacuously through any readout whatsoever. With it, the family is rich enough to detect any source distinction, and descent of all of it becomes equivalent to there being no source distinction to detect. It is satisfiable: $\mathcal{R} = \{\text{id}_S\}$ separates trivially, and in the geometric setting (9) is precisely the (nontrivial, and sometimes false) assertion that the nominated invariants classify the orbit space.

Theorem 3.4 (Source-completeness equivalence). *Let $p: S \rightarrow T_1$ be an exact (hence surjective) readout and let $\mathcal{R}$ satisfy Assumption 3.3. Then the following are equivalent.*

- (C1) p is injective.
- (C2) Every fiber $p^{-1}(t)$, $t \in T_1$, is a singleton.
- (C3) Every $r_\alpha \in \mathcal{R}$ descends to T_1 .
- (C4) The completion projection $u_A: E_A \rightarrow T_1$ of (6) is bijective.
- (C5) p has a two-sided inverse: there is $p^{-1}: T_1 \rightarrow S$ with $p^{-1} \circ p = \text{id}_S$ and $p \circ p^{-1} = \text{id}_{T_1}$.

Proof. (C1) $\Leftrightarrow$ (C2). A map is injective iff each fiber has at most one point; each fiber of the surjection p has at least one point; so injectivity is exactly the singleton condition. (Surjectivity is used here, to rule out empty fibers.)

(C1) $\Leftrightarrow$ (C5). If p is injective and surjective it is a bijection, and its set-theoretic inverse is two-sided. Conversely, a map with a left inverse is injective: $p(s) = p(s') \Rightarrow s = p^{-1}p(s) = p^{-1}p(s') = s'$.

(C1) $\Rightarrow$ (C3). If p is injective then it is bijective, and $\bar{r}_\alpha := r_\alpha \circ p^{-1}$ satisfies $\bar{r}_\alpha \circ p = r_\alpha$. (Equivalently: fibers are singletons, so every function is trivially fiber-constant, and Lemma 3.2 applies.)

(C3) $\Rightarrow$ (C1). This is the only implication that uses Assumption 3.3. Suppose every r_α descends and let $p(s) = p(s')$. By Lemma 3.2, $r_\alpha(s) = r_\alpha(s')$ for every $\alpha \in A$. By (9), contrapositively, $s = s'$. Hence p is injective.

(C3) $\Rightarrow$ (C4). u_A is surjective by Lemma 2.7, so only injectivity is at issue. Let $(t, v), (t, w) \in E_A$ have the same image t under u_A . By definition of E_A there are $s, s' \in S$ with $p(s) = p(s') = t$, $R_A(s) = v$, $R_A(s') = w$. Since every r_α descends, it is fiber-constant (Lemma 3.2), and $p(s) = p(s')$ forces $r_\alpha(s) = r_\alpha(s')$ for all α , i.e. $v = w$. So u_A is injective.

(C4) $\Rightarrow$ (C3). Suppose u_A is bijective. Then $u_A^{-1}: T_1 \rightarrow E_A$ exists, and we set

$$\bar{r}_\alpha := \rho_\alpha \circ u_A^{-1}: T_1 \longrightarrow V_\alpha.$$

For $s \in S$, $u_A^{-1}(p(s)) = u_A^{-1}(u_A(e_A(s))) = e_A(s)$, whence $\bar{r}_\alpha(p(s)) = \rho_\alpha(e_A(s)) = r_\alpha(s)$ by Lemma 2.7. Thus $r_\alpha = \bar{r}_\alpha \circ p$ descends.

The five statements are therefore equivalent. $\square$

Remark 3.5 (Where the hypotheses enter). Surjectivity of p is used in (C1) $\Leftrightarrow$ (C2) (nonempty fibers), in (C1) $\Leftrightarrow$ (C5) (two-sidedness), and inside Lemma 3.2 (existence and uniqueness of $\bar{r}$). Assumption 3.3 is used *only* in (C3) $\Rightarrow$ (C1) — and it is indispensable there. Drop it and the implication fails at once: take $S = \{s_0, s_1\}$, $T_1 = \{t\}$, $\mathcal{R} = \emptyset$ (or any family of constant maps). Every relation descends vacuously, $E_A \cong T_1$ so u_A is bijective, yet p is not injective. So (C3) and (C4) are strictly weaker than (C1) in general, and the theorem’s content is that the separating hypothesis closes exactly this gap.

Corollary 3.6 (Nontrivial completion criterion). *Under Assumption 3.3, E_A properly refines T_1 — i.e. u_A is not injective — if and only if p is not kinematically source-complete.*

Proof. Immediate from $(C1) \Leftrightarrow (C4)$ of Theorem 3.4, since u_A is always surjective, so failure of bijectivity is failure of injectivity. $\square$

Theorem 3.4 thus supplies five interchangeable diagnostic sites for source incompleteness: noninjectivity, nontrivial fibers, failure of relation descent, nontrivial completion, and absence of an inverse. A physical model may find it convenient to test any one of them.

4. THE NON-SOURCE PROJECTION THEOREM

4.1. Targets, essential distinctions, and sufficiency.

Definition 4.1 (Target). A *target* Q is a question posed of the source, formalized by a *correct-answer map*

$$a_Q: S \longrightarrow A_Q.$$

$a_Q(s)$ is the true answer to Q when the source state is s . (Nothing is assumed about A_Q : it may be a two-point set, a space of numbers, a space of future trajectories, or a space of probability measures.)

Definition 4.2 (Essential distinction). A pair $s_0 \neq s_1$ in S is *essential for* Q if $a_Q(s_0) \neq a_Q(s_1)$. It is an *essential fiber distinction* if in addition $p(s_0) = p(s_1)$.

By (2), an essential fiber distinction consists of two states that the model declares physically inequivalent, that the readout cannot tell apart, and that differ in the answer to Q . That triple of conditions is the witness the whole paper is organized around.

Definition 4.3 (Exact extension sufficient for Q). An *exact extension* of the readout is a pair (E, π) where E is a set and $\pi: E \rightarrow T_1$ is a map, together with a surjection $e: S \twoheadrightarrow E$ satisfying

$$p = \pi \circ e. \tag{10}$$

(Thus E refines T_1 : it is a quotient of S that remembers at least as much as T_1 does. Surjectivity of e says E carries no states not realized by the source; (10) says E determines the readout.)

The extension is *sufficient for* Q if a_Q descends through e :

$$\exists \hat{a}: E \rightarrow A_Q \quad \text{with} \quad a_Q = \hat{a} \circ e. \tag{11}$$

Equivalently, by Lemma 3.2, $e(s) = e(s') \Rightarrow a_Q(s) = a_Q(s')$: the extension's state determines the correct answer.

Definition 4.4 (Minimal target completion). Specializing Definition 2.6 to the single relation a_Q , put

$$E_Q := \text{Im}(p, a_Q) = \{(p(s), a_Q(s)) : s \in S\}, \tag{12}$$

with $e_Q(s) = (p(s), a_Q(s))$, $u_Q(t, a) = t$, and $\rho_Q(t, a) = a$. By Lemma 2.7, e_Q is surjective, u_Q is surjective, $p = u_Q \circ e_Q$, and $a_Q = \rho_Q \circ e_Q$. In particular (E_Q, u_Q) is an exact extension sufficient for Q .

4.2. The theorem.

Theorem 4.5 (Non-source projection). *Let $p: S \rightarrow T_1$ be an exact readout on a reduced source (Definition 2.1), let Q be a target, and suppose there exist $s_0, s_1 \in S$ with*

$$p(s_0) = p(s_1) \quad \text{and} \quad a_Q(s_0) \neq a_Q(s_1). \quad (13)$$

Then:

- (N1) $s_0 \neq s_1$ and p is not injective; hence p is not bijective, and the readout does not identify each source state uniquely.
- (N2) Nevertheless T_1 remains the exact quotient: $S/\sim_p \cong T_1$.
- (N3) No deterministic function of the readout answers Q correctly on all of S : there is no $\bar{a}: T_1 \rightarrow A_Q$ with $a_Q = \bar{a} \circ p$.
- (N4) The minimal target completion E_Q properly refines T_1 : $u_Q: E_Q \rightarrow T_1$ is surjective but not injective.
- (N5) Every exact extension sufficient for Q admits a unique surjection onto E_Q commuting with the structure maps (Theorem 4.6).
- (N6) No change of gauge or of coordinates removes the distinction, since all such redundancies were quotiented in forming S .

Accordingly, relative to the model $(S, \simeq_{\text{phys}}, p)$ and the target Q , the readout T_1 is a non-source projection.

Proof. (N1) If $s_0 = s_1$ then $a_Q(s_0) = a_Q(s_1)$, contradicting (13); so $s_0 \neq s_1$. These two distinct points have the same image under p , so p is not injective. Since p is surjective (Lemma 2.3) but not injective, p is not bijective; hence the readout does not identify each source state uniquely — the states s_0 and s_1 share the readout value $p(s_0) = p(s_1)$. We emphasize that this argument is *direct* and invokes no separating hypothesis: Theorem 3.4 is unavailable here, since Assumption 3.3 is not among the hypotheses of the present theorem, and it is not needed — the witness (13) exhibits the noninjectivity outright.

(N2) This is Lemma 2.3, which used only $\simeq_{\text{phys}}$ -invariance and surjectivity of Π_{T_1} , and is untouched by (13). The point deserves emphasis: exactness of the readout and failure of source completeness are perfectly compatible. Losing the fiber's internal structure is not a defect of exactness; it is what exactness of a noninjective quotient *means*.

(N3) Suppose such $\bar{a}$ existed. By Lemma 3.2((1) $\Rightarrow$ (2)), a_Q would be constant on p -fibers. But $p(s_0) = p(s_1)$ places s_0, s_1 in one fiber while $a_Q(s_0) \neq a_Q(s_1)$; contradiction.

(N4) u_Q is surjective by Lemma 2.7. Set $t := p(s_0) = p(s_1)$ and consider

$$e_Q(s_0) = (t, a_Q(s_0)), \quad e_Q(s_1) = (t, a_Q(s_1)).$$

Both lie in E_Q ; they are distinct, because their second coordinates differ by (13); and u_Q sends both to t . Hence u_Q is not injective, i.e. some fiber of u_Q has at least two points, which is what “ E_Q properly refines T_1 ” asserts.

(N5) Proved as Theorem 4.6 below.

(N6) By construction (1), $S = \Omega/\simeq_{\text{phys}}$, so any two raw realizations related by a declared redundancy have already been identified. Since $s_0 \neq s_1$ in S , no such redundancy relates their representatives; a fortiori no relabelling in the declared class maps one to the other. Cf. (2). $\square$

Theorem 4.6 (Universal property of the minimal target completion). *Let (E, π, e) be any exact extension sufficient for Q (Definition 4.3). Then there is a unique map*

$$\phi: E \longrightarrow E_Q \quad \text{with} \quad \phi \circ e = e_Q,$$

and ϕ is surjective and satisfies $u_Q \circ \phi = \pi$ and $\rho_Q \circ \phi = \hat{a}$.

Proof. Existence. Since e is surjective, by Lemma 3.2 it suffices to check that e_Q is constant on the fibers of e . Let $e(s) = e(s')$. Then:

- $p(s) = \pi(e(s)) = \pi(e(s')) = p(s')$, using (10);
- $a_Q(s) = \hat{a}(e(s)) = \hat{a}(e(s')) = a_Q(s')$, using (11).

Hence $e_Q(s) = (p(s), a_Q(s)) = (p(s'), a_Q(s')) = e_Q(s')$. Lemma 3.2 now supplies a unique $\phi: E \rightarrow E_Q$ with $\phi \circ e = e_Q$, which is also the asserted uniqueness.

Surjectivity. e_Q is surjective onto E_Q (Lemma 2.7) and factors as $\phi \circ e$; hence $\text{Im } \phi \supseteq \text{Im}(\phi \circ e) = E_Q$.

Compatibility. For $s \in S$: $u_Q(\phi(e(s))) = u_Q(e_Q(s)) = p(s) = \pi(e(s))$, and since e is surjective, $u_Q \circ \phi = \pi$ on all of E . Likewise $\rho_Q(\phi(e(s))) = \rho_Q(e_Q(s)) = a_Q(s) = \hat{a}(e(s))$, giving $\rho_Q \circ \phi = \hat{a}$. $\square$

Remark 4.7 (Direction of the universal map, and what it means). The map runs *from* an arbitrary sufficient extension *onto* E_Q , not the reverse. The reading is: E_Q is the *coarsest* sufficient extension — it retains exactly the readout plus the answer to Q and nothing else — and any other sufficient extension, being at least as fine, projects onto it. In categorical language E_Q is terminal in the category of exact extensions sufficient for Q (with morphisms the maps commuting with the e 's), and terminality gives both the uniqueness of ϕ and, in the usual way, uniqueness of E_Q up to canonical isomorphism. It is exactly the wrong intuition to expect a map $E_Q \rightarrow E$: a completion that remembers *more* than Q requires cannot be manufactured out of Q 's answer.

4.3. Corollaries: relation, linear, and equivariant forms.

Corollary 4.8 (Relation form). *Let $r: S \rightarrow V$ be an invariant relation and suppose Q is active in r , meaning a_Q factors as $a_Q = \theta \circ r$ for some $\theta: V \rightarrow A_Q$. If r varies on some p -fiber and that variation changes θ , i.e. there are s_0, s_1 with $p(s_0) = p(s_1)$ and $\theta(r(s_0)) \neq \theta(r(s_1))$, then all conclusions of Theorem 4.5 hold, and $E_{\{r\}}$ properly refines T_1 .*

Proof. The hypothesis says exactly $a_Q(s_0) = \theta(r(s_0)) \neq \theta(r(s_1)) = a_Q(s_1)$ with $p(s_0) = p(s_1)$, i.e. (13); apply Theorem 4.5. For the last claim, $r(s_0) \neq r(s_1)$ (else θ would take equal values), so $e_{\{r\}}(s_0) \neq e_{\{r\}}(s_1)$ while $u_{\{r\}}$ identifies them, exactly as in the proof of (N4). $\square$

The qualification in Corollary 4.8 is not decorative. Mere variation of r on a fiber is *not* enough for a target obstruction: if θ is constant on the values r takes in that fiber, the target does not see the variation. Obstruction is always relative to a target.

Corollary 4.9 (Bounded-linear observation form). *Let $K: \mathcal{H}_{\text{src}} \rightarrow \mathcal{H}_{\text{obs}}$ be bounded, let $S = \mathcal{H}_{\text{src}}$, and let $p = K$ (corestricted to $\text{Ran } K$), so that p -fibers are the cosets $\Psi + \ker K$. Let $a_Q: \mathcal{H}_{\text{src}} \rightarrow A_Q$ be a target. If there exist $\Psi \in \mathcal{H}_{\text{src}}$ and $n \in \ker K$ with*

$$a_Q(\Psi + n) \neq a_Q(\Psi), \quad (14)$$

then $\text{Ran } K$ is a non-source projection for Q in the sense of Theorem 4.5: no function of $K\Psi$ alone computes a_Q on all of $\mathcal{H}_{\text{src}}$, and every exact extension sufficient for Q surjects onto E_Q and therefore distinguishes the cosets $\Psi + n$ and Ψ .

Proof. $K(\Psi + n) = K\Psi + Kn = K\Psi$, so $p(\Psi + n) = p(\Psi)$; with (14) this is the witness (13) with $s_0 = \Psi$, $s_1 = \Psi + n$. (Note $n \neq 0$ follows from (14).) Apply Theorem 4.5 and Theorem 4.6. That every sufficient extension separates the two cosets: if $e(\Psi) = e(\Psi + n)$, sufficiency (11) would force $a_Q(\Psi) = a_Q(\Psi + n)$, contradiction. $\square$

We note what Corollary 4.9 does *not* assume: no closed range (that was needed only for the Banach-space isomorphism (8), not for the set-level obstruction), no self-adjointness, no linearity of a_Q . If a_Q happens to be a bounded linear functional $\langle \xi, \cdot \rangle$, then (14) holds for some $n \in \ker K$ iff $\xi \notin (\ker K)^\perp = \overline{\text{Ran } K^*}$ — a clean spectral criterion, and a good way to test the hypothesis in a concrete linear model.

For the equivariant obstruction we first fix notation.

Definition 4.10 (Group action, stabilizer, fixed set). Let a group G act on S and on T_1 , and suppose p is *equivariant*: $p(g \cdot s) = g \cdot p(s)$ for all $g \in G$, $s \in S$. For $t \in T_1$ let

$$G_t := \{g \in G : g \cdot t = t\}$$

be its *stabilizer*. Equivariance implies G_t acts on the fiber $p^{-1}(t)$: if $p(s) = t$ and $g \in G_t$ then $p(g \cdot s) = g \cdot t = t$. Write the *fixed-point set of that action*

$$\text{Fix}_{G_t}(p^{-1}(t)) := \{s \in p^{-1}(t) : g \cdot s = s \ \forall g \in G_t\}.$$

A G -equivariant source selector is a map $\sigma : T_1 \rightarrow S$ with $p \circ \sigma = \text{id}_{T_1}$ and $\sigma(g \cdot t) = g \cdot \sigma(t)$ for all g, t .

Corollary 4.11 (Symmetry-compatible reconstruction obstruction). *In the setting of Definition 4.10, if there exists $t \in T_1$ with*

$$\text{Fix}_{G_t}(p^{-1}(t)) = \emptyset, \tag{15}$$

then no G -equivariant source selector σ exists. In particular, no reconstruction rule that respects the readout symmetry can choose a source representative in every fiber.

Proof. Suppose σ were such a selector and let t satisfy (15). Since $p(\sigma(t)) = t$, we have $\sigma(t) \in p^{-1}(t)$. Let $g \in G_t$. Then, using equivariance of σ and $g \cdot t = t$,

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

As $g \in G_t$ was arbitrary, $\sigma(t)$ is fixed by every element of G_t , i.e. $\sigma(t) \in \text{Fix}_{G_t}(p^{-1}(t))$, contradicting (15). $\square$

Remark 4.12. Corollary 4.11 obstructs *equivariant* selectors only. A non-equivariant selector may well exist — indeed, for p surjective, the axiom of choice always supplies *some* set-theoretic section σ with $p \circ \sigma = \text{id}$. The content is that any such section must *break* the symmetry G : reconstruction is possible only at the cost of a choice that the readout’s own symmetry group does not sanction. This is the precise sense in which the obstruction is a symmetry obstruction and not a cardinality one.

4.4. The scope of Theorem 4.5. Theorem 4.5 is conditional on four supplied objects: the source domain S , the physical equivalence $\simeq_{\text{phys}}$ used to build it, the readout p , and the essential witness (13). It derives none of them. Its strength is precisely delimited: *once* those four are established for a model, the non-source conclusion follows with no further interpretive step — one does not additionally need to argue that the projection “loses something real”, because the witness (13) says so in the model’s own terms. Its weakness is equally precise: it says nothing whatever about whether our universe supplies them. We return to this in §10, and we do not claim more anywhere.

Example 4.13 (Finite source, nontrivial fiber, proper completion). Let

$$S = \{s_0, s_1, s_2\}, \quad T_1 = \{t, t'\}, \quad p(s_0) = p(s_1) = t, \quad p(s_2) = t'.$$

Assume the model has already reduced by $\simeq_{\text{phys}}$, so s_0, s_1, s_2 are pairwise physically inequivalent. Take the target Q with $A_Q = \{0, 1\}$ and

$$a_Q(s_0) = 0, \quad a_Q(s_1) = 1, \quad a_Q(s_2) = 1.$$

Then (13) holds with the pair (s_0, s_1) . Compute the minimal target completion:

$$E_Q = \text{Im}(p, a_Q) = \{(t, 0), (t, 1), (t', 1)\},$$

a three-element set. The projection $u_Q: E_Q \rightarrow T_1$ sends $(t, 0) \mapsto t$, $(t, 1) \mapsto t$, $(t', 1) \mapsto t'$; it is surjective, and its fiber over t has two elements, so it is not injective — E_Q properly refines T_1 , as (N4) requires. Note also that $e_Q: S \rightarrow E_Q$ is here a *bijection*, so in this example E_Q happens to recover S completely; that is an accident of a_Q separating s_0 from s_1 , not a general feature. Indeed, modify the target to $a_{Q'}(s_0) = a_{Q'}(s_1) = 0$, $a_{Q'}(s_2) = 1$. Now $a_{Q'}$ is fiber-constant, hence descends ($\bar{a}_{Q'}(t) = 0$, $\bar{a}_{Q'}(t') = 1$), $E_{Q'} = \{(t, 0), (t', 1)\} \cong T_1$, and T_1 is complete relative to Q' while remaining kinematically incomplete. *Both targets, one readout:* obstruction is target-relative, and Corollary 3.6 bites only for families that separate — $\{a_{Q'}\}$ does not separate s_0 from s_1 .

5. THE PROJECTED-LAW DICHOTOMY

5.1. Descent of dynamics.

Definition 5.1 (Source evolution and autonomy). Let $\Phi_t: S \rightarrow S$, $t \in \mathbb{T}$ (with $\mathbb{T} = \mathbb{R}_{\geq 0}$ or $\mathbb{N}$), be a deterministic source evolution, i.e. $\Phi_0 = \text{id}_S$ and $\Phi_{t+t'} = \Phi_t \circ \Phi_{t'}$. The readout dynamics are *autonomous* if there exist maps $U_t: T_1 \rightarrow T_1$ with

$$p \circ \Phi_t = U_t \circ p \quad \text{for all } t \in \mathbb{T}. \quad (16)$$

(16) is the exact formalization of the statement “the present readout determines the future readout”. We stress that it is a *deterministic* and *instantaneous* demand: U_t takes as input the readout state at one time, and nothing else — no history, no auxiliary variable, no distribution.

Theorem 5.2 (Projected-law dichotomy). *Exactly one of the following holds.*

Case I (fiber-preserving). *For all $t \in \mathbb{T}$ and all $s, s' \in S$,*

$$p(s) = p(s') \implies p(\Phi_t(s)) = p(\Phi_t(s')). \quad (17)$$

Then there exist unique maps $U_t: T_1 \rightarrow T_1$ satisfying (16), and they form a semigroup: $U_0 = \text{id}_{T_1}$, $U_{t+t'} = U_t \circ U_{t'}$. The readout law is autonomous.

Case II (fiber-breaking). *There exist $s_0, s_1 \in S$ and $t > 0$ with*

$$p(s_0) = p(s_1), \quad p(\Phi_t(s_0)) \neq p(\Phi_t(s_1)). \quad (18)$$

Then no family of maps U_t satisfies (16); no deterministic law on the instantaneous readout state reproduces every source trajectory.

Proof. The two cases are the negation of one another, hence exhaustive and mutually exclusive.

Case I. Condition (17) says precisely that, for each fixed t , the map $p \circ \Phi_t: S \rightarrow T_1$ is constant on the fibers of p . Since p is surjective, Lemma 3.2 supplies a unique $U_t: T_1 \rightarrow T_1$ with $p \circ \Phi_t = U_t \circ p$, namely $U_t(p(s)) := p(\Phi_t(s))$, well defined exactly by (17). For the semigroup law, evaluate on an arbitrary $t_1 = p(s)$ (surjectivity again):

$$U_{t+t'}(p(s)) = p(\Phi_{t+t'}(s)) = p(\Phi_t(\Phi_{t'}(s))) = U_t(p(\Phi_{t'}(s))) = U_t(U_{t'}(p(s))),$$

and $U_0(p(s)) = p(\Phi_0(s)) = p(s)$.

Case II. Suppose some U_t satisfied (16), and let s_0, s_1, t be as in (18). Then

$$p(\Phi_t(s_0)) = U_t(p(s_0)) = U_t(p(s_1)) = p(\Phi_t(s_1)),$$

where the middle equality uses $p(s_0) = p(s_1)$ and the fact that U_t is a *function* of the readout state. This contradicts (18). Hence no such U_t exists. $\square$

Corollary 5.3 (Coherence does not imply source completeness). *Case I is compatible with p noninjective. Autonomous, deterministic, internally coherent readout laws establish that the dynamics preserves the readout equivalence classes; they do not establish that those classes are singletons.*

Proof. It suffices to exhibit an instance, which Example 5.5 does. Formally: (17) constrains Φ_t to map $\sim_p$ -classes into $\sim_p$ -classes; it places no constraint whatsoever on the cardinality of those classes. $\square$

Corollary 5.3 is the single most important negative result of this paper, and the one most easily lost. An observer inside T_1 in Case I sees a world with exact deterministic laws, perfectly closed, admitting no anomaly and no residue. That observer’s evidence is entirely consistent with T_1 being the whole of physical reality, and entirely consistent with T_1 being a proper quotient of a strictly larger source. *Dynamical closure is evidence of fiber preservation, not of fiber triviality.*

5.2. What failed descent actually requires. In Case II the instantaneous readout state is insufficient, and one must say for *what*. The three targets below are routinely conflated; they are not the same, and only the first is an exact reconstruction.

- (1) *Exact trajectory prediction.* To recover $t \mapsto p(\Phi_t(s))$ for the individual realized s , one needs strictly more than $p(s)$: either additional state information (a completion variable, i.e. a refinement $e: S \rightarrow E$ sufficient for the target “future readout”, in the sense of Definition 4.3), or an equivalent history-dependent description that encodes the same information in the past of the trajectory. Which of the two one uses is a matter of representation; the information content required is fixed by the fiber structure.
- (2) *Statistical prediction.* One may instead supply a probability distribution over the unresolved states within the fiber $p^{-1}(p(s))$. This yields a distribution over future readouts. But note carefully: this is a prediction *only after an additional probabilistic model is supplied* — namely the within-fiber distribution, which is *not* determined by the deterministic data (S, p, Φ_t) . It is an addition to the model, not a consequence of it. And what it delivers is a law over futures, *not* an exact reconstruction of the individual source trajectory.
- (3) *Memory representation.* One may retain a functional of the past readout. This is exact only when the past readout in fact determines the relevant fiber coordinate; in general it does not, and then memory too is only a partial substitute.

We do not assert that any of these is always available. Case II asserts a negative — that the instantaneous readout state does not suffice — and nothing more.

5.3. The Hilbert-space instance. Specialize to Example 2.9: $S = \mathcal{H}_{\text{src}}$, $p = K$ bounded with $\ker K$ closed, and let the source evolution be a strongly continuous group $\Phi_t = U(t)$ with generator L . The dichotomy takes an operator form.

Proposition 5.4 (Invariance criterion). *The evolution $U(t)$ is fiber-preserving in the sense of (17) if and only if*

$$U(t) \ker K \subseteq \ker K \quad \text{for all } t. \quad (19)$$

Proof. Fibers of K are the cosets $\Psi + \ker K$. Condition (17) reads: $K\Psi = K\Psi' \Rightarrow KU(t)\Psi = KU(t)\Psi'$. Writing $n = \Psi' - \Psi$, the hypothesis is $n \in \ker K$ and the conclusion is $KU(t)n = 0$,

i.e. $U(t)n \in \ker K$. Quantifying over all Ψ and all $n \in \ker K$ gives exactly (19). (Linearity of $U(t)$ and K is used in both directions.) $\square$

Now let P be the orthogonal projection of $\mathcal{H}_{\text{src}}$ onto the closed subspace $(\ker K)^\perp$, the “resolved” subspace, and $Q^\perp := I - P$ the projection onto $\ker K$, the “unresolved” subspace. (We write $Q^\perp$ rather than Q to avoid collision with the target Q and with the generator of §9; in the Mori–Zwanzig literature this operator is universally called Q .) For a bounded generator L , invariance (19) is equivalent to the infinitesimal condition

$$P L Q^\perp = 0, \quad (20)$$

i.e. the generator does not map unresolved directions into resolved ones. Failure of (20) is Case II. Irrespective of whether (20) holds, the standard Nakajima–Zwanzig–Mori projection [13, 15, 12] yields an exact identity for the resolved component, which for bounded L we derive here by variation of constants, so that the manuscript remains self-contained.

Write $\Psi(t) = U(t)\Psi(0)$, so that $\dot{\Psi} = L\Psi$, and split it by the two projections:

$$p(t) := P\Psi(t), \quad q(t) := Q^\perp\Psi(t), \quad \Psi(t) = p(t) + q(t).$$

(The letter p here denotes the resolved component of the state, following the standard notation of the projection literature; it is not the readout map of §2, and no confusion should arise, as the readout does not appear in this derivation.) Since P and $Q^\perp$ are constant projections with $P + Q^\perp = I$, applying each to $\dot{\Psi} = L\Psi = L(p + q)$ gives the coupled pair

$$\dot{p} = PLP p + PLQ^\perp q, \quad (21)$$

$$\dot{q} = Q^\perp L P p + Q^\perp L Q^\perp q. \quad (22)$$

Now read (22) as an inhomogeneous linear equation for q , with generator $Q^\perp L Q^\perp$ and inhomogeneity $Q^\perp L P p(\tau)$ regarded as a given source term. Because L is bounded, so is $Q^\perp L Q^\perp$, and $e^{tQ^\perp L Q^\perp}$ is a norm-continuous semigroup; variation of constants (equivalently, differentiating $e^{-tQ^\perp L Q^\perp} q(t)$ and integrating from 0 to t) gives the unique solution

$$q(t) = e^{tQ^\perp L Q^\perp} q(0) + \int_0^t e^{(t-\tau)Q^\perp L Q^\perp} Q^\perp L P p(\tau) d\tau. \quad (23)$$

Substituting (23) into the term $PLQ^\perp q$ of (21), and using $q(0) = Q^\perp\Psi(0)$ together with $p(\tau) = P\Psi(\tau)$, yields the exact identity

$$\frac{d}{dt} P\Psi(t) = \underbrace{PLP \Psi(t)}_{\text{closed term}} + \underbrace{PLQ^\perp e^{tQ^\perp L Q^\perp} Q^\perp \Psi(0)}_{\text{unresolved-initial-state term}} + \underbrace{\int_0^t PLQ^\perp e^{(t-\tau)Q^\perp L Q^\perp} Q^\perp L P P\Psi(\tau) d\tau}_{\text{memory term}}. \quad (24)$$

Four points must be made about (24), because it is the most frequently over-read equation in this circle of ideas.

First: $P\Psi$ is a source-side coordinate, not the declared readout. The projection formalism evolves $P\Psi$, the component of the source state in $(\ker K)^\perp$. The *declared readout* of our architecture, by contrast, is $K\Psi$. These are different objects living in different spaces, and the identification of the two requires an argument. Under the closed-range hypothesis it is available: the restriction

$$K|_{(\ker K)^\perp} : (\ker K)^\perp \longrightarrow \text{Ran } K$$

is a bounded linear bijection between Banach spaces (injective because $(\ker K)^\perp \cap \ker K = \{0\}$; surjective onto $\text{Ran } K$ by the orthogonal decomposition $\Psi = P\Psi + Q^\perp\Psi$ and $K\Psi = KP\Psi$), and

$\text{Ran } K$ is closed, hence complete; so the bounded inverse theorem supplies a *bounded* inverse. Thus $P\Psi \mapsto K\Psi$ is a topological isomorphism, and the two observable-state descriptions are equivalent: convergence, boundedness, and continuity statements transfer between them. *Without* closed range this stability fails — the inverse need not be bounded — and although the set-theoretic quotient (4) still holds, one may no longer freely trade $P\Psi$ for $K\Psi$ in any topological or approximation-theoretic claim.

Second: (24) is exact and deterministic. No stochasticity, no coarse-graining approximation, and no thermodynamic limit has been invoked. The second and third terms are not noise; they are definite functionals of the initial unresolved data $Q^\perp\Psi(0)$ and of the resolved history $\{P\Psi(\tau)\}_{\tau \leq t}$ respectively.

Third: the exact relation between $PLQ^\perp = 0$ and the vanishing of the two anomalous terms. This must be stated with care, and the naive reading — that the two terms vanish or fail to vanish together — is false. The correct statement is:

- $PLQ^\perp = 0$ is *equivalent* to autonomous closure for every source state (Proposition 5.4 together with (20));
- when $PLQ^\perp = 0$, *both* the unresolved-initial-state term and the memory kernel vanish identically, since each carries $PLQ^\perp$ as a left factor — recovering $\frac{d}{dt}P\Psi = PLP\Psi$. This is Theorem 5.2 in operator dress;
- conversely, if $PLQ^\perp \neq 0$, then the *initial-state* response is nonzero for *some unresolved initial state* — choose $Q^\perp\Psi(0)$ with $PLQ^\perp\Psi(0) \neq 0$ and evaluate at $t = 0$;
- but the *memory kernel* may nevertheless vanish separately. It carries the additional right factor $Q^\perp LP$, so it vanishes identically whenever $Q^\perp LP = 0$ — the case where the generator feeds unresolved directions into resolved ones but not conversely — and it may also vanish through cancellation in the intervening propagator $e^{(t-\tau)Q^\perp LQ^\perp}$.

We therefore do *not* claim that memory must be nonzero whenever closure fails. Failure of closure guarantees a nonzero initial-state response for some unresolved datum; it guarantees nothing about the memory term.

Fourth: what the unresolved-initial-state term signifies. It depends on $Q^\perp\Psi(0)$, which is exactly the fiber coordinate — the information the readout discards. Its presence is the analytic signature of the non-source projection: the projected law is not merely inconvenient but genuinely underdetermined by $P\Psi(t)$ alone. Calling the term “noise” is an additional modelling decision (cf. §5.2(2)), not a theorem; and so, equally, is any replacement of it by a memory kernel.

5.4. Examples.

Example 5.5 (Noninjective but fiber-preserving: coherent laws without source completeness). Let $S = \mathbb{Z}_2 \times \mathbb{Z}_2 = \{(x, y)\}$ with $x, y \in \{0, 1\}$, all four states physically inequivalent. Let the readout forget the second coordinate:

$$p(x, y) = x, \quad T_1 = \mathbb{Z}_2.$$

Fibers are $p^{-1}(x) = \{(x, 0), (x, 1)\}$ — doubletons, so p is not injective. Let the discrete-time source evolution be

$$\Phi(x, y) = (x + 1, y + 1) \pmod{2}.$$

Then $p(\Phi(x, y)) = x + 1$ depends only on x , so (17) holds and the readout law is autonomous:

$$U(x) = x + 1 \pmod{2}.$$

An observer confined to T_1 sees a perfectly deterministic period-2 world, exactly predicted by an autonomous law, with no residue and no anomaly whatever. Yet p is noninjective, and the

hidden coordinate y evolves nontrivially the entire time. This is Corollary 5.3 in miniature: *the readout is closed and the source is strictly larger*. If, moreover, some target Q depends on y — say $a_Q(x, y) = y$ — then (13) holds for $s_0 = (0, 0)$, $s_1 = (0, 1)$, and by Theorem 4.5 the autonomous, coherent readout T_1 is a non-source projection relative to Q . Autonomy buys no protection against this.

Example 5.6 (Fiber-breaking: failure of autonomous projected evolution). Keep $S = \mathbb{Z}_2 \times \mathbb{Z}_2$ and $p(x, y) = x$, but now let

$$\Phi(x, y) = (x + y, y) \pmod{2},$$

so that the hidden coordinate *drives* the visible one. Take

$$s_0 = (0, 0), \quad s_1 = (0, 1), \quad p(s_0) = p(s_1) = 0.$$

One step gives $\Phi(s_0) = (0, 0)$ and $\Phi(s_1) = (1, 1)$, whence

$$p(\Phi(s_0)) = 0 \neq 1 = p(\Phi(s_1)),$$

which is (18). By Theorem 5.2, no map $U: T_1 \rightarrow T_1$ can satisfy (16): it would have to send the single readout value 0 to both 0 and 1. The readout observer sees the state 0 sometimes remain 0 and sometimes become 1, with nothing in the readout to explain the difference.

What repairs it is instructive and matches §5.2 exactly. (i) *Completion*: adjoin y — take $E = S$ with $e = \text{id}$ — and the dynamics is deterministic and closed on E . Indeed the minimal sufficient extension for the target “next readout” is $E_Q = \text{Im}(p, a_Q)$ with $a_Q(x, y) = x + y$, which one computes to be $\{(0, 0), (0, 1), (1, 0), (1, 1)\}$, i.e. all four pairs $(x, x + y)$ — a four-element set properly refining the two-element T_1 , and in this instance recovering S . (ii) *Statistics*: posit a within-fiber distribution, say $y \sim \text{Bernoulli}(\frac{1}{2})$ independent of x , sampled *once* at the initial time. It is essential to describe the resulting readout process correctly, and it is *not* a symmetric random walk with independent increments. Because Φ leaves y fixed, the *same* sampled y controls every subsequent step:

$$x_{n+1} = x_n + y \pmod{2} \quad \text{for all } n,$$

with y drawn once and thereafter constant. Hence:

- the *first* transition stays ($y = 0$) or flips ($y = 1$) with probability $\frac{1}{2}$ each, which coincides with what a symmetric random walk would predict;
- all *subsequent* increments are *perfectly correlated* with the first, since the increment is the same y at every step — the walk either never moves or alternates forever;
- consequently the observed history (x_0, x_1) already reveals y exactly, namely $y = x_1 - x_0$, and every later step is thereafter deterministic;
- the projected process is therefore a *mixture of two deterministic trajectories* — the constant one and the period-two alternating one, each with weight $\frac{1}{2}$ — and not an i.i.d. random walk.

The randomness is entirely in the *initial draw* of the fiber coordinate, not in the transitions. This is precisely the distinction the section is for: a within-fiber *distribution* is not a Markov law on the instantaneous readout. A genuine symmetric random walk on $\mathbb{Z}_2$ would resample an independent increment at every step and would have a well-defined Markov generator on T_1 ; the present process has neither, and indeed by Theorem 5.2 it cannot, since Φ is fiber-breaking. Note also that the Bernoulli law is an *extra assumption*: the deterministic data (S, p, Φ) does not supply it, and a different within-fiber law gives a different projected process. (iii) *Memory*: here y is conserved and, as just computed, is exposed by the readout’s own first transition, so a one-step memory suffices to restore prediction. This is a special feature of this example — y

happens to be both conserved and immediately visible in the readout increment — and not a general remedy; in general the past readout need not determine the fiber coordinate at all.

6. JOINT PROBES AND SOURCE IDENTIFIABILITY

Definition 6.1 (Probe family and joint map). Let I be an index set and $P_i: S \rightarrow Y_i$, $i \in I$, a family of *admissible probes* — maps whose values an observer can in principle obtain. Define the *joint probe map*

$$P_I: S \longrightarrow \prod_{i \in I} Y_i, \quad P_I(s) = (P_i(s))_{i \in I},$$

and corestrict it to its image $\text{Im } P_I$, so that $P_I: S \twoheadrightarrow \text{Im } P_I$ is surjective. Write $s \sim_{P_i} s' \iff P_i(s) = P_i(s')$ and

$$\sim_I := \bigcap_{i \in I} \sim_{P_i},$$

the *joint observational equivalence relation*. The family *identifies the source* if $P_I(s) = P_I(s') \Rightarrow s = s'$.

Theorem 6.2 (Joint-probe identifiability). *The following are equivalent.*

- (I1) *The probe family identifies the source.*
- (I2) *P_I is injective.*
- (I3) *The probes separate the points of S : for all $s \neq s'$ there is $i \in I$ with $P_i(s) \neq P_i(s')$.*
- (I4) *$\sim_I = \Delta_S$, the diagonal (i.e. equality).*
- (I5) *There is a left inverse on the image: a map $R_I: \text{Im } P_I \rightarrow S$ with $R_I \circ P_I = \text{id}_S$.*

Proof. (I1) $\Leftrightarrow$ (I2) is the definition of injectivity. (I2) $\Leftrightarrow$ (I3): $P_I(s) = P_I(s')$ holds iff $P_i(s) = P_i(s')$ for all i ; so P_I is injective iff no distinct pair agrees on every probe, which is (I3). (I3) $\Leftrightarrow$ (I4): $s \sim_I s'$ means $P_i(s) = P_i(s')$ for all i ; so $\sim_I$ reduces to equality iff distinct points are separated by some probe.

(I2) $\Rightarrow$ (I5). $P_I: S \rightarrow \text{Im } P_I$ is injective by (I2) and surjective by corestriction, hence bijective; take $R_I := P_I^{-1}$, which satisfies $R_I \circ P_I = \text{id}_S$.

(I5) $\Rightarrow$ (I2). If $P_I(s) = P_I(s')$ then $s = R_I(P_I(s)) = R_I(P_I(s')) = s'$. $\square$

Corollary 6.3 (Residual ambiguity). *If $\sim_I \neq \Delta_S$, then the maximally identifiable object from the probe family is the quotient*

$$S / \sim_I \xrightarrow{\cong} \text{Im } P_I, \quad (25)$$

the bijection being induced by P_I . No rule whose input is the probe values alone reconstructs individual source states within a nontrivial $\sim_I$ -class.

Proof. $\sim_I$ is precisely the kernel pair of P_I : $s \sim_I s' \iff P_I(s) = P_I(s')$. Hence $[s] \mapsto P_I(s)$ is well defined and injective on $S / \sim_I$, and surjective onto $\text{Im } P_I$ by corestriction, giving (25). For the second claim, let $D: \text{Im } P_I \rightarrow Z$ be *any* map (a “rule”) and let $s \sim_I s'$. Then $D(P_I(s)) = D(P_I(s'))$: the rule returns the same output on both, so it cannot name which of s, s' was realized. $\square$

Proposition 6.4 (Monotonicity under added probes). *If $I \subseteq J$ (the probes of I being among those of J), then*

$$\sim_J \subseteq \sim_I,$$

and there is a canonical surjection

$$S / \sim_J \twoheadrightarrow S / \sim_I, \quad [s]_J \longmapsto [s]_I. \quad (26)$$

Proof. If $s \sim_J s'$ then $P_i(s) = P_i(s')$ for every $i \in J \supseteq I$, hence for every $i \in I$, i.e. $s \sim_I s'$. This inclusion of equivalence relations makes (26) well defined ($[s]_J = [s']_J \Rightarrow s \sim_J s' \Rightarrow s \sim_I s' \Rightarrow [s]_I = [s']_I$), and it is surjective because every class $[s]_I$ is hit by $[s]_J$. $\square$

Remark 6.5 (Direction of (26)). The map goes from the *finer* quotient to the *coarser* one. More probes means a finer partition, hence a larger quotient set, which surjects onto the smaller. There is in general no map back: information present in $S/\sim_J$ cannot be manufactured from $S/\sim_I$.

Thus *tomography is the intersection of probe fibers*: one adds probes, and the joint fiber $\{s' : s' \sim_I s\}$ shrinks. Source recovery occurs precisely when that intersection reaches the diagonal. But we must be exact about what is and is not being claimed. Theorem 6.2 says *when* a family separates; it does not say that any particular physically available family *does*. Whether the probes an observer can actually run separate the source is a contingent fact about the model and the observer, not a theorem. The next section studies the case where, by hypothesis, they do not.

7. THE INTERNAL OBSERVATION CEILING

Assumption 7.1 (Internal-probe factorization). Every admissible probe factors through the readout: for each $i \in I$ there is $\bar{P}_i: T_1 \rightarrow Y_i$ with

$$P_i = \bar{P}_i \circ p. \quad (27)$$

This is an *assumption*, supplied by the model, and it is the entire engine of what follows. It says that a probe may process, compare, transform, amplify, or combine readout information in any way whatever — $\bar{P}_i$ is an arbitrary map, subject to no continuity, computability, or complexity restriction — but that it receives *no input from outside the readout state*. A probe coupling directly to a fiber coordinate violates (27), and the theorem below then simply does not apply to it.

It is tempting to say that an observer “built out of Tier-1 material”, or “interacting only with Tier-1 degrees of freedom”, automatically satisfies (27). That is an unsupported inference, and we do not make it. An observer of that description is represented by Assumption 7.1 *only if the physical model independently establishes that the complete input of every admissible probe factors through p* . Being *called* internal, or being composed of readout-level constituents, does not by itself prove the factorization: whether some residual coupling to a fiber coordinate exists is a substantive question about the model’s dynamics and interactions, and it must be settled there, not by the label. Assumption 7.1 is thus a hypothesis to be discharged, and §10 lists it among the things a physical instantiation must supply.

Theorem 7.2 (Internal identifiability limit). *Under Assumption 7.1,*

$$\sim_p \subseteq \sim_I. \quad (28)$$

Consequently:

- (L1) *If p is noninjective, the internal probe family does not identify the source.*
- (L2) *For every map $D: \text{Im } P_I \rightarrow Z$ — statistic, post-processing rule, inference procedure, deterministic selector, or arbitrary function — and every s, s' with $p(s) = p(s')$, one has $D(P_I(s)) = D(P_I(s'))$.*

Proof. Let $p(s) = p(s')$. For every $i \in I$, using (27),

$$P_i(s) = \bar{P}_i(p(s)) = \bar{P}_i(p(s')) = P_i(s').$$

Since this holds for all i , $P_I(s) = P_I(s')$, i.e. $s \sim_I s'$. This is (28).

(L1): if p is noninjective there are $s \neq s'$ with $p(s) = p(s')$, hence $s \sim_I s'$ by (28), so $\sim_I \neq \Delta_S$ and Theorem 6.2(I4) fails.

(L2): $P_I(s) = P_I(s')$ was just shown; applying the function D to equal arguments yields equal values. Note that no property of D beyond being a function is used — which is exactly why no cleverness in post-processing can evade the conclusion. $\square$

Corollary 7.3 (No internal selector). *If a readout fiber contains two physically inequivalent source states, then no inference procedure whose complete input factors through the readout can determine which member of that fiber was realized.*

Proof. Immediate from (L2) with $Z = S$ and D the putative selector: it returns the same value on s and on s' , so it is wrong on at least one of them. $\square$

Remark 7.4 (Exactly what this ceiling is, and is not). Five clarifications, each of which matters.

(a) *The mechanism is equality of inputs.* The obstruction in Theorem 7.2 arises because the probes' arguments coincide: $P_I(s)$ and $P_I(s')$ are the same object. Nothing is being said about the difficulty, cost, or computability of any inference. A function of equal inputs returns equal outputs; that is all.

(b) *It is not Gödel incompleteness.* There is no self-reference, no arithmetization, no diagonalization, and no undecidable sentence anywhere in the proof. The result is an identifiability theorem in the sense of [9, 10], and conflating it with [6] is a category error.

(c) *It does not apply to a genuinely additional probe.* If some probe P_j does *not* factor through p — if it couples to the fiber coordinate — then Assumption 7.1 fails for the enlarged family, (28) need not hold, and by Proposition 6.4 the joint fibers may strictly shrink, possibly to the diagonal. *The ceiling is a consequence of the assumed input restriction, and it is lifted exactly by lifting that restriction.*

(d) *It is conditional on completeness of the admissible family.* Assumption 7.1 asserts that *every* probe available to the observer factors through p . That assertion is a substantive physical claim about the model, and it can be false. Consequently Theorem 7.2 must *not* be inflated into a general claim that internal observers can never obtain evidence about a source. It concerns exactly those source distinctions that lie within the joint fibers of all allowed probes — and it says that *those* distinctions, and only those, are internally unrecoverable.

(e) *It is a deterministic ceiling, and does not by itself yield the statistical one.* Theorem 7.2 concerns deterministic probe values $P_i(s)$ and deterministic functions D of them. What it rules out is accordingly a deterministic selector, and any post-processing, statistic, or inference rule *applied to those probe values*. It does *not*, on its own, license the far stronger assertions that “no test” or “no amount of data” can discriminate. Those are claims about *outcome laws*, and they require the separate hypothesis $\mu_s = \mathbf{M}_{p(s)}$ of (30) — namely, that the *complete data law* of the experiment factors through p . One may of course *stipulate* that the complete experiment is itself one of the admissible probes, in which case (27) for that probe *is* (30); but that identification is an additional modelling step, and must be made explicitly. Absent it, a random experiment whose law depends on the fiber coordinate is perfectly compatible with Assumption 7.1 holding for every deterministic probe. The two ceilings are stated separately, and proved from separate hypotheses, in Theorem 7.2 and Corollary 8.4; see also Proposition 8.6(R3).

8. STATISTICAL IDENTIFIABILITY

The preceding sections treated probes as deterministic maps, and the ceiling obtained there is correspondingly deterministic. Real experiments return random data, and the ceiling must therefore be established *again*, from its own hypothesis, at the level of probability laws. We

stress that this is not a restatement of Theorem 7.2 but an independent result resting on the separate assumption (30): only at the level of complete outcome laws does one earn the stronger conclusion that no test, at any sample size and with any randomization, can discriminate.

Definition 8.1 (Statistical experiment). Let $(\mathcal{D}, \Sigma)$ be a measurable space, the *data space*. Suppose each source state $s \in S$ induces a probability measure μ_s on $(\mathcal{D}, \Sigma)$, the *complete outcome law* of the experiment (or of a whole family of experiments, in which case $\mathcal{D}$ is the corresponding product space and μ_s the joint law of all of them, including all repetitions). The experiment *identifies the source* if $\mu_s = \mu_{s'} \Rightarrow s = s'$.

We use total variation in the convention

$$\|\mu - \nu\|_{\text{TV}} := \sup_{A \in \Sigma} |\mu(A) - \nu(A)|, \quad (29)$$

so that $\|\mu - \nu\|_{\text{TV}} \in [0, 1]$ for probability measures. (Some authors insert a factor 2; nothing below depends on the choice, but the statements are made for (29).)

Theorem 8.2 (Statistical indistinguishability). *Let $s_0, s_1 \in S$.*

- (S1) *If $\mu_{s_0} = \mu_{s_1}$, then for every measurable statistic $f: (\mathcal{D}, \Sigma) \rightarrow (Z, \mathcal{Z})$ the pushforwards agree: $f_*\mu_{s_0} = f_*\mu_{s_1}$.*
- (S2) *If $\mu_{s_0} = \mu_{s_1}$, then every randomized test $\varphi: \mathcal{D} \rightarrow [0, 1]$ (measurable) has equal rejection probability:*

$$\int_{\mathcal{D}} \varphi d\mu_{s_0} = \int_{\mathcal{D}} \varphi d\mu_{s_1}.$$

- (S3) *If $\mu_{s_0} = \mu_{s_1}$, then $\|\mu_{s_0} - \mu_{s_1}\|_{\text{TV}} = 0$.*
- (S4) *Conversely, if $\mu_{s_0} \neq \mu_{s_1}$, then there exists $A \in \Sigma$ with $\mu_{s_0}(A) \neq \mu_{s_1}(A)$, and $\|\mu_{s_0} - \mu_{s_1}\|_{\text{TV}} > 0$.*

Proof. (S1) For $B \in \mathcal{Z}$, $(f_*\mu_{s_0})(B) = \mu_{s_0}(f^{-1}(B)) = \mu_{s_1}(f^{-1}(B)) = (f_*\mu_{s_1})(B)$, using $f^{-1}(B) \in \Sigma$ by measurability and equality of the measures on Σ .

(S2) φ is a bounded measurable function and the two measures are equal, so the integrals coincide. (Explicitly: the integral of a fixed bounded measurable function is determined by the measure, being the limit of integrals of simple functions, each of which is a finite combination of measure values.)

(S3) Immediate from (29): every summand $|\mu_{s_0}(A) - \mu_{s_1}(A)|$ vanishes.

(S4) Two measures on $(\mathcal{D}, \Sigma)$ are *equal* by definition iff they agree on every $A \in \Sigma$. Hence $\mu_{s_0} \neq \mu_{s_1}$ means precisely that some $A \in \Sigma$ has $\mu_{s_0}(A) \neq \mu_{s_1}(A)$; and then the supremum in (29) is at least $|\mu_{s_0}(A) - \mu_{s_1}(A)| > 0$. $\square$

Remark 8.3 (The three distinctions that must not collapse). Theorem 8.2 is easy to over-read in *both* directions.

Forward. (S1)–(S3) say that when the complete outcome laws are *exactly equal*, statistical discrimination is not difficult but *impossible*: not by any statistic, not by any test, not at any sample size, not with any amount of randomization, not asymptotically. There is no procedure to improve, because the two hypotheses induce literally the same distribution over all possible data. This is the strongest negative available, and it is what Corollary 8.4 exploits.

Backward. (S4) says only that unequal laws *differ on some event*. It emphatically does *not* say that a single observation classifies the state with certainty. If $\mu_{s_0}(A) = 0.51$ and $\mu_{s_1}(A) = 0.49$, the laws are unequal, $\|\mu_{s_0} - \mu_{s_1}\|_{\text{TV}} \geq 0.02 > 0$, and yet one observation of $\mathbf{1}_A$ is nearly worthless. What (S4) supplies is a *nonzero* signal, hence — given i.i.d. repetition, which is an additional assumption about $\mathcal{D}$ and μ_s — consistency of a suitable test as the

sample size grows. One must therefore keep apart: *exact identifiability* (distinct states have distinct laws); *asymptotic identifiability* (some consistent procedure separates them in the limit); and *finite-sample power* (how much data a given error level costs). Only the first is a purely measure-theoretic property, and only the failure of the first is what §7 and Corollary 8.4 are about.

Corollary 8.4 (Statistical readout ceiling). *Suppose the complete outcome law depends on the source only through the readout: there is a family $\{\mathbf{M}_t\}_{t \in T_1}$ of probability measures on $(\mathcal{D}, \Sigma)$ with*

$$\mu_s = \mathbf{M}_{p(s)} \quad \text{for all } s \in S. \quad (30)$$

Then any two source states in a common readout fiber are statistically indistinguishable by that experiment: for $p(s_0) = p(s_1)$, every statistic has the same law, every test has the same power, and the total-variation distance is zero.

Proof. If $p(s_0) = p(s_1)$ then (30) gives $\mu_{s_0} = \mathbf{M}_{p(s_0)} = \mathbf{M}_{p(s_1)} = \mu_{s_1}$. Apply Theorem 8.2(S1)–(S3). $\square$

Remark 8.5 (The two ceilings, stated separately). Corollary 8.4 is the probabilistic analogue of Theorem 7.2, and (30) the probabilistic analogue of Assumption 7.1. It is not a consequence of them. The two results must be kept apart, because their hypotheses differ and so do their conclusions.

Deterministic ceiling. HYPOTHESIS: every admissible probe factors through the readout, $P_i = \tilde{P}_i \circ p$. CONCLUSION: within a readout fiber, all probe values coincide, so no deterministic selector, statistic, or post-processing rule *applied to those values* can distinguish its members (Theorem 7.2, Corollary 7.3).

Statistical ceiling. HYPOTHESIS: the complete outcome law factors through the readout, $\mu_s = \mathbf{M}_{p(s)}$. CONCLUSION: within a readout fiber, the complete data laws are *equal*, so no statistic, no randomized test, no sample size, and no asymptotic procedure can discriminate; the total-variation distance is zero (Corollary 8.4).

The second does not follow from the first. Deterministic factorization constrains probe *values*; it says nothing about the law of a stochastic experiment unless that experiment's complete law is itself identified as a probe factoring through p — an identification which must be made explicitly, and which is exactly the content of (30). Both are hypotheses about the model, to be discharged by the model. Neither is a theorem about observers in general.

8.1. The exact relations among the three obstructions. We can now state precisely how kinematic incompleteness, dynamical nonclosure, and statistical nonidentifiability are related. They are *not* logically independent, and it is worth being exact, since the pattern of implication is what licenses — and what forbids — inferring one from another.

Proposition 8.6 (Logical relations). *Let $p: S \rightarrow T_1$ be an exact readout.*

- (R1) **Dynamical nonclosure implies noninjectivity.** *If Φ_t is fiber-breaking (Case II of Theorem 5.2), then p is noninjective. Thus kinematic incompleteness is necessary for fiber-breaking.*
- (R2) **The converse fails.** *Noninjectivity of p does not imply fiber-breaking: there exist noninjective p with fiber-preserving, autonomous dynamics. Hence kinematic incompleteness is not sufficient for dynamical nonclosure, and, equivalently, autonomy of the projected law does not imply injectivity of p .*
- (R3) **Statistical nonidentifiability is unconditionally unrelated to p .** *Without a hypothesis tying $\{\mu_s\}$ to p , the existence of $s_0 \neq s_1$ with $\mu_{s_0} = \mu_{s_1}$ neither implies nor is implied by noninjectivity of p .*

(R4) ***Under the factorization hypothesis it is implied.*** If $\mu_s = M_{p(s)}$ as in (30) and p is noninjective, then there exist $s_0 \neq s_1$ with $\mu_{s_0} = \mu_{s_1}$: statistical nonidentifiability holds within at least one readout fiber.

Proof. (R1) Case II supplies s_0, s_1 with $p(s_0) = p(s_1)$ and $p(\Phi_t(s_0)) \neq p(\Phi_t(s_1))$ for some t . The latter forces $\Phi_t(s_0) \neq \Phi_t(s_1)$, hence $s_0 \neq s_1$. So p identifies two distinct points and is noninjective.

(R2) Example 5.5 exhibits a noninjective p with $p \circ \Phi_t = U_t \circ p$ for an autonomous U_t . The equivalent formulation is the contrapositive: if autonomy implied injectivity, that example would be impossible.

(R3) Both directions fail, by example. *Noninjectivity without statistical nonidentifiability:* take p noninjective with fiber $\{s_0, s_1\}$, and let the experiment couple to a fiber coordinate — say $\mathcal{D} = \{0, 1\}$ with $\mu_{s_0} = \delta_0$ and $\mu_{s_1} = \delta_1$. The laws are distinct (indeed mutually singular, so a single sample identifies the state), while p is noninjective. *Statistical nonidentifiability without noninjectivity:* take $p = \text{id}_S$, which is injective, and let $\mu_s = M$ be one fixed law for every s — an experiment that measures nothing. Then $\mu_{s_0} = \mu_{s_1}$ for all pairs while p is bijective. So no unconditional implication holds in either direction; the outcome laws are simply extra data, unconstrained by p until a hypothesis links them.

(R4) Noninjectivity supplies $s_0 \neq s_1$ with $p(s_0) = p(s_1)$. Then (30) gives $\mu_{s_0} = M_{p(s_0)} = M_{p(s_1)} = \mu_{s_1}$, and Theorem 8.2(S1)–(S3) applies to that pair. $\square$

The moral of (R3) and (R4) together is that the statistical ceiling is *not* a corollary of the kinematic one. It requires the separate hypothesis (30), which is a claim about what the experiment couples to, and which the deterministic data (S, p, Φ_t) does not supply. We return to this in Remark 8.5.

Example 8.7 (Equal laws versus unequal but overlapping laws). Let $S = \{s_0, s_1, s_2\}$ with $p(s_0) = p(s_1) = t$, $p(s_2) = t'$ as in Example 4.13. Let $\mathcal{D} = \{0, 1\}$ with $\Sigma = 2^{\mathcal{D}}$, and let each state induce a Bernoulli law $\mu_s(\{1\}) = \theta_s$.

(i) *Equal laws.* Suppose the experiment factors through the readout as in (30): $\theta_{s_0} = \theta_{s_1} = 0.3$ and $\theta_{s_2} = 0.8$. Then $\mu_{s_0} = \mu_{s_1}$ exactly. Consider *any* test based on n i.i.d. repetitions — so the data space is $\{0, 1\}^n$ and the law is $\text{Bin}(n, \theta_s)$ coordinatewise. The joint laws under s_0 and s_1 are identical for every n . Hence for every test φ_n , the rejection probabilities coincide, the power against s_1 equals the size under s_0 , and no amount of data helps: the discrimination problem between s_0 and s_1 has *zero* information at every sample size. This is Corollary 8.4. Note that s_2 is *statistically identifiable* against both s_0 and s_1 : its law is different ($\theta_{s_2} = 0.8 \neq 0.3$), so by Theorem 8.2(S4) some event distinguishes it, and under i.i.d. repetition it is consistently distinguishable. It is *not* perfectly discriminable from a finite sample: the Bernoulli laws with parameters 0.3 and 0.8 have the same support $\{0, 1\}$, so every finite data set has positive probability under both hypotheses and no finite-sample test is infallible. The ceiling is confined to the fiber, but what lies outside the fiber is identifiability and asymptotic consistency, not certainty.

(ii) *Unequal but overlapping laws.* Now suppose the model is refined so that $\theta_{s_0} = 0.5$ and $\theta_{s_1} = 0.55$, so (30) fails. The laws are unequal: on the event $A = \{1\}$, $|\mu_{s_0}(A) - \mu_{s_1}(A)| = 0.05$, so $\|\mu_{s_0} - \mu_{s_1}\|_{\text{TV}} = 0.05 > 0$, and Theorem 8.2(S4) is satisfied. But a *single* observation is nearly useless: the optimal single-sample test has error probability

$$\frac{1}{2}(1 - \|\mu_{s_0} - \mu_{s_1}\|_{\text{TV}}) = \frac{1}{2}(1 - 0.05) = 0.475,$$

barely better than a coin flip. With n i.i.d. samples the laws are $\text{Bin}(n, 0.5)$ and $\text{Bin}(n, 0.55)$, which separate as $n \rightarrow \infty$: the states are asymptotically identifiable and a consistent test exists.

The finite-sample power, however, depends quantitatively on the error levels one demands, and we make no numerical claim about the required n here, since that would need a testing criterion and a calculation we do not supply.

The example thus exhibits all four notions that Remark 8.3 insists on keeping apart. *Equal laws* (case (i), the pair s_0, s_1): discrimination is impossible at every sample size. *Unequal but overlapping laws* (case (ii), and also s_2 against s_0 in case (i)): the laws differ, so some event distinguishes them, but their common support means no finite sample is conclusive. *Perfect finite-sample discrimination*: absent here, and it would require laws that are mutually singular — e.g. disjoint supports, as in the degenerate δ_0 versus δ_1 of Proposition 8.6(R3). *Asymptotic consistency*: available in every unequal-law case above. Only the first is a ceiling; the others are matters of cost.

9. EXACT MARKOVIAN PROJECTION: STRONG LUMPABILITY

The dynamical descent problem of §5 has a sharp stochastic counterpart, in which the question “does the projected process close?” admits a complete and checkable answer.

Definition 9.1 (Setting and conventions). Let $S = \{1, \dots, N\}$ be a finite source state space and let $(X_t)_{t \geq 0}$ be a continuous-time, time-homogeneous Markov chain on S with generator $G = (g_{ij})_{i,j \in S}$ satisfying the standard conventions

$$g_{ij} \geq 0 \quad (i \neq j), \quad \sum_{j \in S} g_{ij} = 0 \quad \text{for every } i, \quad (31)$$

so that $g_{ii} = -\sum_{j \neq i} g_{ij} \leq 0$. Transition probabilities are $P(t) = e^{tG}$, and a row-vector law $\pi(t)$ obeys the (forward) master equation

$$\dot{\pi}_j(t) = \sum_{i \in S} \pi_i(t) g_{ij}. \quad (32)$$

For $f: S \rightarrow \mathbb{R}$ the backward generator acts by $(Gf)(i) = \sum_j g_{ij} f(j)$.

Let $p: S \rightarrow T$ be a surjection onto a finite set T , partitioning S into the *blocks*

$$B_a := p^{-1}(a), \quad a \in T.$$

Definition 9.2 (Strong lumpability). The chain is *strongly lumpable* with respect to p if, for every initial source distribution $\pi(0)$, the projected process $(p(X_t))_{t \geq 0}$ is a time-homogeneous Markov chain on T whose transition law depends on the initial data only through the initial block distribution $\bar{\pi}(0)$.

The quantifier “for every initial distribution” is the whole point, and it is what distinguishes *strong* lumpability from the weaker, initial-distribution-dependent notion. A projected process can perfectly well be Markov for one carefully chosen $\pi(0)$ — e.g. one supported on a single state, or a stationary within-block law — while failing to be Markov for another. Such a coincidence gives an observer a Markov description that is valid for their particular preparation and invalid in general. Strong lumpability demands that the projected description be *universally* valid, and that is the property that corresponds to genuine closure of the readout dynamics in the sense of §5.

Theorem 9.3 (Strong lumpability). *The chain is strongly lumpable with respect to p if and only if the block-sum condition holds:*

$$\sum_{j \in B_b} g_{ij} = \sum_{j \in B_b} g_{i'j} \quad \text{for all blocks } B_a, B_b \text{ and all } i, i' \in B_a. \quad (33)$$

In that case the common value

$$\bar{g}_{ab} := \sum_{j \in B_b} g_{ij} \quad (\text{any } i \in B_a) \quad (34)$$

is well defined and $\bar{G} = (\bar{g}_{ab})$ is a valid Markov generator on T , which generates the projected chain.

Proof. Necessity. Suppose the projected process is Markov for every initial distribution, with a transition law determined by $\bar{\pi}(0)$. Fix a block B_a and $i, i' \in B_a$, and consider the two initial distributions δ_i and $\delta_{i'}$. Both project to the same initial block law, namely δ_a on T . Hence by hypothesis the projected chains started from δ_i and from $\delta_{i'}$ have the *same* law, and in particular the same instantaneous transition rate out of a into any block B_b with $b \neq a$.

Compute that rate from the source. Starting at i , for $b \neq a$,

$$\Pr[p(X_h) = b \mid X_0 = i] = \sum_{j \in B_b} (e^{hG})_{ij} = \sum_{j \in B_b} (\delta_{ij} + h g_{ij} + O(h^2)) = h \sum_{j \in B_b} g_{ij} + O(h^2),$$

where $\delta_{ij} = 0$ for $j \in B_b$ since $i \in B_a$ and $B_a \cap B_b = \emptyset$. Dividing by h and letting $h \downarrow 0$, the instantaneous $a \rightarrow b$ rate from i is $\sum_{j \in B_b} g_{ij}$, and likewise from i' it is $\sum_{j \in B_b} g_{i'j}$. Equality of the two projected laws forces

$$\sum_{j \in B_b} g_{ij} = \sum_{j \in B_b} g_{i'j} \quad (b \neq a).$$

For $b = a$ the condition follows from (31): since row sums vanish,

$$\sum_{j \in B_a} g_{ij} = - \sum_{b \neq a} \sum_{j \in B_b} g_{ij} = - \sum_{b \neq a} \sum_{j \in B_b} g_{i'j} = \sum_{j \in B_a} g_{i'j}.$$

This is (33).

Sufficiency. Assume (33), so that $\bar{g}_{ab}$ in (34) is well defined.

$\bar{G}$ is a generator. For $a \neq b$ and $i \in B_a$, every $j \in B_b$ satisfies $j \neq i$, so $g_{ij} \geq 0$ by (31); hence $\bar{g}_{ab} = \sum_{j \in B_b} g_{ij} \geq 0$. And the row sums vanish:

$$\sum_{b \in T} \bar{g}_{ab} = \sum_{b \in T} \sum_{j \in B_b} g_{ij} = \sum_{j \in S} g_{ij} = 0,$$

using that the blocks partition S . So $\bar{G}$ satisfies (31) on T .

(1) *Closure of the block probabilities.* Let $\pi(t)$ solve (32) and set $\bar{\pi}_a(t) := \sum_{i \in B_a} \pi_i(t)$. Then

$$\begin{aligned} \dot{\bar{\pi}}_b(t) &= \sum_{j \in B_b} \dot{\pi}_j(t) = \sum_{j \in B_b} \sum_{i \in S} \pi_i(t) g_{ij} = \sum_{a \in T} \sum_{i \in B_a} \pi_i(t) \left(\sum_{j \in B_b} g_{ij} \right) \\ &= \sum_{a \in T} \sum_{i \in B_a} \pi_i(t) \bar{g}_{ab} = \sum_{a \in T} \bar{\pi}_a(t) \bar{g}_{ab}, \end{aligned}$$

where the fourth equality is exactly (33): the inner sum is the *same number* $\bar{g}_{ab}$ for every $i \in B_a$, so it factors out of the sum over $i \in B_a$ regardless of how $\pi(t)$ distributes mass *within* the block. Thus $\bar{\pi}$ obeys the closed master equation $\dot{\bar{\pi}} = \bar{\pi} \bar{G}$, for every initial $\pi(0)$.

(2) *Process-level Markovianity.* Closure of the one-time marginals is *not* by itself the Markov property, and we now establish the latter. Let $\mathcal{F} := \{f: S \rightarrow \mathbb{R} : f \text{ is constant on each block}\}$, and for $f \in \mathcal{F}$ write f_b for its value on B_b . For $i \in B_a$,

$$(Gf)(i) = \sum_{j \in S} g_{ij} f(j) = \sum_{b \in T} \sum_{j \in B_b} g_{ij} f_b = \sum_{b \in T} f_b \left(\sum_{j \in B_b} g_{ij} \right) = \sum_{b \in T} \bar{g}_{ab} f_b, \quad (35)$$

using (33) in the last step. The right-hand side depends on i only through its block a . Hence $Gf \in \mathcal{F}$: the subspace $\mathcal{F}$ of block-constant functions is invariant under G . Being finite-dimensional and G -invariant, $\mathcal{F}$ is invariant under e^{tG} for all $t \geq 0$ (expand the exponential; each partial sum preserves $\mathcal{F}$, and $\mathcal{F}$ is closed).

Apply this with $f = \mathbf{1}_{B_b} \in \mathcal{F}$. Invariance gives that $i \mapsto (e^{tG} \mathbf{1}_{B_b})(i) = \Pr[p(X_t) = b \mid X_0 = i]$ is constant on blocks: writing $\bar{P}_{ab}(t)$ for its common value on B_a ,

$$\Pr[p(X_t) = b \mid X_0 = i] = \bar{P}_{ab}(t) \quad \text{for every } i \in B_a. \quad (36)$$

Moreover (35) identifies the generator of $\bar{P}(t)$: reading (35) in $\mathcal{F} \cong \mathbb{R}^T$ shows $G|_{\mathcal{F}}$ acts as $\bar{G}$, so $\bar{P}(t) = e^{t\bar{G}}$.

Now take any times $0 \leq t_1 < \dots < t_n < t_{n+1}$, any blocks $b_1, \dots, b_{n+1}$, and any initial law $\pi(0)$. Write

$$H := \{p(X_{t_1}) = b_1, \dots, p(X_{t_n}) = b_n\}$$

for the projected past, and assume $\Pr[H] > 0$; histories of zero probability impose no condition, since the Markov property is a statement about conditional probabilities, which are defined only on positive-probability events. (Equivalently, one may phrase the argument with regular conditional distributions, which exist automatically here because S is finite; nothing changes.)

Let ν denote the conditional law of X_{t_n} given H , i.e. $\nu(i) := \Pr[X_{t_n} = i \mid H]$. It is a probability measure supported in B_{b_n} , and it may depend intricately on the whole past H and on $\pi(0)$. The Markov property of X at time t_n , together with (36), gives

$$\begin{aligned} \Pr[p(X_{t_{n+1}}) = b_{n+1} \mid H] &= \sum_{i \in B_{b_n}} \nu(i) \Pr[p(X_{t_{n+1}-t_n}) = b_{n+1} \mid X_0 = i] \\ &= \sum_{i \in B_{b_n}} \nu(i) \bar{P}_{b_n b_{n+1}}(t_{n+1} - t_n) \\ &= \bar{P}_{b_n b_{n+1}}(t_{n+1} - t_n), \end{aligned} \quad (37)$$

where the second equality is (36) — the summand is *independent of i* — and the third follows because ν is a probability measure on B_{b_n} , so its masses sum to 1. The right-hand side of (37) depends on the past only through b_n , and on the times only through $t_{n+1} - t_n$. This is exactly the Markov property with time-homogeneous transition semigroup $\bar{P}(t) = e^{t\bar{G}}$, and it holds for *every* $\pi(0)$ — which is the assertion of strong lumpability.

The crux, and the reason marginal closure alone would not suffice, is visible in (37): the unknown within-block conditional law ν — which is precisely the information the readout discards, and which *does* depend on the whole past — drops out only because $\bar{P}_{b_n b_{n+1}}$ does not depend on which $i \in B_{b_n}$ carries the mass. That independence is (36), i.e. (33). $\square$

Corollary 9.4 (Failure of lumpability). *Suppose (33) fails: there are a block B_a , states $i, i' \in B_a$, and a block B_b with $\sum_{j \in B_b} g_{ij} \neq \sum_{j \in B_b} g_{i'j}$. Then the two initial distributions δ_i and $\delta_{i'}$ — which have the same projected initial state a — generate different instantaneous projected transition rates into b . Consequently the instantaneous readout state is insufficient for a universal Markov description of the projected process.*

Proof. The instantaneous rates were computed in the necessity half of Theorem 9.3 to be $\sum_{j \in B_b} g_{ij}$ and $\sum_{j \in B_b} g_{i'j}$, which differ by hypothesis. A time-homogeneous Markov generator on T would assign a single rate $\bar{g}_{ab}$ to the projected state a , contradiction. $\square$

Remark 9.5 (What failure of lumpability does and does not establish). Corollary 9.4 is a *negative* result: the instantaneous projected state does not support a universal Markov law. It

does *not* by itself establish that the projected process is governed by any *particular* memory kernel or non-Markovian equation. Producing such a kernel requires an additional reduction construction — a Nakajima–Zwanzig projection with a chosen within-block reference law, or a completion in the sense of Definition 4.3 — and different choices give different kernels.

It is worth being exact about what determines the within-block conditional law ν , since this is the discarded information. The generator G *together with an initial distribution* $\pi(0)$ determines it: the pair $(G, \pi(0))$ fixes the full law of the source process, hence the conditional law of X_t given any positive-probability projected history. The generator *alone* does not: without an initial law — equivalently, without a specification of how the system was prepared — there is no such thing as “the” within-block distribution, and different preparations consistent with the same projected state yield different ν , which is exactly why Corollary 9.4 bites. Compare §5.2(2): in the deterministic setting likewise, (S, p, Φ_t) does not determine the within-fiber distribution, which must be supplied. What is theorem here is the impossibility; what is modelling is the replacement.

Example 9.6 (A lumpable and a non-lumpable partition). Let $S = \{1, 2, 3\}$ with blocks $B_a = \{1, 2\}$ and $B_b = \{3\}$, so $T = \{a, b\}$.

(i) *Lumpable*. Take

$$G = \begin{pmatrix} -3 & 2 & 1 \\ 1 & -2 & 1 \\ 4 & 0 & -4 \end{pmatrix},$$

which satisfies (31) (nonnegative off-diagonals, zero row sums). Check (33) for the pair $i = 1, i' = 2 \in B_a$:

$$\sum_{j \in B_b} g_{1j} = g_{13} = 1, \quad \sum_{j \in B_b} g_{2j} = g_{23} = 1.$$

These agree, and the B_a column sums then agree automatically by the row-sum identity ($g_{11} + g_{12} = -3 + 2 = -1 = 1 - 2 = g_{21} + g_{22}$). So the chain is strongly lumpable, with projected generator

$$\bar{G} = \begin{pmatrix} \bar{g}_{aa} & \bar{g}_{ab} \\ \bar{g}_{ba} & \bar{g}_{bb} \end{pmatrix} = \begin{pmatrix} -1 & 1 \\ 4 & -4 \end{pmatrix},$$

a valid generator. Note the internal states 1 and 2 exchange probability at rates 2 and 1 and are in no sense equivalent — the readout is genuinely noninjective, its fiber $B_a = \{1, 2\}$ is a genuine two-element fiber (it is a set of two source states, not a vector space, and no dimension is being asserted of it), and the internal dynamics within the fiber is nontrivial — yet the projected process is *exactly* Markov for every initial condition. This is the stochastic analogue of Example 5.5 and of Corollary 5.3: closure of the projected law is no evidence of source completeness.

(ii) *Non-lumpable*. Now perturb a single entry:

$$G' = \begin{pmatrix} -3 & 2 & 1 \\ 5 & -8 & 3 \\ 4 & 0 & -4 \end{pmatrix},$$

again a valid generator. Test (33):

$$\sum_{j \in B_b} g'_{1j} = g'_{13} = 1, \quad \sum_{j \in B_b} g'_{2j} = g'_{23} = 3,$$

and $1 \neq 3$. By Corollary 9.4, the projected process is not Markov for all initial distributions: started at δ_1 the instantaneous $a \rightarrow b$ rate is 1; started at δ_2 it is 3. Both have the same projected initial state a . A readout observer preparing what they take to be “the state a ” will

see different escape statistics depending on a source variable they cannot access, and no single Markov generator on $\{a, b\}$ reproduces both.

Observe finally the sense in which this is only a *negative* result, per Remark 9.5: it tells us no Markov law on T works universally. It does not tell us what *does* govern the projected process — that depends on the within-block law, which is precisely the discarded information.

10. PHYSICAL INSTANTIATION AND ITS BOUNDARY

10.1. What a physical model must supply. Everything above is a theorem about a supplied structure. The structure becomes a physical model when the following objects are specified *independently* — that is, on physical grounds internal to the sector being modelled, and not reverse-engineered to produce the conclusion:

$$(S, p, \mathcal{R}, \Phi_t) \quad \text{or, in the statistical formulation,} \quad (S, p, \mathcal{R}, \{\mu_s\}).$$

Concretely, this means specifying: a raw realization space Ω ; a physical equivalence $\simeq_{\text{phys}}$ with a justification for calling exactly those differences redundant; a readout Π_{T_1} with a justification for calling exactly that the observable tier; a family of invariant relations; and a dynamics or an outcome-law assignment.

If the model additionally wishes to invoke either identifiability ceiling, it must discharge the corresponding hypothesis, and these are two hypotheses, not one (Remark 8.5). For the deterministic ceiling it must establish Assumption 7.1: that the complete input of *every* admissible probe factors through p . For the statistical ceiling it must establish (30): that the *complete outcome law* of the experiment factors through p . Neither is established by calling the observer “internal”, and neither follows from the other.

10.2. What counts as a witness. We restate, in its strongest form, the standard we have been holding to. The decisive witness for Theorem 4.5 is *not*:

- a map merely *named* a projection;
- noninjectivity manufactured by arbitrarily restricting the descriptor (any map becomes noninjective if one deletes enough of its output; that is a fact about the deletion, not about the world);
- a nontrivial fiber whose points are related by a redundancy that the model failed to quotient (this is gauge, and (1) exists precisely to exclude it);
- a fiber distinction that no target, dynamics, or outcome law is sensitive to (this is an idle refinement, and by Lemma 3.2 everything of interest descends anyway).

It *is*: a pair of admissible, gauge-inequivalent source states with the *same physically justified readout* and a difference that changes *a target, a future response, or a complete outcome law*. That is exactly the hypothesis (13), and each of its three clauses excludes one of the four failures above.

10.3. The alternatives the framework generates. Given an instantiated model, the theorems above become a determinate list of things to compute or measure.

- (1) *Static descent* (Lemma 3.2): is each proposed source relation constant on readout fibers?
- (2) *Minimal completion* (Definition 2.6, Theorem 4.6): construct $E_A = \text{Im}(p, R_A)$ and identify what a sufficient extension must contain.
- (3) *Dynamical closure* (Theorem 5.2): does the source evolution preserve readout fibers? In the linear case, test $U(t) \ker K \subseteq \ker K$, or infinitesimally $PLQ^\perp = 0$.

- (4) *Statistical closure* (Corollary 8.4, Theorem 9.3): do the projected transition rates or outcome laws depend only on the readout state? For a finite chain this is the checkable block-sum condition (33).
- (5) *Discrimination* (Theorem 6.2, Proposition 6.4): add probes and determine whether the joint fibers shrink to physical equality.

Two consequences of this list are worth naming, and both must be stated with restraint.

First, failure of lumpability (33), or observation of history-dependent response at fixed present readout state, *rules out* the universal Markov or autonomous sufficiency of the *chosen instantaneous state description*, under the stated model assumptions. This is a genuine empirical inference, and it is the most accessible test the framework offers — but its conclusion is a *rejection*, not a construction. It does *not* follow, from data alone, that a hidden source of any particular kind exists, still less that the proposed one does. The same observations are compatible with an incorrectly chosen state space (the readout was simply the wrong variable, and a different one closes); with genuine time dependence of the dynamics; with measurement or preparation effects; and with model misspecification of any other sort. What is licensed is exactly this: *the proposed instantaneous closure is false*. Which of the competing explanations for its failure is correct is a further question, to be settled by further modelling and further data, and Theorem 4.5 does not settle it.

Second, a source model earns real physical force only when it *predicts the form and magnitude* of the missing-state contribution — the second term of (24), and where applicable the memory kernel — rather than merely permitting one. A framework that can accommodate any residue predicts nothing.

10.4. The boundary of the result. The non-source-projection theorem is a model-relative mathematical result. Once a reduced source S , a justified physical equivalence relation $\simeq_{\text{phys}}$, an exact readout p , a target Q , and a witness pair satisfying (13) are supplied, the conclusions (N1)–(N6) follow: T_1 is a non-source projection relative to that model and target. This is Theorem 4.5, and it is proved — not conjectural, not heuristic, and not contingent on anything outside the stated hypotheses.

Applying the theorem to nature requires more. A concrete physical theory must supply Ω , $\simeq_{\text{phys}}$, Π_{T_1} , and a witness pair, with physical justification for each, and must establish that the witness distinction is physically substantive rather than an artifact of description. That work is not performed here for any theory; the theorems above constrain it but do not carry it out. The distinctive predictions of an instantiated model — residual memory terms, failure of lumpability, history dependence at fixed readout, or the shrinking of joint probe fibers under a new probe — then require empirical testing against data and preference to alternatives, which is likewise not performed here.

Neither the abstract theorem nor a successful physical instantiation entails an unrestricted ontological conclusion about reality or the universe. Even a fully instantiated and empirically successful model licenses only the relativized statement that, relative to the physically justified source model, readout, and target, the readout is a non-source projection; the move from a successful model to an unrestricted claim about what exists is not a theorem of this paper, and we do not make it. In particular, we do not claim that our universe has been proved to instantiate the hypotheses of Theorem 4.5, and no result above should be read as such a claim.

11. SYNTHESIS

The architecture, assembled, reads as follows. Each line is a separate theorem, and each contributes something the others do not.

physical reduction	$\Omega \longrightarrow S = \Omega / \simeq_{\text{phys}}$	(Def. 2.1)
exact readout	$S \xrightarrow{p} T_1 \cong S / \sim_p$	(Lem. 2.3)
descent test	$r = \bar{r} \circ p \iff r \text{ fiber-constant}$	(Lem. 3.2)
completeness	$(C1) \Leftrightarrow \dots \Leftrightarrow (C5)$	(Thm. 3.4)
target obstruction	$p(s_0) = p(s_1), a_Q(s_0) \neq a_Q(s_1)$	(Thm. 4.5)
minimal completion	$E_Q = \text{Im}(p, a_Q), \text{ terminal}$	(Thm. 4.6)
geometric realization	$E_A^{\text{geo}} = \text{Im}(p_{\text{geo}}, J_A)$	(Ex. 2.8)
observable quotient	$\mathcal{H}_{\text{src}} / \ker K \cong \text{Ran } K$	(Ex. 2.9)
dynamical descent	$p\Phi_t = U_t p \iff \text{fibers preserved}$	(Thm. 5.2)
linear form	$U(t) \ker K \subseteq \ker K, \quad PLQ^\perp = 0$	(Prop. 5.4)
identifiability	$\bigcap_i \sim_{P_i} = \Delta_S$	(Thm. 6.2)
internal ceiling	$P_i = \bar{P}_i \circ p \implies \sim_p \subseteq \sim_I$	(Thm. 7.2)
statistical ceiling	$\mu_s = \mathbf{M}_{p(s)} \implies \ \mu_{s_0} - \mu_{s_1}\ _{\text{TV}} = 0$	(Cor. 8.4)
Markov projection	$\sum_{j \in B_b} g_{ij} \text{ independent of } i \in B_a$	(Thm. 9.3)

(38)

Read as a whole, (38) makes four points that no single line makes alone.

The source/readout distinction is not inferred from the word “projection.” It is established, for a specified model, by a physically reduced nontrivial fiber carrying an essential distinction — and by nothing less.

Completion is not arbitrary extra structure. E_Q is the terminal object among exact extensions sufficient for Q (Theorem 4.6): it is forced, not chosen, and every adequate extension surjects onto it.

Geometric realization is a restriction, not an automatic step. In a geometrically realized source, the completion takes the form $E_A^{\text{geo}} = \text{Im}(p_{\text{geo}}, J_A)$ of (7), whose elements are exactly the completed signatures *attained by compatible global geometric fields*. Not every abstract signature is so realized: field equations, global compatibility, and topological obstructions restrict the image. Which signatures survive is a geometric question, answered by geometry and not by the quotient formalism.

The three failures are ordered, not independent. They stand in the exact relations of Proposition 8.6. Fiber-breaking dynamical nonclosure *requires* two distinct source states sharing a readout, so it *implies* kinematic incompleteness: noninjectivity of p is necessary for it. The converse fails — Example 5.5 and Example 9.6(i) are kinematically incomplete yet dynamically closed — so noninjectivity is not sufficient, and, equivalently, autonomy of the projected law does not imply injectivity of p . Statistical nonidentifiability, by contrast, has no unconditional relation to p in either direction; it acquires one only under the outcome-law factorization hypothesis (30), and under that hypothesis noninjectivity of p implies statistical nonidentifiability within at least one fiber. The one moral that survives all three cases is negative: autonomy of the projected law is evidence of *fiber preservation*, never of *fiber triviality*.

The ceiling is exactly as wide as the input restriction. Theorem 7.2 bites only under Assumption 7.1, and a genuinely new probe — one not factoring through p — lifts it. Source

identification is possible exactly when the joint probe family separates physical source states, and not otherwise.

12. CONCLUSION

We have converted a structural question — is the observable tier the physical state space, or a faithful quotient of a larger one? — into a sequence of exact theorems, and we have marked the boundary at which the mathematics stops.

The strongest justified result is this. *In any model that supplies a source domain reduced by its own declared physical equivalence, an exact readout, and a witness pair of gauge-inequivalent source states with the same readout but different correct answers to a nominated target, the readout is — relative to that model and that target — a non-source projection: no deterministic function of the readout answers the target, the minimal completion properly refines the readout, and every sufficient extension surjects onto it.* Two ceilings then follow, from two separate hypotheses, and we have been careful to derive neither from the other. If every admissible internal probe factors through the readout, then within a fiber all probe values coincide, and no deterministic selector or post-processing of those values can recover the distinction. If, separately, the *complete outcome law* factors through the readout, then within a fiber the data laws are equal, and the stronger conclusion follows: no statistic, no randomized test, and no sample size can discriminate. If instead some further probe fails to factor through the readout — or some experiment’s law depends on the fiber coordinate — the distinction becomes, in principle, empirically discriminable, and Theorem 6.2 says precisely when.

What remains outside is equally definite. The mathematics does not supply the source domain, does not certify the physical equivalence, does not adjudicate which map deserves to be called the readout, and does not exhibit a witness. Those are the tasks of physical instantiation and empirical testing, as delimited in §10.4. We prove the model-relative theorem, sharply; we do not draw the unrestricted ontological conclusion, and no argument above supports it. The value of a boundary theorem lies in both of its sides.

REFERENCES

- [1] F. Ball and G. F. Yeo, *Lumpability and marginalisability for continuous-time Markov chains*, J. Appl. Probab. **30** (1993), 518–528.
- [2] D. Blackwell, *Comparison of experiments*, in: Proc. Second Berkeley Symp. on Math. Statist. and Prob., Univ. of California Press, Berkeley, 1951, pp. 93–102.
- [3] P. Buchholz, *Exact and ordinary lumpability in finite Markov chains*, J. Appl. Probab. **31** (1994), 59–75.
- [4] T. tom Dieck, *Transformation Groups*, de Gruyter Studies in Mathematics 8, Walter de Gruyter, Berlin, 1987.
- [5] K.-J. Engel and R. Nagel, *One-Parameter Semigroups for Linear Evolution Equations*, Graduate Texts in Mathematics 194, Springer, New York, 2000.
- [6] K. Gödel, *Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I*, Monatsh. Math. Phys. **38** (1931), 173–198.
- [7] R. E. Kalman, *Mathematical description of linear dynamical systems*, J. SIAM Control Ser. A **1** (1963), 152–192.
- [8] J. G. Kemeny and J. L. Snell, *Finite Markov Chains*, Van Nostrand, Princeton, 1960.
- [9] L. Le Cam, *Sufficiency and approximate sufficiency*, Ann. Math. Statist. **35** (1964), 1419–1455.
- [10] E. L. Lehmann and J. P. Romano, *Testing Statistical Hypotheses*, 3rd ed., Springer, New York, 2005.
- [11] S. Mac Lane, *Categories for the Working Mathematician*, 2nd ed., Graduate Texts in Mathematics 5, Springer, New York, 1998.
- [12] H. Mori, *Transport, collective motion, and Brownian motion*, Progr. Theoret. Phys. **33** (1965), 423–455.
- [13] S. Nakajima, *On quantum theory of transport phenomena: steady diffusion*, Progr. Theoret. Phys. **20** (1958), 948–959.

- [14] K. Yosida, *Functional Analysis*, 6th ed., Grundlehren der mathematischen Wissenschaften 123, Springer, Berlin, 1980.
- [15] R. Zwanzig, *Ensemble method in the theory of irreversibility*, J. Chem. Phys. **33** (1960), 1338–1341.
- [16] J. Rodgers, *Source–Readout Non-Equivalence: Descent and Equivariant Reconstruction Obstructions*, Zenodo, 2026. doi:10.5281/zenodo.21369967.
- [17] J. Rodgers, *Target-Relative Necessity of Completion: When Readout Loss Obstructs, and What a Sufficient Extension Must Retain*, Zenodo, 2026. doi:10.5281/zenodo.21370501.
- [18] J. Rodgers, *Canonical Minimal Source Completion: The Coarsest Readout Extension on which a Nominated Family of Source Relations Becomes Well Defined*, Zenodo, 2026. doi:10.5281/zenodo.21370763.
- [19] J. Rodgers, *Geometric Realization of Completed Source Relations: Descent, Orbit Spaces, Invariant Relations, and Variational Response in Shadow Theory*, Zenodo, 2026. doi:10.5281/zenodo.21370985.
- [20] J. Rodgers, *Observable Quotients and Exact Projected Dynamics: Closure, Memory, Minimal Dynamical Completion, and Effective Field Operators*, Zenodo, 2026. doi:10.5281/zenodo.21371251.

INDEPENDENT RESEARCHER