

OBSERVABLE QUOTIENTS AND EXACT PROJECTED DYNAMICS: CLOSURE, MEMORY, MINIMAL DYNAMICAL COMPLETION, AND EFFECTIVE FIELD OPERATORS

JEREMY RODGERS

ABSTRACT. We develop the observable and dynamical layer of Shadow Theory. A bounded observation operator $K : \mathcal{H}_{\text{src}} \rightarrow \mathcal{H}_{\text{obs}}$ identifies source states differing by $\ker K$ and produces an exact observable quotient $\mathcal{H}_K = \mathcal{H}_{\text{src}} / \ker K$, canonically and boundedly bijective onto $\text{Ran } K$. We prove that this bijection has bounded inverse in the norm inherited from $\mathcal{H}_{\text{obs}}$ exactly when $\text{Ran } K$ is closed, and we keep sharply apart three distinct phenomena: exact quotient loss ($\ker K \neq \{0\}$), controlled low-rank approximation (singular-value truncation, with the operator-norm optimality $\|K - K_r\|_{\text{op}} = \sigma_{r+1}$), and unstable inversion (failure of K to be bounded below on $(\ker K)^\perp$, equivalently nonclosed range). Compactness is exactly what supplies the controlled finite-rank approximation, but by itself it implies neither a nontrivial kernel nor unstable inversion.

For a strongly continuous source semigroup $U(t)$ we prove that an induced evolution on $\text{Ran } K$ exists if and only if $U(t) \ker K \subseteq \ker K$ for all $t \geq 0$; failure produces initially indistinguishable source states with later distinct readouts, so that no single-valued autonomous law on the instantaneous observable state can reproduce all source trajectories. For an orthogonal resolved/unresolved splitting with bounded generator L we derive the exact projected equation

$$\dot{x}(t) = PLP x(t) + PLQ e^{tQLQ} z(0) + \int_0^t PLQ e^{(t-s)QLQ} QLP x(s) ds,$$

whose second and third terms are, respectively, a deterministic unresolved-initial-state term and an exact memory kernel; no stochastic, Markovian, or timescale approximation is made anywhere. Autonomous first-order closure valid for *every* source initial state is proved equivalent to $PLQ = 0$, i.e. to invariance of $\ker P$ under L .

We then construct the minimal dynamical completion: the unresolved response map $z_0 \mapsto PLQ e^{QLQ} z_0$ descends to an injective map on the Hilbert quotient $E_{\text{dyn}} = \mathcal{H}_Q / N_Q$, which is universal among surjective extensions determining all resolved responses, and which in finite dimensions is computed by a Kalman-type observability matrix. A fiberwise criterion for smooth surjective submersions gives the nonlinear analogue. Finally, eliminating a genuine source sector from a bounded quadratic action yields Schur-complement effective operators $A - BC^{-1}B^*$, whose correction can be nonlocal when the composed operator $BC^{-1}B^*$ retains nonlocal Green-operator directions; and stationary elimination in a diffeomorphism-invariant action yields a covariant effective stress-energy defined as a variational response, together with total—not sectorwise—covariant conservation. Throughout, the results are necessary-and-sufficient criteria and exact reduction identities for models in which a source/readout architecture is declared; no empirical claim is made.

Date: June 2026.

2020 Mathematics Subject Classification. Primary 47A05, 47D06; Secondary 47B07, 93B07, 82C31, 83C99, 47A62.

Key words and phrases. Observable quotient, bounded observation operator, quotient Hilbert space, closed range, singular value decomposition, Eckart–Young–Mirsky theorem, C_0 -semigroup, invariant subspace, Mori–Zwanzig projection, memory kernel, autonomous closure, observability, submersion, projectable vector field, Schur complement, effective action, variational stress-energy.

Website: everythingequation.com.

1. INTRODUCTION

1.1. The problem. Suppose a physical model is specified on a state space $\mathcal{H}_{\text{src}}$ (the *source*), while every measurement one can actually perform is the image of a fixed bounded linear map $K : \mathcal{H}_{\text{src}} \rightarrow \mathcal{H}_{\text{obs}}$ (the *readout*). Two questions then arise, and they are not the same question.

- (Q1) *Statics.* Which source quantities are functions of the readout alone?
- (Q2) *Dynamics.* If the source evolves, does the readout evolve autonomously—i.e. is there a law of motion on the observable state alone that reproduces the readout of every source trajectory?

An affirmative answer to (Q1) for a given functional is a purely kinematic statement about the fibers of K . An affirmative answer to (Q2) is a compatibility statement between the fibers of K and the flow. The central point of this paper is that (Q2) can fail while the quotient of (Q1) remains a perfectly coherent, complete Hilbert space with a perfectly well-defined observable state. A theory can be an *exact quotient theory* and nevertheless fail to be an *autonomous presentation* of its own observable degrees of freedom. When that happens, the exact projection of the source dynamics is not a first-order autonomous law but an integro-differential law carrying two extra structures: dependence on unresolved initial data, and dependence on the resolved history through a memory kernel.

1.2. Four distinctions that must not be merged. Much informal discussion of “information loss” under coarse-graining conflates phenomena that are mathematically independent. We insist on four:

- (D1) **Exact nonidentifiability.** $\ker K \neq \{0\}$. Distinct source states are *identically* equal in readout, for all time-zero purposes. No amount of numerical precision recovers them.
- (D2) **Unstable inversion.** K may be injective yet have unbounded inverse on its range. In general this is the failure of K to be bounded below on $(\ker K)^\perp$, equivalently the failure of $\text{Ran } K$ to be closed (Theorem 2.3). *Under the additional hypothesis that K is compact* with infinite-dimensional range, it takes the familiar form that the singular values tend to 0; but singular values are available only in the compact setting, and the general criterion is the one just stated. Here the source state *is* uniquely determined by the readout, but not continuously.
- (D3) **Approximation error.** If K is compact one may truncate its singular system at rank r and incur exactly σ_{r+1} in operator norm. This is a controlled, quantified deficit, and it exists even when K is injective and even when its inverse is bounded.
- (D4) **Failure of autonomous dynamics.** $U(t)\ker K \not\subseteq \ker K$. This is a genuinely dynamical obstruction, not determined by the static facts (D1)–(D3).

The logical relation between (D1) and (D4) must be stated with care, and we do so once here and abide by it throughout. Since $\{0\}$ is invariant under every linear map, $\ker K = \{0\}$ forces $U(t)\ker K \subseteq \ker K$ trivially. Hence

$$(D4) \text{ holds} \implies (D1) \text{ holds:} \quad U(t_0)\ker K \not\subseteq \ker K \implies \ker K \neq \{0\}. \quad (1.1)$$

A *nontrivial kernel is therefore necessary for failure of descent*. It is not sufficient, and this is the whole point: a nontrivial kernel may be invariant under the source semigroup, in which case dynamics descend exactly and the observable theory is autonomous despite losing exact source distinctions; or it may fail to be invariant, in which case descent fails. The static fact (D1) does not decide which. So (D4) is an *additional* dynamical obstruction, logically *stronger* than (D1) and not deducible from it—but it is not logically independent of it. (D2) and (D3),

by contrast, are genuinely independent of both: Example 4.5 exhibits (D2) and (D3) with (D1), and hence also (D4), failing.

Section 2 isolates (D1) and (D2), Section 4 isolates (D3), and Sections 5–6 isolate (D4). The compact injective diagonal operator of Example 4.5 exists precisely to show that (D1), (D2), (D3) can be dissociated; Example 6.6 exists to show (D4) in a completely computed finite-dimensional block—where, in accordance with (1.1), the kernel $\ker P = \mathcal{H}_Q$ is indeed nontrivial, and the content of the failure is that it is not L -invariant.

1.3. Contributions.

- **Exact quotient (§2).** $\tilde{K} : \mathcal{H}_K \rightarrow \text{Ran } K$ is a bounded linear bijection; $\tilde{K}^{-1}$ is bounded *in the norm inherited from $\mathcal{H}_{\text{obs}}$* if and only if $\text{Ran } K$ is closed (Theorem 2.3). We are careful never to identify $\mathcal{H}_K$ isometrically with $\text{Ran } K$ carrying the inherited norm: what is always true is that $\tilde{K}$ transports the quotient norm to $\text{Ran } K$, and the transported norm agrees topologically with the inherited one exactly in the closed-range case.
- **Descent of functionals (§3).** A factors through K iff A is constant on cosets of $\ker K$ (Proposition 3.2). For continuous linear functionals this is $A|_{\ker K} = 0$, and $(\ker K)^\perp = \overline{\text{Ran } K^*}$. The closure is essential: an observable linear functional need not be K^* applied to any *exact* observable probe when $\text{Ran } K^*$ fails to be closed (Remark 3.5).
- **Optimal reduction (§4).** Schmidt’s singular expansion and the Eckart–Young–Mirsky operator-norm bound $\|K - K_r\|_{\text{op}} = \sigma_{r+1}$ (Theorem 4.4), together with the compact injective example (Example 4.5).
- **Semigroup descent (§5).** $KU(t) = V(t)K$ for a well-defined $V(t)$ on $\text{Ran } K$ iff $U(t)N \subseteq N$, $N = \ker K$ (Theorem 5.2), with an explicit account of which topology the induced semigroup is strongly continuous in (Proposition 5.3) and the failure corollary (Corollary 5.4).
- **Exact projected dynamics (§6).** The Mori–Zwanzig-type identity (Theorem 6.2), obtained by variation of constants—exactly, with no ensemble, no random force, no Markov limit—and the closure criterion $PLQ = 0$ *quantified over all source states* (Theorem 6.4).
- **Minimal dynamical completion (§7).** The Hilbert quotient $E_{\text{dyn}} = \mathcal{H}_Q/N_Q$ is the coarsest extension of the resolved state determining every resolved response function (Theorem 7.4), it closes the resolved trajectory through a Volterra equation (Theorem 7.6), and in finite dimensions N_Q is a finite intersection of kernels computable by Cayley–Hamilton (Proposition 7.7), linking the construction to Kalman observability.
- **Nonlinear projectability (§8).** For a surjective submersion $p : S \rightarrow T$, a projected field Y with $Dp \circ X = Y \circ p$ exists iff $Dp_s X_s$ is constant along fibers (Theorem 8.2); this is the nonlinear analogue of $PLQ = 0$.
- **Elimination (§9–§10).** Schur-complement effective operators for bounded quadratic actions with boundedly invertible C (Theorem 9.2), and stationary elimination in a diffeomorphism-invariant action producing a covariant effective stress-energy with total covariant conservation (Proposition 10.2, Theorem 10.5).

1.4. What is established mathematics, and what is new here. The component theories on which this paper draws are established mathematics, and we claim none of them. Closed-range theory and the bounded inverse theorem are classical functional analysis [1, 13]; the singular system and the operator-norm optimality of rank- r truncation are due to Schmidt, Eckart–Young, and Mirsky [14, 2, 9]; the theory of C_0 -semigroups and of invariant subspaces under them is standard [3, 11]; the projection identity of Theorem 6.2 is the Mori–Zwanzig

identity [20, 10, 21]; the finite rank criterion of Proposition 7.7 is Kalman observability [6, 15]; the Schur complement and block-operator elimination are classical [19]; and Noether identities and the variational definition of stress-energy are textbook general relativity [16]. Where these results appear below they are proved for completeness and stated in the form needed, not asserted as new.

What this paper contributes is their rigorous integration into a single source–readout architecture, and the results that only that integration makes visible. Specifically: the exact descent/closure obstruction—that source dynamics descend to the observable quotient if and only if the readout kernel is invariant (Theorem 5.2), equivalently, in the projected setting, if and only if $PLQ = 0$ *quantified over all source states* (Theorem 6.4)—together with the sharp logical relation this obstruction bears to the static phenomena (D1)–(D3), which is neither independence nor implication in the direction usually assumed (1.1); the *target-relative minimal dynamical completion* $E_{\text{dyn}} = \mathcal{H}_Q/N_Q$, characterized by a universal property among all sufficient extensions (Theorem 7.4) and shown to close the resolved trajectory (Theorem 7.6), whose minimality is relative to a declared target and changes when the target changes (Remark 7.5); the accompanying separation of algebraic descent from stable readout (Theorem 3.4); and the unified statement of the relation among unresolved initial data, memory, and completion. On that last point the precise claim is the following, and it is deliberately weaker than the slogan it replaces. The vanishing $PLQ = 0$ makes the initial-state term, the memory kernel, and the minimal completion all trivial simultaneously, and is equivalent to autonomous closure for every source state (Theorem 6.4). If $PLQ \neq 0$, then some unresolved initial data do affect the resolved evolution and E_{dyn} is nontrivial. The memory kernel $\mathcal{K}(\tau) = PLQ e^{\tau QLQ} QLP$, however, can vanish identically even when $PLQ \neq 0$ —most simply when $QLP = 0$, so that the resolved sector never feeds back into the unresolved one, but also through cancellation in the composition. Initial-state dependence, memory, and retained completion are therefore three structurally distinct phenomena which need not all be nonzero at once; the exact projection identity is what exhibits them as separate, and the completion E_{dyn} is the precise price of restoring an initial-value problem when the resolved state alone will not close. The field-theoretic sections then carry the same architecture into elimination: Schur-complement effective operators and a variational, covariant effective stress-energy, with their scope and their failure modes marked.

1.5. What is *not* claimed. We prove no empirical proposition. We do not claim that the physical universe possesses hidden source variables, memory, or any particular effective correction. We supply necessary-and-sufficient criteria and exact reduction identities for models *in which a source/readout architecture has been declared*. A formal memory kernel, a compact approximation, a Schur complement, or a named effective tensor is, by itself, a piece of mathematics and not a physical discovery; its physical content, if any, lies entirely in the model that supplies K , L , and the action. We also do not claim that the memory term is stochastic. It is not: at the level of exactness maintained here, both η_{z_0} and $\mathcal{K}$ are deterministic linear objects. They acquire a probabilistic reading only after an *additional*, independent, probabilistic model of $z(0)$ is imposed, which we do not impose.

1.6. Relation to the companion papers. This paper is the dynamical layer of a seven-paper foundational sequence. Its companions in that sequence are the preceding companion papers, cited as [22, 23, 24, 25]: descent and equivariant reconstruction obstructions [22]; target-relative necessity of completion [23]; canonical minimal source completion [24]; and geometric realization with variational response [25]. The present paper is self-contained: every definition it needs is restated locally, and every theorem it states is proved here, including the minimal-completion

theorem, which we prove directly (Theorem 7.4) rather than importing. Cross-references to [22, 23, 24, 25] are for context only.

1.7. Notation and standing conventions. All Hilbert spaces are complex unless stated otherwise, with inner products $\langle \cdot, \cdot \rangle$ linear in the second argument. $\mathcal{B}(\mathcal{H}, \mathcal{K})$ denotes bounded linear operators. For $T \in \mathcal{B}(\mathcal{H}, \mathcal{K})$ we write $\ker T$, $\text{Ran } T$, T^* , and $\|T\|_{\text{op}}$. Orthogonal projections are denoted P , $Q = I - P$, with ranges $\mathcal{H}_P = P\mathcal{H}$ and $\mathcal{H}_Q = Q\mathcal{H}$; both are closed, and $\mathcal{H} = \mathcal{H}_P \oplus \mathcal{H}_Q$. We freely regard PLQ as an operator $\mathcal{H}_Q \rightarrow \mathcal{H}_P$, QLQ as an operator $\mathcal{H}_Q \rightarrow \mathcal{H}_Q$, and so on, the identification being the obvious restriction/corestriction. Semigroups are one-parameter families $U(t)$, $t \geq 0$, of bounded operators with $U(0) = I$, $U(t+s) = U(t)U(s)$, and $t \mapsto U(t)\Psi$ continuous for each Ψ (the C_0 property) [3, 11]. In Sections 6–7 the generator L is *bounded* and $U(t) = e^{tL}$, which makes all exponentials norm-convergent power series and all manipulations below elementary; unbounded generators are addressed only in explicitly flagged remarks, and never silently.

2. THE EXACT OBSERVABLE QUOTIENT

Assumption 2.1. $\mathcal{H}_{\text{src}}$ and $\mathcal{H}_{\text{obs}}$ are Hilbert spaces and $K \in \mathcal{B}(\mathcal{H}_{\text{src}}, \mathcal{H}_{\text{obs}})$.

Definition 2.2 (Observational equivalence). For $\Psi, \Psi' \in \mathcal{H}_{\text{src}}$ set

$$\Psi \sim_K \Psi' \quad :\Longleftrightarrow \quad K\Psi = K\Psi' \quad \Longleftrightarrow \quad \Psi - \Psi' \in \ker K. \quad (2.1)$$

Since K is continuous, $N := \ker K$ is a closed subspace of $\mathcal{H}_{\text{src}}$. The quotient vector space

$$\mathcal{H}_K := \mathcal{H}_{\text{src}} / \ker K \quad (2.2)$$

therefore carries the quotient norm

$$\|[\Psi]\|_K := \inf_{n \in \ker K} \|\Psi + n\|_{\mathcal{H}_{\text{src}}} = \text{dist}(\Psi, \ker K), \quad (2.3)$$

and $(\mathcal{H}_K, \|\cdot\|_K)$ is a Hilbert space: the canonical map $(\ker K)^\perp \rightarrow \mathcal{H}_K$, $\Psi \mapsto [\Psi]$, is a linear bijection, and it is isometric because the infimum in (2.3) is attained at the orthogonal projection of Ψ onto $(\ker K)^\perp$. Thus the quotient norm is induced by the inner product $\langle [\Psi], [\Psi'] \rangle_K := \langle \Pi\Psi, \Pi\Psi' \rangle_{\mathcal{H}_{\text{src}}}$, where Π is the orthogonal projection onto $(\ker K)^\perp$, and completeness is inherited from that of the closed subspace $(\ker K)^\perp$.

The *exact observable state* associated with a source state Ψ is the coset

$$[\Psi]_K = \Psi + \ker K. \quad (2.4)$$

It is essential to the logic of what follows that the quotient distinguishes no representative within this coset. The orthogonal representative $\Pi\Psi$ is a choice made by the ambient inner product of $\mathcal{H}_{\text{src}}$, not by the readout.

Theorem 2.3 (Observable quotient theorem). *Under Assumption 2.1, the map*

$$\tilde{K} : \mathcal{H}_K \longrightarrow \text{Ran } K, \quad \tilde{K}[\Psi] := K\Psi, \quad (2.5)$$

is well defined, linear, bijective, and bounded, with $\|\tilde{K}\|_{\text{op}} = \|K\|_{\text{op}}$, where $\text{Ran } K$ carries the norm inherited from $\mathcal{H}_{\text{obs}}$. Moreover the following are equivalent:

- (1) $\tilde{K}^{-1} : \text{Ran } K \rightarrow \mathcal{H}_K$ is bounded when $\text{Ran } K$ carries the inherited norm;
- (2) $\text{Ran } K$ is closed in $\mathcal{H}_{\text{obs}}$;
- (3) there exists $c > 0$ with $\|K\Psi\|_{\mathcal{H}_{\text{obs}}} \geq c \text{dist}(\Psi, \ker K)$ for all $\Psi \in \mathcal{H}_{\text{src}}$.

Proof. Well-definedness and bijectivity. If $[\Psi] = [\Psi']$ then $\Psi - \Psi' \in \ker K$, so $K\Psi = K\Psi'$; hence (2.5) does not depend on the representative. Linearity is immediate. Surjectivity onto $\text{Ran } K$ holds by definition of the range. If $\tilde{K}[\Psi] = 0$ then $K\Psi = 0$, i.e. $\Psi \in \ker K$, i.e. $[\Psi] = 0$; hence $\tilde{K}$ is injective.

Boundedness. For every $n \in \ker K$,

$$\|\tilde{K}[\Psi]\|_{\mathcal{H}_{\text{obs}}} = \|K\Psi\| = \|K(\Psi + n)\| \leq \|K\|_{\text{op}} \|\Psi + n\|_{\mathcal{H}_{\text{src}}}.$$

Taking the infimum over $n \in \ker K$ gives $\|\tilde{K}[\Psi]\| \leq \|K\|_{\text{op}} \|\Psi\|_K$, so $\|\tilde{K}\|_{\text{op}} \leq \|K\|_{\text{op}}$. Conversely $\|K\Psi\| = \|\tilde{K}[\Psi]\| \leq \|\tilde{K}\|_{\text{op}} \|\Psi\|_K \leq \|\tilde{K}\|_{\text{op}} \|\Psi\|$, so the two operator norms coincide.

(1) $\Rightarrow$ (2). Suppose $\|\tilde{K}^{-1}\|_{\text{op}} =: M < \infty$. Then $\tilde{K}$ is a linear homeomorphism from the complete space $\mathcal{H}_K$ onto $\text{Ran } K$ with its inherited norm; hence $(\text{Ran } K, \|\cdot\|_{\mathcal{H}_{\text{obs}}})$ is complete. A complete subspace of a normed space is closed, so $\text{Ran } K$ is closed in $\mathcal{H}_{\text{obs}}$.

(2) $\Rightarrow$ (1). If $\text{Ran } K$ is closed in the Hilbert space $\mathcal{H}_{\text{obs}}$, then $(\text{Ran } K, \|\cdot\|_{\mathcal{H}_{\text{obs}}})$ is itself a Banach space. Since $\tilde{K}$ is a bounded linear bijection between the Banach spaces $\mathcal{H}_K$ and $\text{Ran } K$, the bounded inverse theorem [1, Cor. 2.7] yields boundedness of $\tilde{K}^{-1}$.

(1) $\Leftrightarrow$ (3). Statement (1) says exactly that there is $M < \infty$ with $\|[\Psi]\|_K \leq M\|K\Psi\|$ for all Ψ ; by (2.3) this is (3) with $c = 1/M$. $\square$

Remark 2.4 (The two norms on $\text{Ran } K$). Theorem 2.3 always allows us to *transport* the quotient norm to $\text{Ran } K$ by declaring $\|K\Psi\|_{\sharp} := \|[\Psi]\|_K$; with this transported norm $\tilde{K}$ is by construction an isometric isomorphism, and $(\text{Ran } K, \|\cdot\|_{\sharp})$ is a Hilbert space. What is *not* generally true is that $\|\cdot\|_{\sharp}$ coincides with, or is even equivalent to, the norm inherited from $\mathcal{H}_{\text{obs}}$. The inequality $\|\cdot\|_{\mathcal{H}_{\text{obs}}} \leq \|K\|_{\text{op}} \|\cdot\|_{\sharp}$ always holds on $\text{Ran } K$; the reverse inequality holds, with some constant, exactly in the closed-range case. We therefore never write $\mathcal{H}_K \cong \text{Ran } K$ without saying which norm $\text{Ran } K$ carries. Every later statement about topologies on the observable side is subject to this distinction.

Corollary 2.5 (Exact loss criterion). *The observation map K loses exact source distinctions—i.e. there exist $\Psi \neq \Psi'$ with $\Psi \sim_K \Psi'$ —if and only if $\ker K \neq \{0\}$. If $\ker K = \{0\}$, then K is injective and the source state is uniquely determined by its readout, even if $\tilde{K}^{-1} = K^{-1}$ is unbounded on $\text{Ran } K$.*

Proof. Immediate from (2.1): $\Psi \sim_K \Psi'$ with $\Psi \neq \Psi'$ iff $\Psi - \Psi' \in \ker K \setminus \{0\}$. The second sentence is the content of Theorem 2.3: injectivity is the statement $\ker K = \{0\}$ and is logically independent of closedness of $\text{Ran } K$, which is what governs boundedness of the inverse. $\square$

Corollary 2.5 is the formal separation of (D1) from (D2). Nonidentifiability is a statement about the kernel; instability is a statement about the range. An injective compact operator exhibits the second without the first, as we show in Example 4.5.

3. DESCENT OF OBSERVABLE FUNCTIONALS

Definition 3.1 (Exact observability of a functional). Let V be a set and $A : \mathcal{H}_{\text{src}} \rightarrow V$. We say A is *exactly observable through K* if there exists $\bar{A} : \text{Ran } K \rightarrow V$ with

$$A = \bar{A} \circ K. \tag{3.1}$$

Proposition 3.2 (Functional descent). *A is exactly observable through K if and only if*

$$A(\Psi + n) = A(\Psi) \quad \text{for all } \Psi \in \mathcal{H}_{\text{src}}, n \in \ker K. \tag{3.2}$$

When (3.2) holds, $\bar{A}$ is unique.

Proof. If (3.1) holds and $n \in \ker K$, then $A(\Psi + n) = \bar{A}(K\Psi + Kn) = \bar{A}(K\Psi) = A(\Psi)$, which is (3.2).

Conversely, assume (3.2). Given $\omega \in \text{Ran } K$, choose any Ψ with $K\Psi = \omega$ and set $\bar{A}(\omega) := A(\Psi)$. If $K\Psi = K\Psi' = \omega$ then $\Psi' = \Psi + n$ with $n \in \ker K$, so (3.2) gives $A(\Psi') = A(\Psi)$; the definition is therefore independent of the chosen preimage, and (3.1) holds by construction. Uniqueness is forced because K is onto $\text{Ran } K$: any $\bar{A}'$ satisfying (3.1) agrees with $\bar{A}$ at every $\omega = K\Psi$. $\square$

Proposition 3.2 is the fiberwise descent statement of [22] specialized to the affine fibers $\Psi + \ker K$; we have proved it here so that no external result is load-bearing.

Corollary 3.3 (Linear descent and the adjoint). *Let $A \in \mathcal{H}_{\text{src}}^*$ be a continuous linear functional. Then A descends through K , in the sense of Definition 3.1, if and only if*

$$A|_{\ker K} = 0. \quad (3.3)$$

Identifying $\mathcal{H}_{\text{src}}^* \cong \mathcal{H}_{\text{src}}$ by the Riesz representation $A = \langle a, \cdot \rangle$, condition (3.3) says $a \in (\ker K)^\perp$, and

$$(\ker K)^\perp = \overline{\text{Ran } K^*}. \quad (3.4)$$

Proof. Since A is linear, (3.2) reads $A(n) = 0$ for all $n \in \ker K$, which is (3.3). For (3.4): $\Psi \in \ker K$ iff $K\Psi = 0$ iff $\langle K\Psi, \omega \rangle = 0$ for all $\omega \in \mathcal{H}_{\text{obs}}$ iff $\langle \Psi, K^*\omega \rangle = 0$ for all ω , i.e. iff $\Psi \perp \text{Ran } K^*$. Hence $\ker K = (\text{Ran } K^*)^\perp$, and taking orthogonal complements and using $M^{\perp\perp} = \overline{M}$ for a subspace M gives $(\ker K)^\perp = \overline{\text{Ran } K^*}$. $\square$

It is essential to notice what Definition 3.1 does and does not require. It requires a *factorization*, $A = \bar{A} \circ K$, and nothing more: $\bar{A}$ is a map on the *set* $\text{Ran } K$, and no continuity of $\bar{A}$ with respect to any particular norm on $\text{Ran } K$ has been demanded. Corollary 3.3 therefore characterizes *algebraic* descent. For a continuous linear $A = \langle a, \cdot \rangle$ with $a \in (\ker K)^\perp$, the induced $\bar{A}$ is automatically continuous in the *transported quotient norm* $\|\cdot\|_\#$ of Remark 2.4 (because $\bar{A} \circ \tilde{K}$ is the continuous functional $[\Psi] \mapsto \langle a, \Psi \rangle$ on the Hilbert space $\mathcal{H}_K$). But continuity in the *inherited* observable norm—the norm an experimenter actually has access to on $\mathcal{H}_{\text{obs}}$ —is a strictly stronger demand, and it is governed by a different, sharper condition.

Theorem 3.4 (Stable observability of a linear functional). *Let $A = \langle a, \cdot \rangle \in \mathcal{H}_{\text{src}}^*$ with $a \in (\ker K)^\perp$, so that A descends algebraically, with unique $\bar{A} : \text{Ran } K \rightarrow \mathbb{C}$ satisfying $A = \bar{A} \circ K$. Then*

$$\bar{A} \text{ is bounded on } (\text{Ran } K, \|\cdot\|_{\mathcal{H}_{\text{obs}}}) \iff a \in \text{Ran } K^*, \quad (3.5)$$

i.e. if and only if there exists an exact observable probe $\omega \in \mathcal{H}_{\text{obs}}$ with $a = K^*\omega$. In that case $\bar{A}$ extends to the bounded functional $\langle \omega, \cdot \rangle$ on all of $\mathcal{H}_{\text{obs}}$, and $A(\Psi) = \langle \omega, K\Psi \rangle_{\mathcal{H}_{\text{obs}}}$ for every Ψ : the source quantity is read off the observable state by a bounded measurement.

Proof. ($\Leftarrow$) Suppose $a = K^*\omega$ with $\omega \in \mathcal{H}_{\text{obs}}$. Then for every Ψ ,

$$A(\Psi) = \langle a, \Psi \rangle = \langle K^*\omega, \Psi \rangle = \langle \omega, K\Psi \rangle_{\mathcal{H}_{\text{obs}}},$$

so $\bar{A}(y) = \langle \omega, y \rangle_{\mathcal{H}_{\text{obs}}}$ for $y \in \text{Ran } K$, which satisfies $|\bar{A}(y)| \leq \|\omega\| \|y\|_{\mathcal{H}_{\text{obs}}}$ and is therefore bounded in the inherited norm.

($\Rightarrow$) Suppose $\bar{A}$ is bounded on $(\text{Ran } K, \|\cdot\|_{\mathcal{H}_{\text{obs}}})$, say $|\bar{A}(y)| \leq M\|y\|_{\mathcal{H}_{\text{obs}}}$ for all $y \in \text{Ran } K$. Since $\text{Ran } K$ is a subspace of $\mathcal{H}_{\text{obs}}$, the Hahn–Banach theorem [1, Cor. 1.2] extends $\bar{A}$ to a bounded linear functional Λ on all of $\mathcal{H}_{\text{obs}}$ with the same norm. By the Riesz representation theorem there is $\omega \in \mathcal{H}_{\text{obs}}$ with $\Lambda = \langle \omega, \cdot \rangle$. Then for every $\Psi \in \mathcal{H}_{\text{src}}$,

$$\langle a, \Psi \rangle = A(\Psi) = \bar{A}(K\Psi) = \Lambda(K\Psi) = \langle \omega, K\Psi \rangle = \langle K^*\omega, \Psi \rangle,$$

so $\langle a - K^*\omega, \Psi \rangle = 0$ for all Ψ , i.e. $a = K^*\omega \in \text{Ran } K^*$. $\square$

Remark 3.5 (Algebraic descent versus stable readout: exact preimage and closure). Equation (3.4) contains a closure and one must not delete it, and Theorem 3.4 says exactly what the closure costs. Combining the two:

- $a \in \overline{\text{Ran } K^*} = (\ker K)^\perp \iff A$ descends *algebraically* through K (Corollary 3.3), with $\bar{A}$ continuous in the transported quotient norm;
- $a \in \text{Ran } K^* \iff A$ descends to a functional that is *bounded in the inherited observable norm* (Theorem 3.4).

The two conditions coincide precisely when $\text{Ran } K^*$ is closed, which by the closed range theorem is equivalent to $\text{Ran } K$ being closed [1, Thm. 2.19]—the same dichotomy as in Theorem 2.3. So for a compact K with infinite-dimensional range—for instance Example 4.5— $\text{Ran } K^*$ is dense in $(\ker K)^\perp$ but *proper*, and there exist source functionals which descend perfectly well as maps of the observable state, yet possess no exact representing probe ω with $a = K^*\omega$. Such a functional is a limit of representable ones, but the approximating probes ω_k have $\|\omega_k\| \rightarrow \infty$: reading the quantity off the observable state amplifies observational error without bound. This is (D2) reappearing on the dual side. Accordingly we say that the elements of $(\ker K)^\perp$ are the *observable* linear source directions, and reserve *stably observable* for those in $\text{Ran } K^*$; the phrase should never be read as asserting inherited-readout stability where only algebraic descent has been established.

4. COMPACT OBSERVATION AND CONTROLLED REDUCTION

Assumption 4.1. In this section $\mathcal{H}_{\text{src}}$ and $\mathcal{H}_{\text{obs}}$ are separable and $K \in \mathcal{B}(\mathcal{H}_{\text{src}}, \mathcal{H}_{\text{obs}})$ is compact.

Under Assumption 4.1, K^*K is a positive, self-adjoint, compact operator on $\mathcal{H}_{\text{src}}$. By the spectral theorem for compact self-adjoint operators [13, Ch. VI] its nonzero spectrum consists of positive eigenvalues of finite multiplicity, which can accumulate only at 0. Enumerate them, repeated according to multiplicity and in nonincreasing order, as $\sigma_1^2 \geq \sigma_2^2 \geq \dots > 0$ with $\sigma_n > 0$; the index set is $\{1, \dots, R\}$ if K has finite rank R , and $\mathbb{N}$ otherwise. Let $\{v_n\}$ be a corresponding orthonormal family of eigenvectors of K^*K , so $K^*Kv_n = \sigma_n^2 v_n$, and put

$$u_n := \sigma_n^{-1} K v_n. \quad (4.1)$$

Lemma 4.2 (Schmidt singular system). *The families $\{v_n\}$ and $\{u_n\}$ are orthonormal, $\{v_n\}$ is an orthonormal basis of $(\ker K)^\perp = \overline{\text{Ran } K^*}$, $\{u_n\}$ is an orthonormal basis of $\overline{\text{Ran } K}$, and*

$$K v_n = \sigma_n u_n, \quad K^* u_n = \sigma_n v_n, \quad (4.2)$$

$$K \Psi = \sum_{n \geq 1} \sigma_n \langle v_n, \Psi \rangle u_n \quad (\Psi \in \mathcal{H}_{\text{src}}), \quad (4.3)$$

the series converging in the norm of $\mathcal{H}_{\text{obs}}$. If the index set is infinite then $\sigma_n \rightarrow 0$.

Proof. Orthonormality of $\{u_n\}$: $\langle u_m, u_n \rangle = (\sigma_m \sigma_n)^{-1} \langle K v_m, K v_n \rangle = (\sigma_m \sigma_n)^{-1} \langle v_m, K^* K v_n \rangle = (\sigma_n / \sigma_m) \langle v_m, v_n \rangle = \delta_{mn}$. From (4.1), $K v_n = \sigma_n u_n$; and $K^* u_n = \sigma_n^{-1} K^* K v_n = \sigma_n v_n$, which is (4.2).

Since $\ker K = \ker K^* K$ (indeed $K \Psi = 0$ iff $\langle \Psi, K^* K \Psi \rangle = \|K \Psi\|^2 = 0$), the spectral theorem gives the orthogonal decomposition $\mathcal{H}_{\text{src}} = \ker K \oplus \overline{\text{span}\{v_n\}}$, and by Corollary 3.3 the second summand is $(\ker K)^\perp = \overline{\text{Ran } K^*}$. Thus every Ψ decomposes as $\Psi = \Psi_0 + \sum_n \langle v_n, \Psi \rangle v_n$ with $\Psi_0 \in \ker K$, the series converging in $\mathcal{H}_{\text{src}}$. Applying the bounded operator K term by term and

using (4.2) yields (4.3) with convergence in $\mathcal{H}_{\text{obs}}$. In particular $\text{Ran } K \subseteq \overline{\text{span}\{u_n\}}$ and each $u_n = \sigma_n^{-1} K v_n \in \text{Ran } K$, so $\{u_n\}$ is an orthonormal basis of $\overline{\text{Ran } K}$.

Finally, if infinitely many $\sigma_n > 0$ existed with $\sigma_n \geq \varepsilon > 0$ along a subsequence, then $\{K v_n\} = \{\sigma_n u_n\}$ would satisfy $\|K v_n - K v_m\|^2 = \sigma_n^2 + \sigma_m^2 \geq 2\varepsilon^2$, so the bounded sequence $\{v_n\}$ would have image with no convergent subsequence, contradicting compactness. Hence $\sigma_n \rightarrow 0$. $\square$

Definition 4.3 (Rank- r truncation). For $r \geq 0$ set

$$K_r := \sum_{n=1}^r \sigma_n u_n \langle v_n, \cdot \rangle, \quad (4.4)$$

that is, $K_r \Psi = \sum_{n=1}^r \sigma_n \langle v_n, \Psi \rangle u_n$. If K has finite rank $R \leq r$, the sum runs only over the available indices and $K_r = K$. We adopt the convention $\sigma_{r+1} := 0$ whenever the index $r+1$ exceeds the number of positive singular values.

Theorem 4.4 (Optimal compact reduction; Eckart–Young–Mirsky, operator norm). *Let K be compact and $r \geq 0$. Then for every $B \in \mathcal{B}(\mathcal{H}_{\text{src}}, \mathcal{H}_{\text{obs}})$ with $\text{rank } B \leq r$,*

$$\|K - B\|_{\text{op}} \geq \sigma_{r+1}, \quad (4.5)$$

and equality is attained at $B = K_r$. Consequently

$$\|K - K_r\|_{\text{op}} = \sigma_{r+1} = \min\{\|K - B\|_{\text{op}} : \text{rank } B \leq r\}. \quad (4.6)$$

Proof. Upper bound. By (4.3) and (4.4), $(K - K_r)\Psi = \sum_{n>r} \sigma_n \langle v_n, \Psi \rangle u_n$. By orthonormality of $\{u_n\}$ and Bessel's inequality for $\{v_n\}$,

$$\|(K - K_r)\Psi\|^2 = \sum_{n>r} \sigma_n^2 |\langle v_n, \Psi \rangle|^2 \leq \sigma_{r+1}^2 \sum_{n>r} |\langle v_n, \Psi \rangle|^2 \leq \sigma_{r+1}^2 \|\Psi\|^2,$$

using $\sigma_n \leq \sigma_{r+1}$ for $n > r$. Hence $\|K - K_r\|_{\text{op}} \leq \sigma_{r+1}$. Testing on $\Psi = v_{r+1}$ (when $\sigma_{r+1} > 0$) gives $\|(K - K_r)v_{r+1}\| = \sigma_{r+1}$, so equality holds.

Lower bound. We may assume $\sigma_{r+1} > 0$, else (4.5) is vacuous. Let $\text{rank } B \leq r$ and put $W := \text{span}\{v_1, \dots, v_{r+1}\}$, of dimension $r+1$. The restriction $B|_W : W \rightarrow \mathcal{H}_{\text{obs}}$ has rank at most $r < \dim W$, so by rank-nullity there exists $v \in W$ with $\|v\| = 1$ and $Bv = 0$. Write $v = \sum_{n=1}^{r+1} c_n v_n$ with $\sum_{n=1}^{r+1} |c_n|^2 = 1$. Then by (4.3),

$$\|(K - B)v\|^2 = \|Kv\|^2 = \left\| \sum_{n=1}^{r+1} c_n \sigma_n u_n \right\|^2 = \sum_{n=1}^{r+1} \sigma_n^2 |c_n|^2 \geq \sigma_{r+1}^2 \sum_{n=1}^{r+1} |c_n|^2 = \sigma_{r+1}^2,$$

since $\sigma_n \geq \sigma_{r+1}$ for $n \leq r+1$. Therefore $\|K - B\|_{\text{op}} \geq \|(K - B)v\| \geq \sigma_{r+1}$. $\square$

Theorem 4.4 is the operator-norm form of the Eckart–Young–Mirsky theorem [14, 2, 9]. It quantifies (D3) and nothing else: it is an approximation statement about *how well a finite-rank operator can imitate K* . It says nothing about $\ker K$ and nothing about invertibility.

Example 4.5 (Compact, injective, unstably invertible). On $\mathcal{H}_{\text{src}} = \mathcal{H}_{\text{obs}} = \ell^2(\mathbb{N})$ define

$$K(x_1, x_2, x_3, \dots) := \left(x_1, \frac{x_2}{2}, \frac{x_3}{3}, \dots\right), \quad \text{i.e.} \quad (Kx)_n = \frac{x_n}{n}. \quad (4.7)$$

Let $\{e_n\}$ be the standard basis. Then:

- (1) **K is compact.** Let K_r be as in Definition 4.3, which here is $(K_r x)_n = x_n/n$ for $n \leq r$ and 0 for $n > r$; this has rank r . For any x with $\|x\| \leq 1$, $\|(K - K_r)x\|^2 = \sum_{n>r} |x_n|^2/n^2 \leq (r+1)^{-2}$, so $\|K - K_r\|_{\text{op}} \leq 1/(r+1) \rightarrow 0$. A norm limit of finite-rank operators is compact [13, Thm. VI.12]. (This also exhibits the singular system explicitly:

- $v_n = u_n = e_n$, $\sigma_n = 1/n$, and Theorem 4.4 gives the exact value $\|K - K_r\|_{\text{op}} = \sigma_{r+1} = 1/(r+1)$.)
- (2) **K is injective.** If $Kx = 0$ then $x_n/n = 0$ for every n , so $x = 0$. Hence $\ker K = \{0\}$ and, by Corollary 2.5, no exact source distinction is lost: the observational equivalence relation $\sim_K$ is equality.
 - (3) **The inverse on the range is unbounded, and $\text{Ran } K$ is not closed.** On $\text{Ran } K$ the inverse is $(K^{-1}y)_n = n y_n$. Then $K^{-1}e_n = n e_n$ has norm n while $\|e_n\| = 1$, so K^{-1} is unbounded. By Theorem 2.3, $\text{Ran } K$ is therefore not closed. Concretely, $y = (1/n)_{n \geq 1} \in \ell^2$ is the limit of $K((1, \dots, 1, 0, 0, \dots))$ but is not in $\text{Ran } K$, since $K^{-1}y = (1, 1, 1, \dots) \notin \ell^2$.

Example 4.5 decouples the three static phenomena completely. K is compact, so it admits arbitrarily good, quantitatively controlled finite-rank approximations (D3)—and, more sharply, Theorem 4.4 gives the exact approximation error. K is injective, so (D1) fails: nothing is exactly identified, and the quotient $\mathcal{H}_K$ is $\mathcal{H}_{\text{src}}$ itself. Yet K^{-1} is unbounded, so (D2) holds in full force: the source is determined by the readout, but not continuously, and no stability estimate as in Theorem 2.3(3) is available. *Compactness does not imply a nontrivial kernel*—what it does imply is the controlled finite-rank approximation of Theorem 4.4—and a trivial kernel does not imply stable inversion. Note also that for this K , (D4) cannot occur under *any* source semigroup whatever: $\ker K = \{0\}$ is invariant trivially, so by (1.1) the dynamics always descend, however violently the inversion is ill-conditioned. Instability of the inverse and failure of dynamical descent are different diseases. Any argument that slides between “the observation is compact”, “information is lost”, and “the source is nonidentifiable” has already conflated three different theorems.

5. DESCENT OF SOURCE DYNAMICS

We now add dynamics. Throughout this section, K satisfies Assumption 2.1, $N := \ker K$, and $\{U(t)\}_{t \geq 0}$ is a C_0 -semigroup of bounded operators on $\mathcal{H}_{\text{src}}$ representing source evolution $\Psi(t) = U(t)\Psi_0$.

Definition 5.1 (Induced observable evolution). A family $\{V(t)\}_{t \geq 0}$ of linear maps $\text{Ran } K \rightarrow \text{Ran } K$ is an *induced observable evolution* if

$$K U(t) = V(t) K \quad \text{on } \mathcal{H}_{\text{src}}, \text{ for every } t \geq 0. \quad (5.1)$$

Equation (5.1) is the precise sense in which the readout “has its own law of motion”: the observable state at time t is a function of the observable state at time 0 alone, uniformly in the source state realizing it.

Theorem 5.2 (Semigroup descent theorem). *The following are equivalent.*

- (1) *There exists an induced observable evolution $\{V(t)\}_{t \geq 0}$ in the sense of Definition 5.1.*
- (2) *Observational equivalence is preserved by source evolution:*

$$U(t)N \subseteq N \quad \text{for every } t \geq 0. \quad (5.2)$$

When these hold, $V(t)$ is unique, is given by

$$V(t)(K\Psi) := K U(t)\Psi \quad (\Psi \in \mathcal{H}_{\text{src}}), \quad (5.3)$$

and satisfies the semigroup laws $V(0) = I_{\text{Ran } K}$ and $V(t+s) = V(t)V(s)$. Equivalently,

$$\bar{U}(t)[\Psi] := [U(t)\Psi] \quad (5.4)$$

is a well-defined one-parameter semigroup on the quotient Hilbert space $\mathcal{H}_K$, and $\tilde{K} \bar{U}(t) = V(t) \tilde{K}$.

Proof. (1) $\Rightarrow$ (2). Let $n \in N$ and $t \geq 0$. Then $Kn = 0$, so by (5.1),

$$KU(t)n = V(t)Kn = V(t)0 = 0,$$

i.e. $U(t)n \in \ker K = N$. As $n \in N$ and $t \geq 0$ were arbitrary, (5.2) holds.

(2) $\Rightarrow$ (1). Fix $t \geq 0$ and suppose $K\Psi = K\Psi'$. Then $\Psi - \Psi' \in N$, so by (5.2) $U(t)(\Psi - \Psi') \in N$, hence

$$KU(t)\Psi - KU(t)\Psi' = KU(t)(\Psi - \Psi') = 0.$$

Therefore the prescription (5.3) depends only on the value $K\Psi \in \text{Ran } K$ and not on the representative Ψ ; it defines a map $V(t) : \text{Ran } K \rightarrow \mathcal{H}_{\text{obs}}$, whose image lies in $\text{Ran } K$ because $V(t)K\Psi = K(U(t)\Psi) \in \text{Ran } K$. Linearity of $V(t)$ follows from linearity of K and $U(t)$ together with surjectivity of K onto $\text{Ran } K$. Equation (5.1) holds by construction.

Semigroup laws and uniqueness. $V(0)K\Psi = KU(0)\Psi = K\Psi$, so $V(0) = I$ on $\text{Ran } K$. For $s, t \geq 0$,

$$V(t)V(s)K\Psi = V(t)K(U(s)\Psi) = KU(t)U(s)\Psi = KU(t+s)\Psi = V(t+s)K\Psi,$$

and since K is onto $\text{Ran } K$ this gives $V(t)V(s) = V(t+s)$. Uniqueness is likewise forced: any $V'(t)$ satisfying (5.1) agrees with $V(t)$ at every point $K\Psi$ of $\text{Ran } K$.

Quotient statement. Well-definedness of $\bar{U}(t)$ in (5.4) is exactly the implication just proved: $[\Psi] = [\Psi']$ means $\Psi - \Psi' \in N$, so $U(t)\Psi - U(t)\Psi' \in N$, so $[U(t)\Psi] = [U(t)\Psi']$. The semigroup laws for $\bar{U}$ follow from those of U applied to representatives, and the intertwining $\tilde{K}\bar{U}(t)[\Psi] = KU(t)\Psi = V(t)K\Psi = V(t)\tilde{K}[\Psi]$ is immediate. Conversely, if $\bar{U}(t)$ is well defined by (5.4) for all t , then applying it to $[n] = 0$ for $n \in N$ gives $[U(t)n] = 0$, i.e. (5.2). $\square$

Proposition 5.3 (Topology of the induced evolution). *Assume (5.2). Then:*

- (1) $\bar{U}(t)$ is a C_0 -semigroup of bounded operators on the quotient Hilbert space $\mathcal{H}_K$, with $\|\bar{U}(t)\|_{\mathcal{B}(\mathcal{H}_K)} \leq \|U(t)\|_{\mathcal{B}(\mathcal{H}_{\text{src}})}$.
- (2) Consequently $V(t) = \tilde{K}\bar{U}(t)\tilde{K}^{-1}$ is a C_0 -semigroup of bounded operators on $\text{Ran } K$ equipped with the transported quotient norm $\|\cdot\|_{\sharp}$ of Remark 2.4.
- (3) If $\text{Ran } K$ is closed, then $\|\cdot\|_{\sharp}$ and the inherited norm $\|\cdot\|_{\mathcal{H}_{\text{obs}}}$ are equivalent on $\text{Ran } K$, and hence $V(t)$ is a C_0 -semigroup of bounded operators on $\text{Ran } K$ in the inherited observable norm as well.
- (4) Orbit continuity in the inherited norm is automatic, for every K , closed range or not. For each $y = K\Psi \in \text{Ran } K$ the orbit $t \mapsto V(t)y$ is continuous in $\|\cdot\|_{\mathcal{H}_{\text{obs}}}$.
- (5) What can fail when $\text{Ran } K$ is not closed is boundedness of the operator $V(t)$ on the normed space $(\text{Ran } K, \|\cdot\|_{\mathcal{H}_{\text{obs}}})$. Precisely, $V(t)$ is bounded in the inherited norm if and only if there is a finite constant C_t with

$$\|KU(t)\Psi\|_{\mathcal{H}_{\text{obs}}} \leq C_t \|K\Psi\|_{\mathcal{H}_{\text{obs}}} \quad \text{for every } \Psi \in \mathcal{H}_{\text{src}}, \quad (5.5)$$

and this is an additional hypothesis: it does not follow from (5.2) alone. If (5.5) does hold for every $t \geq 0$, then $\{V(t)\}$ is a strongly continuous semigroup of bounded operators on the possibly incomplete normed space $(\text{Ran } K, \|\cdot\|_{\mathcal{H}_{\text{obs}}})$.

- (6) Local uniform boundedness does not follow from (5) in the nonclosed case, because the Banach–Steinhaus theorem requires completeness of the domain and $\text{Ran } K$ is then in general incomplete. If one additionally assumes

$$\sup_{0 \leq t \leq T} \|V(t)\|_{\text{op}} < \infty \quad \text{for every } T > 0, \quad (5.6)$$

then each $V(t)$ extends uniquely to a bounded operator $\widehat{V}(t)$ on the completion $\overline{\text{Ran } K}$, and $\{\widehat{V}(t)\}$ is a C_0 -semigroup there. In the closed-range case $\text{Ran } K$ is already complete, (5.6) is automatic, and the ordinary C_0 -semigroup conclusions of (3) apply unchanged.

Proof. (1) Boundedness: for any $n \in N$, (5.2) gives $U(t)n \in N$, hence

$$\|\bar{U}(t)[\Psi]\|_K = \inf_{m \in N} \|U(t)\Psi + m\| \leq \inf_{n \in N} \|U(t)(\Psi + n)\| \leq \|U(t)\| \inf_{n \in N} \|\Psi + n\| = \|U(t)\| \|\Psi\|_K,$$

where in the first inequality we restricted the infimum to the subfamily $m = U(t)n$, $n \in N$, which is legitimate precisely because $U(t)N \subseteq N$. Strong continuity: for fixed Ψ , $\|\bar{U}(t)[\Psi] - [\Psi]\|_K \leq \|U(t)\Psi - \Psi\|_{\mathcal{H}_{\text{src}}} \rightarrow 0$ as $t \downarrow 0$ by strong continuity of U , and the semigroup laws are Theorem 5.2. Strong continuity at $t = 0$ plus local boundedness gives strong continuity on $[0, \infty)$ by the standard argument [3, Prop. I.5.3].

(2) By construction $\tilde{K} : (\mathcal{H}_K, \|\cdot\|_K) \rightarrow (\text{Ran } K, \|\cdot\|_{\#})$ is an isometric isomorphism, and conjugation by an isometric isomorphism carries C_0 -semigroups to C_0 -semigroups.

(3) If $\text{Ran } K$ is closed, Theorem 2.3 makes $\tilde{K}$ and $\tilde{K}^{-1}$ bounded in the inherited norm, so $\|\cdot\|_{\#}$ and $\|\cdot\|_{\mathcal{H}_{\text{obs}}}$ are equivalent on $\text{Ran } K$; equivalent norms have the same bounded operators and the same strongly convergent limits.

(4) Let $y = K\Psi \in \text{Ran } K$ and $t, t_0 \geq 0$. By (5.3) and boundedness of K ,

$$\|V(t)y - V(t_0)y\|_{\mathcal{H}_{\text{obs}}} = \|K(U(t)\Psi - U(t_0)\Psi)\|_{\mathcal{H}_{\text{obs}}} \leq \|K\|_{\text{op}} \|U(t)\Psi - U(t_0)\Psi\|_{\mathcal{H}_{\text{src}}} \xrightarrow{t \rightarrow t_0} 0$$

by strong continuity of U . Thus every orbit is continuous in the inherited norm, whatever $\text{Ran } K$ may be. (Note that the estimate is made through the *representative* Ψ , which is legitimate because $V(t)y$ does not depend on the choice of representative, by Theorem 5.2.)

(5) The stated equivalence is the definition of boundedness of the linear map $V(t)$ from $(\text{Ran } K, \|\cdot\|_{\mathcal{H}_{\text{obs}}})$ to itself, once one writes a general element of $\text{Ran } K$ as $K\Psi$ and recalls $V(t)K\Psi = KU(t)\Psi$: the inequality $\|V(t)y\| \leq C_t\|y\|$ for all $y \in \text{Ran } K$ is exactly (5.5). It is not implied by (5.2): in the nonclosed case $\tilde{K}^{-1}$ is unbounded in the inherited norm, so the conjugation argument of (2) is unavailable, and no such C_t is produced.

Suppose now (5.5) holds for every $t \geq 0$. Then each $V(t)$ is a bounded operator on the normed space $\text{Ran } K$, the semigroup laws hold by Theorem 5.2, and by part (4) the map $t \mapsto V(t)y$ is continuous for every $y \in \text{Ran } K$; so $\{V(t)\}$ is a strongly continuous semigroup of bounded operators on $\text{Ran } K$. That is all that follows, and we stop there.

(6) We do *not* obtain (5.6) from (5). The usual route—pointwise boundedness of $\{V(t) : 0 \leq t \leq T\}$ on each vector (which does hold, each orbit being the continuous image of a compact interval) plus Banach–Steinhaus—is unavailable, because Banach–Steinhaus requires the domain to be complete, and when $\text{Ran } K$ is not closed it is an incomplete normed space. Nor may one evade this by passing to the completion $\overline{\text{Ran } K}$ first: pointwise boundedness has been established only on the dense subspace $\text{Ran } K$, not at every vector of $\overline{\text{Ran } K}$, and a family of operators that is pointwise bounded on a dense subspace need not be pointwise bounded on the whole space absent a uniform bound—which is precisely what is at issue. So (5.6) is an extra hypothesis.

Assume it. Then for $0 \leq t \leq T$ the operator $V(t)$ is bounded on the dense subspace $\text{Ran } K$ of $\overline{\text{Ran } K}$ with $\|V(t)\|_{\text{op}} \leq M_T := \sup_{0 \leq t \leq T} \|V(t)\|_{\text{op}} < \infty$, so it extends uniquely to a bounded operator $\widehat{V}(t)$ on $\overline{\text{Ran } K}$ with the same norm (uniformly continuous extension from a dense subspace of a complete space). The semigroup laws pass to the closure by density and continuity. For strong continuity, let $\hat{y} \in \overline{\text{Ran } K}$ and $\varepsilon > 0$; pick $y \in \text{Ran } K$ with $\|\hat{y} - y\| < \varepsilon$; then for $0 \leq t \leq T$,

$$\|\widehat{V}(t)\hat{y} - \hat{y}\| \leq \|\widehat{V}(t)(\hat{y} - y)\| + \|V(t)y - y\| + \|y - \hat{y}\| \leq (M_T + 1)\varepsilon + \|V(t)y - y\|,$$

and the middle term tends to 0 as $t \downarrow 0$ by part (4). Hence $\widehat{V}(t)\hat{y} \rightarrow \hat{y}$, and $\{\widehat{V}(t)\}$ is a C_0 -semigroup on $\overline{\text{Ran } K}$. In the closed-range case $\text{Ran } K = \overline{\text{Ran } K}$ is complete, (5.6) follows from the standard local boundedness of the C_0 -semigroup already produced in (3), and nothing new is needed. $\square$

The three-way distinction just drawn—orbit continuity (free), operator boundedness (5.5) (an extra hypothesis), local uniform boundedness (5.6) (a further one)—collapses entirely when $\text{Ran } K$ is closed, and it is only the nonclosed case that forces one to keep the three apart.

Corollary 5.4 (Failure of autonomous descent). *Suppose there exist $n \in \ker K$ and $t_0 > 0$ with*

$$KU(t_0)n \neq 0. \quad (5.7)$$

Then the two source states $\Psi = 0$ and $\Psi' = n$ satisfy $K\Psi = K\Psi' = 0$ at time zero but $KU(t_0)\Psi = 0 \neq KU(t_0)\Psi'$. Consequently there is no single-valued map $V(t_0)$ on $\text{Ran } K$ satisfying (5.1): no autonomous law on the present observable state reproduces the readouts of both source trajectories. More generally, for any Ψ , the pair Ψ and $\Psi + n$ are observationally identical at $t = 0$ and observationally distinct at $t = t_0$.

Proof. If such a $V(t_0)$ existed, (5.1) applied to n would give $KU(t_0)n = V(t_0)Kn = 0$, contradicting (5.7). The general statement follows by linearity: $K(\Psi + n) = K\Psi$, while $KU(t_0)(\Psi + n) - KU(t_0)\Psi = KU(t_0)n \neq 0$. $\square$

Corollary 5.4 is (D4) in its bluntest form, and it should be contrasted carefully with Corollary 2.5. Corollary 2.5 says the readout at a single time cannot distinguish Ψ from $\Psi + n$. Corollary 5.4 says something strictly stronger: the *future* readouts *can* distinguish them. The hypothesis of Corollary 5.4 presupposes a nonzero $n \in \ker K$, so it can only be met when $\ker K \neq \{0\}$ —this is (1.1) again, and it is why (D4) is not logically independent of (D1) but strictly stronger than it. What (D1) leaves open, and what Theorem 5.2 decides, is whether the nontrivial kernel is invariant. Hidden structure that is dynamically inert ($U(t)N \subseteq N$) costs the observable theory nothing, however large N is; hidden structure that is dynamically active makes the observable theory nonautonomous, however tidy its instantaneous state space may be. The whole of Sections 6–7 is a quantitative account of exactly how much it costs.

Proposition 5.5 (Bounded generator). *Suppose $L \in \mathcal{B}(\mathcal{H}_{\text{src}})$ and $U(t) = e^{tL}$. Then (5.2) holds if and only if*

$$LN \subseteq N. \quad (5.8)$$

Proof. ($\Leftarrow$) If $LN \subseteq N$ then $L^k N \subseteq N$ for all $k \geq 0$ by induction. For $n \in N$, the partial sums $\sum_{k=0}^m \frac{t^k}{k!} L^k n$ lie in N and converge in norm to $e^{tL}n$; since N is closed, $e^{tL}n \in N$.

($\Rightarrow$) If $e^{tL}N \subseteq N$ for all $t \geq 0$ and $n \in N$, then $t^{-1}(e^{tL}n - n) \in N$ for $t > 0$, and N closed gives $Ln = \lim_{t \downarrow 0} t^{-1}(e^{tL}n - n) \in N$, the limit existing in norm because L is bounded. $\square$

Remark 5.6 (Unbounded generators). If L is the (generally unbounded) generator of $U(t)$ with domain $D(L)$, the algebraic condition

$$L(N \cap D(L)) \subseteq N$$

does *not* by itself settle (5.2). Two separate issues arise: domain invariance ($U(t)$ need not map $N \cap D(L)$ into $D(L)$ in a way compatible with the difference-quotient argument above) and the density/core structure of $N \cap D(L)$ in N (which may be trivial in pathological cases). Correct criteria for invariance of a closed subspace under a C_0 -semigroup are formulated in terms of the resolvent—for instance, $U(t)N \subseteq N$ for all $t \geq 0$ if and only if $(\lambda - L)^{-1}N \subseteq N$ for all sufficiently large real λ —and require verification in each application; see [3, §II.2 and §IV.2].

We do not use unbounded generators anywhere in the theorems below, and we do not assert their invariance theory here.

6. THE EXACT PROJECTED-DYNAMICS IDENTITY

Assumption 6.1. For the remainder of Sections 6–7: $\mathcal{H}$ is a Hilbert space; $P \in \mathcal{B}(\mathcal{H})$ is an orthogonal projection onto a closed *resolved* subspace $\mathcal{H}_P := P\mathcal{H}$; $Q := I - P$ is the orthogonal projection onto the *unresolved* subspace $\mathcal{H}_Q := Q\mathcal{H}$; and $L \in \mathcal{B}(\mathcal{H})$ generates the source evolution

$$\dot{\Psi}(t) = L\Psi(t), \quad \Psi(0) = \Psi_0 \in \mathcal{H}, \quad (6.1)$$

whose unique solution is $\Psi(t) = e^{tL}\Psi_0$.

This is the case $K = P$ of the previous section: the readout literally *is* the resolved component, $\mathcal{H}_{\text{obs}} = \mathcal{H}_P$, $\text{Ran } P = \mathcal{H}_P$ is closed, and $\ker P = \mathcal{H}_Q$. It is the setting in which the cost of failed descent can be computed in closed form. Write

$$x(t) := P\Psi(t) \in \mathcal{H}_P, \quad z(t) := Q\Psi(t) \in \mathcal{H}_Q, \quad (6.2)$$

so that $\Psi(t) = x(t) + z(t)$ and $x(0) = P\Psi_0$, $z(0) = z_0 = Q\Psi_0$. Applying P and Q to (6.1) and inserting $I = P + Q$ gives the exact block system

$$\dot{x}(t) = PLP x(t) + PLQ z(t), \quad (6.3)$$

$$\dot{z}(t) = QLP x(t) + QLQ z(t). \quad (6.4)$$

Note that $PLQ : \mathcal{H}_Q \rightarrow \mathcal{H}_P$ measures precisely how the unresolved sector drives the resolved one, and $QLP : \mathcal{H}_P \rightarrow \mathcal{H}_Q$ how the resolved sector feeds back into the unresolved one.

Theorem 6.2 (Exact projected-dynamics equation). *Under Assumption 6.1, for every source initial state $\Psi_0 = x(0) + z_0$ the resolved component $x(t) = Pe^{tL}\Psi_0$ satisfies, for all $t \geq 0$,*

$$\boxed{\dot{x}(t) = PLP x(t) + PLQ e^{tQLQ} z_0 + \int_0^t PLQ e^{(t-s)QLQ} QLP x(s) ds.} \quad (6.5)$$

Defining the unresolved-initial-state term

$$\eta_{z_0}(t) := PLQ e^{tQLQ} z_0 \in \mathcal{H}_P \quad (6.6)$$

and the memory kernel

$$\mathcal{K}(\tau) := PLQ e^{\tau QLQ} QLP \in \mathcal{B}(\mathcal{H}_P), \quad (6.7)$$

equation (6.5) reads

$$\dot{x}(t) = PLP x(t) + \eta_{z_0}(t) + \int_0^t \mathcal{K}(t-s) x(s) ds. \quad (6.8)$$

Both η_{z_0} and $\mathcal{K}$ are continuous in t , with the bounds $\|\eta_{z_0}(t)\| \leq \|PLQ\| e^{t\|QLQ\|} \|z_0\|$ and $\|\mathcal{K}(\tau)\|_{\text{op}} \leq \|PLQ\| \|QLP\| e^{\tau\|QLQ\|}$.

Proof. Equation (6.4) is an inhomogeneous linear ODE in the Banach space $\mathcal{H}_Q$ with bounded generator $QLQ \in \mathcal{B}(\mathcal{H}_Q)$ and continuous inhomogeneity $t \mapsto QLP x(t)$ (continuous because $t \mapsto e^{tL}\Psi_0$ is norm-continuous and QLP is bounded). Variation of constants gives the unique solution

$$z(t) = e^{tQLQ} z_0 + \int_0^t e^{(t-s)QLQ} QLP x(s) ds, \quad (6.9)$$

the integral being a Riemann (equivalently Bochner) integral of a continuous $\mathcal{H}_Q$ -valued integrand. Indeed, put $w(t) := e^{-tQLQ} z(t)$; then w is differentiable with $\dot{w}(t) = e^{-tQLQ} (\dot{z}(t) -$

$QLQ z(t) = e^{-tQLQ}QLP x(t)$ by (6.4), and integrating from 0 to t and multiplying by e^{tQLQ} yields (6.9). Uniqueness holds because the difference of two solutions solves the homogeneous equation with zero data.

Substituting (6.9) into (6.3) and applying the bounded operator PLQ under the integral sign gives

$$\dot{x}(t) = PLP x(t) + PLQ e^{tQLQ} z_0 + \int_0^t PLQ e^{(t-s)QLQ} QLP x(s) ds,$$

which is (6.5); (6.8) is a rewriting using (6.6) and (6.7), noting $\mathcal{K}(t-s)$ has argument $t-s \geq 0$ throughout the integration range. The stated bounds follow from $\|e^{\tau QLQ}\|_{\text{op}} \leq e^{\tau\|QLQ\|}$ and submultiplicativity; continuity in t follows from norm-continuity of $\tau \mapsto e^{\tau QLQ}$. $\square$

Remark 6.3 (What (6.5) is, and what it is not). Equation (6.5) is an *identity*, valid for every source solution, obtained by nothing more than variation of constants. It is the projection-operator identity of Mori and Zwanzig [20, 10, 21] in the special case of a bounded generator and an orthogonal projection. In that literature the second term is often called the “fluctuating force” or “noise”, and $\mathcal{K}$ the “friction kernel”. We use neither word without qualification, and we record explicitly what has *not* been done here:

- (1) η_{z_0} is *not stochastic*. It is a deterministic, linear, continuous function of the single vector $z_0 \in \mathcal{H}_Q$, given by the closed formula (6.6). It becomes a random process only if one additionally posits a probability measure on the unresolved initial data—an equilibrium ensemble, say—and that is an extra modelling hypothesis with its own physical content. We impose no such measure.
- (2) $\mathcal{K}$ is *not a friction coefficient*. It is the exact kernel (6.7). No fluctuation–dissipation relation has been derived, because no ensemble has been introduced with respect to which one could be stated.
- (3) *No Markov approximation has been made*. The convolution in (6.8) is over the entire history $[0, t]$. Replacing $\mathcal{K}$ by δ -correlated kernel, or taking a timescale-separation limit, would produce a different (approximate) equation; we take no such limit.
- (4) *No ensemble average or coarse-graining functional appears*. P is an orthogonal projection on a Hilbert space, chosen by the modeller, and nothing else.

The physical content of (6.8) is therefore modest and exact: if the resolved sector does not close, then the resolved trajectory is determined by resolved initial data *plus* an inhomogeneity carrying unresolved initial data *plus* a convolution over its own past. These are three structurally distinct dependences, and (6.8) says they are the only ones.

Theorem 6.4 (Autonomous-closure criterion). *Under Assumption 6.1, the following are equivalent.*

- (1) *There exists a single operator $Y \in \mathcal{B}(\mathcal{H}_P)$ such that, for every source initial state $\Psi_0 \in \mathcal{H}$ —that is, for every $x(0) \in \mathcal{H}_P$ and every $z_0 \in \mathcal{H}_Q$ —the resolved component $x(t) = Pe^{tL}\Psi_0$ satisfies the autonomous first-order equation*

$$\dot{x}(t) = Y x(t), \quad t \geq 0. \tag{6.10}$$

- (2) $PLQ = 0$.
- (3) $L(\ker P) \subseteq \ker P$, i.e. $L\mathcal{H}_Q \subseteq \mathcal{H}_Q$.

When they hold, $Y = PLP|_{\mathcal{H}_P}$ is unique, and both $\eta_{z_0} \equiv 0$ (for every z_0) and $\mathcal{K} \equiv 0$.

Proof. (2) $\Rightarrow$ (1) and the value of Y . If $PLQ = 0$ then (6.3) reduces immediately to $\dot{x}(t) = PLP x(t)$ for every source solution, which is (6.10) with $Y = PLP|_{\mathcal{H}_P}$. Moreover $\eta_{z_0}(t) =$

$PLQ e^{tQLQ} z_0 = 0$ for all t and all z_0 , and $\mathcal{K}(\tau) = PLQ e^{\tau QLQ} QLP = 0$, so the last two terms of (6.8) vanish identically, consistently.

(1) $\Rightarrow$ (2). This is where the quantifier “for every source initial state” does the work. Fix any $x_0 \in \mathcal{H}_P$ and any two unresolved data $z_0, z'_0 \in \mathcal{H}_Q$. The two source initial states $\Psi_0 = x_0 + z_0$ and $\Psi'_0 = x_0 + z'_0$ have the same resolved component $x(0) = x'(0) = x_0$. By hypothesis both resolved trajectories obey (6.10) with the *same* Y , so evaluating at $t = 0$,

$$\dot{x}(0) = Yx_0 = \dot{x}'(0).$$

On the other hand (6.3) at $t = 0$ gives the true initial velocities

$$\dot{x}(0) = PLP x_0 + PLQ z_0, \quad \dot{x}'(0) = PLP x_0 + PLQ z'_0.$$

Equating and cancelling $PLP x_0$ yields $PLQ(z_0 - z'_0) = 0$. Since z_0, z'_0 range over all of $\mathcal{H}_Q$, the difference $z_0 - z'_0$ ranges over all of $\mathcal{H}_Q$, whence $PLQ = 0$ as an operator $\mathcal{H}_Q \rightarrow \mathcal{H}_P$. Taking $z_0 = 0$ in (6.10) and (6.3) then forces $Yx_0 = PLP x_0$ for all $x_0 \in \mathcal{H}_P$, i.e. $Y = PLP|_{\mathcal{H}_P}$, which also gives uniqueness of Y .

(2) $\Leftrightarrow$ (3). $\ker P = \mathcal{H}_Q = \text{Ran } Q$. Now $L\mathcal{H}_Q \subseteq \mathcal{H}_Q$ means $PL\zeta = 0$ for every $\zeta \in \mathcal{H}_Q$, i.e. $PLQ\Psi = 0$ for every $\Psi \in \mathcal{H}$, i.e. $PLQ = 0$. Conversely $PLQ = 0$ gives $PL\zeta = PLQ\zeta = 0$ for $\zeta \in \mathcal{H}_Q$, so $L\zeta = (P + Q)L\zeta = QL\zeta \in \mathcal{H}_Q$. $\square$

Remark 6.5 (The quantifier is not decorative). Theorem 6.4 is a statement about *all* source states. Three distinct weaker phenomena must not be mistaken for autonomous closure:

- (1) *Closure on a prepared state.* If $z_0 \in N_Q$ —see (7.2)—then $\eta_{z_0} \equiv 0$ and the *initial-state* term drops out *for that state*. The memory kernel $\mathcal{K}$ is unaffected: by (6.7) it is determined by L and P alone and does not depend on z_0 at all, so no preparation of the unresolved initial data can remove it. The resolved motion of such a state is therefore still history-dependent, and other initial states fail even the weaker condition. Vanishing of the initial-state term in one specially prepared state is not closure.
- (2) *Static quotient.* That $\mathcal{H}_P \cong \mathcal{H}/\mathcal{H}_Q$ is a perfectly good Hilbert space says nothing whatever about PLQ . This is (D1)/(D4) again.
- (3) *Small singular values.* That PLQ has small norm makes the coupling weak and may justify a perturbative approximation, with a controllable error. It does not make the projected equation autonomous, and no truncation of (6.8) is exact unless $PLQ = 0$ identically. This is (D3)/(D4).

Finally, observe that Theorem 6.4(3) is precisely the invariance condition (5.8) of Proposition 5.5 for the observation operator $K = P$, whose kernel is $\mathcal{H}_Q$. Theorems 5.2 and 6.4 therefore agree, as they must: the semigroup descends exactly when the projected equation closes.

Example 6.6 (A completely calculated block). Take $\mathcal{H} = \mathbb{C}^3$ with the standard basis e_1, e_2, e_3 , let P project onto $\mathcal{H}_P = \text{span}\{e_1\}$ and Q onto $\mathcal{H}_Q = \text{span}\{e_2, e_3\}$, and let

$$L = \begin{pmatrix} a & b_1 & b_2 \\ c_1 & d & 0 \\ c_2 & 0 & \mu \end{pmatrix},$$

so that

$$PLP = (a), \quad PLQ = (b_1 \ b_2), \quad QLP = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}, \quad QLQ = \begin{pmatrix} d & 0 \\ 0 & \mu \end{pmatrix},$$

with $d \neq \mu$. Write $x(t) \in \mathbb{C}$ for the resolved component and $z_0 = (z_0^{(2)}, z_0^{(3)})^\top$.

(i) *Same readout, different projected velocity.* Suppose $b_1 \neq 0$. The two source states $\Psi_0 = x_0 e_1$ and $\Psi'_0 = x_0 e_1 + e_2$ have identical resolved component x_0 , hence identical readout at $t = 0$. By

(6.3),

$$\dot{x}(0) = ax_0, \quad \dot{x}'(0) = ax_0 + b_1 \neq \dot{x}(0).$$

No autonomous law $\dot{x} = Yx$ can produce both, confirming Theorem 6.4: here $PLQ = (b_1 \ b_2) \neq 0$.

(ii) *The initial-state term and memory kernel.* Since QLQ is diagonal, $e^{\tau QLQ} = \text{diag}(e^{\tau d}, e^{\tau \mu})$, and (6.6), (6.7) give in closed form

$$\eta_{z_0}(t) = b_1 e^{td} z_0^{(2)} + b_2 e^{t\mu} z_0^{(3)}, \quad \mathcal{K}(\tau) = b_1 c_1 e^{\tau d} + b_2 c_2 e^{\tau \mu},$$

so the exact projected equation (6.8) is the scalar Volterra equation

$$\dot{x}(t) = a x(t) + b_1 e^{td} z_0^{(2)} + b_2 e^{t\mu} z_0^{(3)} + \int_0^t \left(b_1 c_1 e^{(t-s)d} + b_2 c_2 e^{(t-s)\mu} \right) x(s) ds.$$

Everything here is deterministic: η_{z_0} is a sum of two exponentials with coefficients read off from z_0 ; there is no noise in sight.

(iii) N_Q and E_{dyn} . Here $m = \dim \mathcal{H}_Q = 2$, so Proposition 7.7 gives $N_Q = \ker(PLQ) \cap \ker(PLQ \cdot QLQ)$, i.e. the solutions of

$$\begin{pmatrix} b_1 & b_2 \\ b_1 d & b_2 \mu \end{pmatrix} \begin{pmatrix} z^{(2)} \\ z^{(3)} \end{pmatrix} = 0.$$

The observability matrix has determinant $b_1 b_2 (\mu - d)$. *Case A:* $b_1, b_2 \neq 0$ and $d \neq \mu$. Then the determinant is nonzero, so $N_Q = \{0\}$ and $E_{\text{dyn}} = \mathcal{H}_Q \cong \mathbb{C}^2$: both unresolved directions are dynamically distinguishable and must be retained. *Case B:* $b_2 = 0$, $b_1 \neq 0$. Then both rows are multiples of $(1 \ 0)$, so $N_Q = \text{span}\{e_3\}$ and $E_{\text{dyn}} \cong \mathbb{C}$: the direction e_3 never influences the resolved trajectory through the initial-state term, and the minimal completion retains only $z^{(2)}$ (indeed $\mathcal{K}(\tau) = b_1 c_1 e^{\tau d}$ also loses its e_3 contribution). Note that e_3 is not dynamically trivial—it evolves as $e^{t\mu}$ —it is merely *unobservable in the resolved response*, which is exactly the relativization Theorem 7.4 makes precise.

7. MINIMAL DYNAMICAL COMPLETION

Theorem 6.4 says that when $PLQ \neq 0$ the resolved state $x(0)$ is not enough to determine the resolved future. The natural question is then: *how much* unresolved information must be adjoined? The answer is not “all of $\mathcal{H}_Q$ ”. Distinct unresolved initial states may produce identical resolved responses, and directions that never affect the resolved trajectory need not be retained. We now identify exactly the right quotient.

Fix the target: the resolved response function generated by unresolved initial data alone.

Definition 7.1 (Unresolved response map). Let $C([0, \infty), \mathcal{H}_P)$ carry the topology of uniform convergence on compact subsets of $[0, \infty)$, i.e. the locally convex topology generated by the seminorms $p_T(f) := \sup_{0 \leq t \leq T} \|f(t)\|_{\mathcal{H}_P}$, $T > 0$. Define the linear map

$$\mathcal{O}_Q : \mathcal{H}_Q \longrightarrow C([0, \infty), \mathcal{H}_P), \quad (\mathcal{O}_Q z_0)(t) := PLQ e^{tQLQ} z_0 = \eta_{z_0}(t). \quad (7.1)$$

$\mathcal{O}_Q$ is well defined (its values are continuous by Theorem 6.2) and continuous: $p_T(\mathcal{O}_Q z_0) \leq \|PLQ\| e^{T\|QLQ\|} \|z_0\|$ for every T .

Definition 7.2 (Dynamically invisible directions). Set

$$N_Q := \ker \mathcal{O}_Q = \left\{ z_0 \in \mathcal{H}_Q : PLQ e^{tQLQ} z_0 = 0 \text{ for all } t \geq 0 \right\} = \bigcap_{t \geq 0} \ker(PLQ e^{tQLQ}). \quad (7.2)$$

Lemma 7.3. N_Q is a closed subspace of $\mathcal{H}_Q$, and consequently

$$E_{\text{dyn}} := \mathcal{H}_Q / N_Q \quad (7.3)$$

is a Hilbert space with the quotient norm $\|[z_0]\|_{E_{\text{dyn}}} = \text{dist}(z_0, N_Q)$.

Proof. For each fixed $t \geq 0$ the operator $PLQ e^{tQLQ} : \mathcal{H}_Q \rightarrow \mathcal{H}_P$ is bounded, so its kernel is a closed subspace. By (7.2), N_Q is the intersection of a family of closed subspaces, hence a closed subspace. The quotient of a Hilbert space by a closed subspace is a Hilbert space with the quotient norm, exactly as established after (2.2). $\square$

Two unresolved initial states are called *dynamically equivalent for resolved evolution* when their difference lies in N_Q .

Theorem 7.4 (Minimal unresolved-state completion). *Under Assumption 6.1:*

- (1) $\mathcal{O}_Q$ is constant exactly on the cosets of N_Q : for $z_0, z'_0 \in \mathcal{H}_Q$, $\mathcal{O}_Q z_0 = \mathcal{O}_Q z'_0$ if and only if $z_0 - z'_0 \in N_Q$.
- (2) Consequently $\mathcal{O}_Q$ descends to an injective continuous linear map

$$\overline{\mathcal{O}} : E_{\text{dyn}} \longrightarrow C([0, \infty), \mathcal{H}_P), \quad \overline{\mathcal{O}}[z_0] := \mathcal{O}_Q z_0, \quad (7.4)$$

with $\mathcal{O}_Q = \overline{\mathcal{O}} \circ \pi$, where $\pi : \mathcal{H}_Q \rightarrow E_{\text{dyn}}$ is the quotient map.

- (3) (**Universality / minimality.**) Let E be any set and $q : \mathcal{H}_Q \rightarrow E$ any surjective map which determines every resolved response function, meaning that there exists a map $F : E \rightarrow C([0, \infty), \mathcal{H}_P)$ with

$$\mathcal{O}_Q = F \circ q. \quad (7.5)$$

Then there is a unique map $\rho : E \rightarrow E_{\text{dyn}}$ with

$$\pi = \rho \circ q, \quad (7.6)$$

and ρ is surjective. If moreover E is a vector space and q is linear, then ρ is linear.

Thus E_{dyn} is the coarsest extension of the resolved data that retains every unresolved distinction visible in the resolved response, and every sufficient extension maps onto it.

Proof. (1) $\mathcal{O}_Q$ is linear, so $\mathcal{O}_Q z_0 = \mathcal{O}_Q z'_0$ iff $\mathcal{O}_Q(z_0 - z'_0) = 0$ iff $z_0 - z'_0 \in \ker \mathcal{O}_Q = N_Q$.

(2) By (1), the prescription (7.4) is independent of the representative, hence well defined; it is linear because π and $\mathcal{O}_Q$ are; and it is injective because $\overline{\mathcal{O}}[z_0] = 0$ forces $z_0 \in N_Q$, i.e. $[z_0] = 0$. Continuity: for each T , $p_T(\overline{\mathcal{O}}[z_0]) = p_T(\mathcal{O}_Q(z_0 + n)) \leq \|PLQ\| e^{T\|QLQ\|} \|z_0 + n\|$ for every $n \in N_Q$; taking the infimum over n gives $p_T(\overline{\mathcal{O}}[z_0]) \leq \|PLQ\| e^{T\|QLQ\|} \|[z_0]\|_{E_{\text{dyn}}}$.

(3) *Existence of ρ .* We first show q separates no more coarsely than π ; that is, $q(z_0) = q(z'_0) \Rightarrow \pi(z_0) = \pi(z'_0)$. Indeed if $q(z_0) = q(z'_0)$ then by (7.5)

$$\mathcal{O}_Q z_0 = F(q(z_0)) = F(q(z'_0)) = \mathcal{O}_Q z'_0,$$

so by (1) $z_0 - z'_0 \in N_Q$, i.e. $\pi(z_0) = \pi(z'_0)$. Now define ρ as follows: given $e \in E$, use surjectivity of q to choose z_0 with $q(z_0) = e$, and set $\rho(e) := \pi(z_0)$. The implication just proved shows this is independent of the choice of z_0 , so ρ is a well-defined map $E \rightarrow E_{\text{dyn}}$, and (7.6) holds by construction. It is surjective because π is: any $[z_0] \in E_{\text{dyn}}$ equals $\rho(q(z_0))$.

Uniqueness of ρ . If ρ' also satisfies $\pi = \rho' \circ q$, then for every $e \in E$ pick z_0 with $q(z_0) = e$; then $\rho'(e) = \rho'(q(z_0)) = \pi(z_0) = \rho(q(z_0)) = \rho(e)$. Hence $\rho' = \rho$. (Uniqueness uses surjectivity of q , which is why it is included in the hypothesis: a nonsurjective extension leaves ρ undetermined on the unreachable part of E .)

Linearity. If E is a vector space and q is linear and surjective, then for $e = q(z_0)$, $e' = q(z'_0)$ and scalars α, β we have $\alpha e + \beta e' = q(\alpha z_0 + \beta z'_0)$, so $\rho(\alpha e + \beta e') = \pi(\alpha z_0 + \beta z'_0) = \alpha \pi(z_0) + \beta \pi(z'_0) = \alpha \rho(e) + \beta \rho(e')$. $\square$

We have stated the universal property in plain terms of maps and quotients of sets and vector spaces; no categorical apparatus is needed, and none is used. Theorem 7.4(3) is the minimal-completion statement of [24] for the specific target relation $z_0 \mapsto \mathcal{O}_Q z_0$, proved here directly so that the present paper stands alone.

Remark 7.5 (Minimality is relative to the target). E_{dyn} may be strictly smaller than $\mathcal{H}_Q$, as Case B of Example 6.6 shows. The directions in N_Q are not “physically absent” in any absolute sense: they may evolve nontrivially under e^{tQLQ} , they may carry energy, they may be visible to a *different* readout. They are invisible *to the resolved response function*, and that is the target relative to which minimality is asserted. Change the target—ask instead for the resolved response *and* some other functional of the unresolved sector—and the minimal completion changes accordingly. This target-relativity is exactly the phenomenon isolated in [23]; it is not a defect but the content of the theorem.

Theorem 7.6 (Completed initial data close the resolved trajectory). *Under Assumption 6.1, define the completed initial state space*

$$\widehat{\mathcal{H}}_P := \mathcal{H}_P \times E_{\text{dyn}}. \quad (7.7)$$

Then the pair $(x(0), [z_0]) \in \widehat{\mathcal{H}}_P$ determines the resolved trajectory $x(\cdot)$ uniquely on $[0, \infty)$. Precisely: for each $T > 0$, the linear Volterra equation of the second kind

$$x(t) = x(0) + \int_0^t \left(PLP x(s) + \eta_{z_0}(s) + \int_0^s \mathcal{K}(s-r)x(r) dr \right) ds, \quad 0 \leq t \leq T, \quad (7.8)$$

equivalent to (6.8) with datum $x(0)$, has exactly one continuous solution $x \in C([0, T], \mathcal{H}_P)$, and this solution depends on z_0 only through $[z_0] \in E_{\text{dyn}}$.

Proof. Dependence only on $[z_0]$. The kernel $\mathcal{K}$ is determined by L and P alone, and the only occurrence of z_0 in (7.8) is through $\eta_{z_0} = \mathcal{O}_Q z_0$, which by Theorem 7.4(1) depends only on the class $[z_0]$.

Existence and uniqueness. Fix $T > 0$. Equation (7.8) has the form $x = g + \mathcal{V}x$ on the Banach space $X := C([0, T], \mathcal{H}_P)$ with the sup-norm, where $g(t) := x(0) + \int_0^t \eta_{z_0}(s) ds$ is continuous, and $\mathcal{V}$ is the bounded linear Volterra operator

$$(\mathcal{V}x)(t) := \int_0^t \left(PLP x(s) + \int_0^s \mathcal{K}(s-r)x(r) dr \right) ds.$$

Both PLP and $\mathcal{K}$ are bounded, with $M := \|PLP\|_{\text{op}} + T \sup_{0 \leq \tau \leq T} \|\mathcal{K}(\tau)\|_{\text{op}} < \infty$ (the supremum is finite because $\mathcal{K}$ is norm-continuous on the compact $[0, T]$). An elementary induction gives the standard Volterra estimate

$$\|(\mathcal{V}^k x)(t)\| \leq \frac{(Mt)^k}{k!} \|x\|_X, \quad 0 \leq t \leq T, \quad k \geq 0,$$

so $\|\mathcal{V}^k\|_{\mathcal{B}(X)} \leq (MT)^k/k! \rightarrow 0$; hence $\mathcal{V}^k$ is a contraction for k large, $I - \mathcal{V}$ is invertible with $(I - \mathcal{V})^{-1} = \sum_{k \geq 0} \mathcal{V}^k$ convergent in operator norm, and $x = (I - \mathcal{V})^{-1}g$ is the unique solution in X [5, Ch. 2, Thm. 3.5]. Since $T > 0$ was arbitrary and solutions on nested intervals agree by uniqueness, x is uniquely determined on $[0, \infty)$. $\square$

Theorems 7.4 and 7.6 together give the precise cost of failed descent. The resolved state alone is insufficient (Theorem 6.4). The resolved state plus *all* of $\mathcal{H}_Q$ is sufficient but wasteful. The resolved state plus E_{dyn} is sufficient (Theorem 7.6) and minimal (Theorem 7.4). This is the exact sense in which the unresolved sector must be “completed” into the observable theory, and it is the exact sense in which only part of it need be.

Proposition 7.7 (Finite-dimensional observability criterion). *Suppose $\dim \mathcal{H}_Q = m < \infty$. Then*

$$N_Q = \bigcap_{k=0}^{m-1} \ker \left(PLQ (QLQ)^k \right), \quad (7.9)$$

and $\dim E_{\text{dyn}} = \text{rank } \mathcal{O}_m$, where $\mathcal{O}_m$ is the observability map

$$\mathcal{O}_m : \mathcal{H}_Q \longrightarrow \mathcal{H}_P^m, \quad \mathcal{O}_m z = \left(PLQ z, PLQ (QLQ)z, \dots, PLQ (QLQ)^{m-1}z \right). \quad (7.10)$$

Proof. Fix $z \in \mathcal{H}_Q$. Since QLQ is bounded, $t \mapsto PLQ e^{tQLQ}z$ is given by the everywhere-convergent power series

$$PLQ e^{tQLQ}z = \sum_{k=0}^{\infty} \frac{t^k}{k!} PLQ (QLQ)^k z, \quad (7.11)$$

which is a real-analytic (indeed entire) $\mathcal{H}_P$ -valued function of t . Such a function vanishes identically on $[0, \infty)$ if and only if all its Taylor coefficients at $t = 0$ vanish, i.e. iff

$$PLQ (QLQ)^k z = 0 \quad \text{for all } k \geq 0. \quad (7.12)$$

(One direction is trivial; for the other, differentiate (7.11) k times at $t = 0$, which is legitimate by uniform convergence of the differentiated series on compacts.)

It remains to see that in (7.12) the range $0 \leq k \leq m - 1$ suffices. By the Cayley–Hamilton theorem applied to the operator QLQ on the m -dimensional space $\mathcal{H}_Q$, $(QLQ)^m$ is a linear combination of $I, QLQ, \dots, (QLQ)^{m-1}$; by induction, every power $(QLQ)^k$ with $k \geq m$ is a linear combination of $I, \dots, (QLQ)^{m-1}$. Hence if $PLQ (QLQ)^j z = 0$ for $0 \leq j \leq m - 1$, then applying PLQ to the corresponding linear combination gives $PLQ (QLQ)^k z = 0$ for every $k \geq m$ as well. This proves (7.9).

Finally, (7.9) says exactly $N_Q = \ker \mathcal{O}_m$, so by rank–nullity $\dim E_{\text{dyn}} = \dim \mathcal{H}_Q - \dim N_Q = \text{rank } \mathcal{O}_m$. $\square$

Remark 7.8 (Relation to Kalman observability). Equation (7.10) is the observability map of linear control theory [6, 15] for the finite-dimensional linear system with state $z \in \mathcal{H}_Q$, state matrix QLQ , and output matrix PLQ : the pair (PLQ, QLQ) is observable exactly when $\text{rank } \mathcal{O}_m = m$, i.e. exactly when $N_Q = \{0\}$, i.e. exactly when *every* unresolved direction leaves a trace in the resolved response and the whole of $\mathcal{H}_Q$ must be retained. The dictionary must be read with care, however. What (7.10) tests is observability of the *unresolved sector through the resolved-response output*. It is *not* a claim that the full source system $(\mathcal{H}, L)$ is observable from the resolved readout, which would additionally require that the resolved sector’s own dynamics and the coupling QLP be taken into account. We assert only what (7.9) states: N_Q , and hence $\dim E_{\text{dyn}}$, is computable by a finite rank condition. The general Hilbert space case has no such finite reduction; there N_Q is defined by (7.2) and closed by Lemma 7.3, and no analogue of Cayley–Hamilton is available.

8. NONLINEAR DYNAMICAL DESCENT

The criterion $PLQ = 0$ has an exact nonlinear counterpart, obtained by replacing the linear quotient map by a surjective submersion and the linear generator by a smooth vector field.

Assumption 8.1. S and T are smooth manifolds, $p : S \rightarrow T$ is a smooth *surjective submersion* (so $Dp_s : T_s S \rightarrow T_{p(s)} T$ is surjective for every $s \in S$, and p is onto), and X is a smooth vector field on S .

The hypothesis that p be a surjective submersion is not cosmetic. It guarantees (i) that T is a genuine smooth manifold and $p^{-1}(t)$ an embedded submanifold for each t , and (ii) the existence of smooth local sections, which is what makes the smoothness argument in Theorem 8.2 work. We make no claim about arbitrary set-theoretic quotients: a quotient of a manifold by an arbitrary equivalence relation need not be a smooth manifold at all, and nothing below should be read as asserting otherwise.

Theorem 8.2 (Nonlinear projectability criterion). *Under Assumption 8.1, there exists a smooth vector field Y on T with*

$$Dp \circ X = Y \circ p, \quad \text{i.e.} \quad Dp_s X_s = Y_{p(s)} \quad \text{for every } s \in S, \quad (8.1)$$

if and only if Dp applied to X is constant along every fiber:

$$p(s) = p(s') \implies Dp_s X_s = Dp_{s'} X_{s'} \quad \text{in } T_{p(s)}T. \quad (8.2)$$

When it exists, Y is unique.

Proof. Necessity. Suppose Y exists and $p(s) = p(s') =: t$. Then by (8.1), $Dp_s X_s = Y_t = Dp_{s'} X_{s'}$, and both sides lie in the same tangent space $T_t T$, so (8.2) holds.

Sufficiency. Assume (8.2). For $t \in T$, use surjectivity of p to pick $s \in p^{-1}(t)$ and define

$$Y_t := Dp_s X_s \in T_t T. \quad (8.3)$$

By (8.2) this is independent of the chosen $s \in p^{-1}(t)$, so $Y : T \rightarrow TT$ is a well-defined section of TT : $Y_t \in T_t T$ by construction. Equation (8.1) holds by definition.

Smoothness. Fix $t_0 \in T$ and $s_0 \in p^{-1}(t_0)$. Because p is a submersion, the local normal form for submersions [8, Thm. 4.12] provides an open neighbourhood $W \ni t_0$ in T and a smooth local section $\sigma : W \rightarrow S$ with $p \circ \sigma = \text{id}_W$ and $\sigma(t_0) = s_0$. On W we then have, from (8.3) applied with $s = \sigma(t) \in p^{-1}(t)$,

$$Y|_W = Dp \circ X \circ \sigma, \quad (8.4)$$

which is a composition of smooth maps (σ smooth, X smooth as a section of TS , $Dp : TS \rightarrow TT$ smooth as the tangent map of a smooth map), hence smooth on W . Different choices of section give the same Y on the overlap by fiber constancy (8.2), so these local formulas patch to a globally well-defined smooth field. Since t_0 was arbitrary, Y is smooth on T .

Uniqueness. If Y and Y' both satisfy (8.1), then for any $t \in T$ pick $s \in p^{-1}(t)$ (possible by surjectivity); then $Y_t = Dp_s X_s = Y'_t$. $\square$

Remark 8.3 (The nonlinear analogue of $PLQ \neq 0$). Failure of (8.2) says: there are two source states s, s' with the same readout $p(s) = p(s')$ whose projected velocities $Dp_s X_s \neq Dp_{s'} X_{s'}$ differ. This is verbatim the phenomenon of Theorem 6.4, where p is the linear map P , the fibers are the affine subspaces $x_0 + \mathcal{H}_Q$, and the projected velocity at $x_0 + z_0$ is $PL(x_0 + z_0) = PLP x_0 + PLQ z_0$, which is fiber-constant precisely when $PLQ = 0$. Thus the linear closure criterion is the linearization, and (8.2) its exact generalization: *the flow descends to the observable manifold exactly when equal present readouts force equal projected velocities*. Nothing in Theorem 8.2 produces a memory kernel; in the nonlinear setting the exact reduction identity of Theorem 6.2 has no elementary counterpart, and the theorem should be read as supplying only the sharp criterion for descent, not a reduction formula.

9. ELIMINATION OF A SOURCE SECTOR: SCHUR-COMPLEMENT EFFECTIVE OPERATORS

We now turn from state-space projection to *field-theoretic elimination*: the source variables split into a visible field φ and a completion field χ , and χ is removed by solving its own stationarity condition and substituting back. In the quadratic case this can be done exactly, and the result is a Schur complement.

Assumption 9.1. $\mathcal{H}_\varphi$ and $\mathcal{H}_\chi$ are (real or complex) Hilbert spaces; $A \in \mathcal{B}(\mathcal{H}_\varphi)$ and $C \in \mathcal{B}(\mathcal{H}_\chi)$ are self-adjoint; $B \in \mathcal{B}(\mathcal{H}_\chi, \mathcal{H}_\varphi)$, with adjoint $B^* \in \mathcal{B}(\mathcal{H}_\varphi, \mathcal{H}_\chi)$; and C is boundedly invertible, i.e. $C^{-1} \in \mathcal{B}(\mathcal{H}_\chi)$ exists.

Define the quadratic action functional

$$S[\varphi, \chi] := \frac{1}{2} \langle \varphi, A\varphi \rangle + \operatorname{Re} \langle \varphi, B\chi \rangle + \frac{1}{2} \langle \chi, C\chi \rangle, \quad (\varphi, \chi) \in \mathcal{H}_\varphi \times \mathcal{H}_\chi. \quad (9.1)$$

(In the real case Re is vacuous and we simply write $\langle \varphi, B\chi \rangle$; we keep it so that S is real-valued in the complex case, which is what a physical action must be. Self-adjointness of A and C makes the first and third terms real automatically.)

Theorem 9.2 (Schur-complement effective action). *Under Assumption 9.1, for each fixed $\varphi \in \mathcal{H}_\varphi$ the stationarity condition*

$$\frac{\partial S}{\partial \chi} = 0 \quad \Longleftrightarrow \quad C\chi + B^*\varphi = 0 \quad (9.2)$$

has the unique solution

$$\chi_* = \chi_*[\varphi] = -C^{-1}B^*\varphi, \quad (9.3)$$

and substitution yields the exact reduced action

$$\boxed{S_{\text{eff}}[\varphi] := S[\varphi, \chi_*[\varphi]] = \frac{1}{2} \langle \varphi, (A - BC^{-1}B^*)\varphi \rangle.} \quad (9.4)$$

The operator $A - BC^{-1}B^* \in \mathcal{B}(\mathcal{H}_\varphi)$ is self-adjoint; it is the Schur complement of C in the block operator matrix $\begin{pmatrix} A & B \\ B^* & C \end{pmatrix}$.

Proof. The χ -equation. Fix φ and let $\delta \in \mathcal{H}_\chi$. Using $C = C^*$,

$$\left. \frac{d}{d\varepsilon} \right|_{\varepsilon=0} S[\varphi, \chi + \varepsilon\delta] = \operatorname{Re} \langle \varphi, B\delta \rangle + \operatorname{Re} \langle \chi, C\delta \rangle = \operatorname{Re} \langle B^*\varphi + C\chi, \delta \rangle.$$

(For the middle term: $\frac{1}{2} \langle \chi + \varepsilon\delta, C(\chi + \varepsilon\delta) \rangle$ has ε -derivative at 0 equal to $\frac{1}{2} (\langle \delta, C\chi \rangle + \langle \chi, C\delta \rangle) = \operatorname{Re} \langle \chi, C\delta \rangle$ by self-adjointness.) This vanishes for all δ (taking also $i\delta$ in the complex case) iff $C\chi + B^*\varphi = 0$, which is (9.2). As C is boundedly invertible, the unique solution is (9.3), and χ_* depends boundedly and linearly on φ .

Substitution. Insert (9.3) into (9.1). Term by term:

$$\operatorname{Re} \langle \varphi, B\chi_* \rangle = \operatorname{Re} \langle \varphi, -BC^{-1}B^*\varphi \rangle = -\langle \varphi, BC^{-1}B^*\varphi \rangle,$$

the last step because $BC^{-1}B^*$ is self-adjoint: indeed

$$(BC^{-1}B^*)^* = B(C^{-1})^*B^* = B(C^*)^{-1}B^* = BC^{-1}B^*,$$

using $C^* = C$ and $(C^{-1})^* = (C^*)^{-1}$, so that $\langle \varphi, BC^{-1}B^*\varphi \rangle$ is real. Next,

$$\frac{1}{2} \langle \chi_*, C\chi_* \rangle = \frac{1}{2} \langle -C^{-1}B^*\varphi, C(-C^{-1}B^*\varphi) \rangle = \frac{1}{2} \langle C^{-1}B^*\varphi, B^*\varphi \rangle = \frac{1}{2} \langle \varphi, BC^{-1}B^*\varphi \rangle,$$

where the last equality moves B^* across as B and uses self-adjointness of C^{-1} , which makes the quantity $\langle C^{-1}B^*\varphi, B^*\varphi \rangle$ real and equal to $\langle \varphi, BC^{-1}B^*\varphi \rangle$. Summing,

$$S_{\text{eff}}[\varphi] = \frac{1}{2} \langle \varphi, A\varphi \rangle - \langle \varphi, BC^{-1}B^*\varphi \rangle + \frac{1}{2} \langle \varphi, BC^{-1}B^*\varphi \rangle = \frac{1}{2} \langle \varphi, (A - BC^{-1}B^*)\varphi \rangle,$$

which is (9.4). Self-adjointness of $A - BC^{-1}B^*$ was verified above. $\square$

Remark 9.3 (Nonlocality of the effective operator). The correction to the visible operator is

$$\Delta := -BC^{-1}B^*, \quad (9.5)$$

and it is derived—not postulated. What the theorem establishes is exactly the *form* of Δ ; whether Δ is nonlocal is a further question about the composed operator, and must not be read off from C^{-1} alone. In field theory, when $\mathcal{H}_\chi$ is a space of functions on a spacetime region and C is (the bounded realization of) a differential operator, C^{-1} is an integral operator with a Green kernel $G(y, y')$, and

$$(\Delta\varphi)(x) = - \int \int \int B(x, y) G(y, y') B^*(y', x') \varphi(x') dy dy' dx'.$$

This composition *can* be *nonlocal*: the effective equation for φ at a point then involves φ elsewhere, weighted by the Green function of the eliminated sector—even though the coupled source equations for (φ, χ) were strictly local.

The qualification is not pedantic, because nonlocality of C^{-1} is *not* sufficient. The correction depends on the coupling as well: if $B = 0$ there is no correction at all, and more generally B may fail to couple to the nonlocal directions of G , or the two compositions B and B^* may conspire so that $BC^{-1}B^*$ collapses to a local (e.g. multiplication or differential) operator. A concrete instance: if B is a rank-one coupling $B = |f\rangle\langle h|$, then $BC^{-1}B^* = \langle h, C^{-1}h \rangle |f\rangle\langle f|$, whose kernel is the product $\overline{f(x')}f(x)$ times a single scalar—the whole Green function G has entered only through the number $\langle h, C^{-1}h \rangle$. We stress what this example does and does not show. The rank-one operator $|f\rangle\langle f|$ is itself generally *nonlocal* (its kernel $\overline{f(x')}f(x)$ does not concentrate on the diagonal), so the example does not exhibit a local effective operator; what it exhibits is that the detailed structure of G —the very thing one might expect to be imprinted on the correction—can collapse to a single number, leaving a correction whose nonlocal structure is dictated entirely by B and not by the eliminated sector’s propagator. Nonlocality of the effective operator is therefore a property of $BC^{-1}B^*$, requiring that B couple to the relevant Green-operator directions without cancellation; it is an additional verification, not a corollary of Theorem 9.2.

Subject to that verification, elimination of a sector converts locality-with-extra-fields into nonlocality-with-fewer-fields, and there is an exact structural analogy with the memory kernel of Theorem 6.2: there, eliminating $\mathcal{H}_Q$ converted a first-order ODE into an equation nonlocal *in time*, and there too the kernel $\mathcal{K} = PLQ e^{\tau QLQ} QLP$ can degenerate—it vanishes identically when $PLQ = 0$ or $QLP = 0$ —so that nonlocality in time likewise depends on the couplings and not on the eliminated sector’s propagator alone. In both settings nonlocality is a possible price of the reduction, not an automatic one, and never a feature of the source.

Remark 9.4 (Unbounded operators, and what would have to be supplied). Theorem 9.2 is stated for *bounded* A , B , C with C boundedly invertible, and the proof uses boundedness at every step (in the differentiation, in the existence of C^{-1} as a bounded operator, in the adjoint manipulations). It does *not* automatically cover the case where A and C are unbounded differential operators, which is the case of primary interest in field theory. To extend it one must supply, and verify: a common dense domain, or a formulation in terms of closed, densely defined, semibounded quadratic forms with form domains [7, Ch. VI]; boundary conditions making C self-adjoint and rendering the integration by parts in the variation legitimate; invertibility of C (equivalently, that $0 \notin \sigma(C)$, which fails whenever the eliminated sector has zero modes, e.g. a massless field on a compact region with the wrong boundary conditions); and the boundedness of C^{-1} from the form dual back to the form domain, so that $BC^{-1}B^*$ is defined on the form domain of A . None of this follows from the bounded proof, and we do not assert it here.

Where the physics literature performs such eliminations with unbounded operators, the Green operator $C^{-1} = G$ is being used, and the validity of the manipulation is exactly the validity of the associated elliptic/hyperbolic solvability theory, which must be supplied case by case.

Remark 9.5 (Tree level, not the quantum effective action). Theorem 9.2 is a statement of *classical stationary elimination*: χ is replaced by the solution of its classical field equation (9.2). Equivalently, in a functional-integral language, (9.4) is the tree-level (saddle-point) contribution. It is *not* the quantum effective action for φ . Integrating χ out of a Gaussian functional integral with weight $e^{iS[\varphi, \chi]}$ produces, in addition to (9.4), a φ -independent determinant factor $(\det C)^{-1/2}$; if C depends on φ (a nonquadratic coupling, absent from Assumption 9.1) the determinant becomes φ -dependent and contributes a genuine one-loop term $\frac{i}{2} \text{Tr} \log C[\varphi]$; and beyond one loop there are further corrections [12, 17]. Nothing in Theorem 9.2 computes or bounds these. The theorem is exactly what it says: substitution of a stationary point into a quadratic functional.

10. EFFECTIVE METRIC RESPONSE

We finally allow the metric to participate, and ask what elimination of χ does to the gravitational field equations. The answer is a stationary-elimination (envelope) identity, which converts the eliminated sector into a *covariant variational response* rather than an unexplained additive term.

Assumption 10.1. Let M be a smooth spacetime manifold. Let the metrics, visible fields, and completion fields range over open subsets $\mathcal{G}, \mathcal{F}_\varphi, \mathcal{F}_\chi$ of Banach spaces (equivalently, over Banach manifolds modelled on them—for definiteness one may take suitable Sobolev or weighted Hölder sections of the relevant bundles over M , with the boundary or fall-off conditions specified below built into the space). Let

$$S : \mathcal{G} \times \mathcal{F}_\varphi \times \mathcal{F}_\chi \longrightarrow \mathbb{R}$$

be a real, diffeomorphism-invariant action functional which is *continuously Fréchet differentiable* (C^1) on its domain, and suppose the boundary/fall-off conditions defining the spaces are such that all integrations by parts performed below produce no boundary contributions, and such that the Fréchet derivatives are represented by the densities $\delta S / \delta g^{\mu\nu}$, $\delta S / \delta \varphi$, $\delta S / \delta \chi$ in the usual way,

$$DS[g, \varphi, \chi](h, \psi, \theta) = \int_M \left(\frac{\delta S}{\delta g^{\mu\nu}} h^{\mu\nu} + \frac{\delta S}{\delta \varphi} \psi + \frac{\delta S}{\delta \chi} \theta \right) d^4x.$$

Suppose a *branch*

$$(g, \varphi) \longmapsto \chi_*[g, \varphi] \tag{10.1}$$

of solutions of the eliminated field equation

$$\left. \frac{\delta S}{\delta \chi(x)} \right|_{\chi=\chi_*[g, \varphi]} = 0 \quad \text{for all } x \in M \tag{10.2}$$

has been selected, and that this branch is a *continuously Fréchet differentiable* map

$$\chi_* : \mathcal{G} \times \mathcal{F}_\varphi \longrightarrow \mathcal{F}_\chi.$$

(In practice such a branch is produced, where it exists, by the implicit function theorem in Banach spaces [18, Ch. 4], applied to the Euler map $(g, \varphi, \chi) \mapsto \delta S / \delta \chi$ at a solution where the linearization $\delta^2 S / \delta \chi^2$ is a Banach-space isomorphism. That theorem requires the Euler map itself to be continuously Fréchet differentiable, which does *not* follow from $S \in C^1$; so for this remark—and for this remark only—assume additionally that $S \in C^2$, or equivalently for present purposes that $(g, \varphi, \chi) \mapsto \delta S / \delta \chi$ is itself continuously Fréchet differentiable. Under that

extra hypothesis the implicit function theorem also delivers the C^1 dependence of χ_* assumed above. We emphasize that this strengthening is *not* imposed on Assumption 10.1 or used in Proposition 10.2: the envelope identity needs only $S \in C^1$ together with the separately assumed C^1 branch χ_* , however that branch is obtained. We do not verify any of these hypotheses for a particular action; we assume them.) Define the *reduced action*

$$\Gamma_{\text{eff}}[g, \varphi] := S[g, \varphi, \chi_*[g, \varphi]]. \quad (10.3)$$

We stress the three hypotheses that are doing real work and that are frequently left tacit: (i) *differentiability of the branch* χ_* —a solution of (10.2) need not depend differentiably on (g, φ) , and where several branches exist a choice has been made; (ii) *boundary conditions*—the variation $\delta\chi_*$ induced by a variation of g must be an admissible variation, i.e. satisfy the same boundary conditions, or the envelope argument fails; (iii) *equivariance of the branch*—the assignment (10.1) must intertwine the diffeomorphism action, $\chi_*[\phi^*g, \phi^*\varphi] = \phi^*\chi_*[g, \varphi]$ for diffeomorphisms ϕ , if Γ_{eff} is to inherit the diffeomorphism invariance of S . Without (iii), S can be invariant while Γ_{eff} is not, because the selection of the branch may break the symmetry.

Proposition 10.2 (Stationary-elimination / envelope identity). *Under Assumption 10.1,*

$$\boxed{\frac{\delta\Gamma_{\text{eff}}[g, \varphi]}{\delta g^{\mu\nu}(x)} = \frac{\delta S[g, \varphi, \chi]}{\delta g^{\mu\nu}(x)} \Big|_{\chi=\chi_*[g, \varphi]}} \quad (10.4)$$

and likewise $\frac{\delta\Gamma_{\text{eff}}}{\delta\varphi(x)} = \frac{\delta S}{\delta\varphi(x)} \Big|_{\chi_*}$. That is, the metric (and visible-field) derivative of the reduced action is computed by differentiating S at fixed χ and then evaluating on the branch: the implicit dependence of χ_* on g contributes nothing.

Proof. Since S and χ_* are continuously Fréchet differentiable, so is the composite $\Gamma_{\text{eff}} = S \circ (\text{id}, \text{id}, \chi_*)$, and the standard Fréchet chain rule for a composition of continuously Fréchet differentiable maps between Banach spaces [18, Ch. 4] applies without further hypotheses:

$$D\Gamma_{\text{eff}}[g, \varphi] = D_{g, \varphi} S|_{\chi_*} + D_{\chi} S|_{\chi_*} \circ D\chi_*[g, \varphi].$$

Writing this in terms of the representing densities—with the spacetime integration and field indices explicit, and with χ carrying whatever index structure I it possesses—it reads

$$\frac{\delta\Gamma_{\text{eff}}[g, \varphi]}{\delta g^{\mu\nu}(x)} = \frac{\delta S}{\delta g^{\mu\nu}(x)} \Big|_{\chi_*} + \int_M d^4y \frac{\delta S}{\delta \chi^I(y)} \Big|_{\chi_*} \frac{\delta \chi_*^I(y)}{\delta g^{\mu\nu}(x)}, \quad (10.5)$$

with summation over I . The induced variation $\delta\chi_*$ is an admissible variation—it lands in the tangent space to $\mathcal{F}_\chi$, hence satisfies the boundary conditions built into that space—which is what licenses writing the second term as an integral against $\delta S/\delta \chi^I$ with no boundary contribution. By the eliminated field equation (10.2), the factor $\delta S/\delta \chi^I(y)|_{\chi_*}$ vanishes identically in y , so the entire integral in (10.5) vanishes, leaving (10.4). The argument for $\delta/\delta\varphi$ is identical, replacing $g^{\mu\nu}(x)$ by $\varphi(x)$. $\square$

Proposition 10.2 is what makes the reduction *usable*: to compute the gravitational response of the reduced theory one does not need $\delta\chi_*/\delta g$ at all—an object which is typically the solution of a linearized equation and thoroughly inaccessible.

Definition 10.3 (Effective stress-energy). Split the reduced action into a visible part and a correction,

$$\Gamma_{\text{eff}}[g, \varphi] = S_{\text{visible}}[g, \varphi] + \Gamma_{\text{corr}}[g, \varphi], \quad (10.6)$$

where, by declaration, S_{visible} contains the gravitational sector and the visible matter sector in the conventions

$$S_{\text{visible}}[g, \varphi] = \frac{1}{16\pi G_N} \int_M (R - 2\Lambda) \sqrt{-g} d^4x + S_{\text{mat}}[g, \varphi], \quad T_{\mu\nu}^{\text{visible}} := -\frac{2}{\sqrt{-g}} \frac{\delta S_{\text{mat}}}{\delta g^{\mu\nu}}, \quad (10.7)$$

and $\Gamma_{\text{corr}} := \Gamma_{\text{eff}} - S_{\text{visible}}$ is whatever the elimination of χ has added. We assume, as part of the declaration (10.7), that the visible matter action S_{mat} is itself diffeomorphism invariant; the Einstein–Hilbert term is invariant automatically, so S_{visible} is an invariant functional. Consequently, whenever Γ_{eff} is invariant—which, by Assumption 10.1 and the equivariance of the branch, is the case under the hypotheses of Theorem 10.5—the correction $\Gamma_{\text{corr}} = \Gamma_{\text{eff}} - S_{\text{visible}}$ is a difference of invariant functionals and is therefore separately invariant as well. This is what entitles us to call *both* $T_{\mu\nu}^{\text{visible}}$ and $T_{\mu\nu}^{\text{eff}}$ covariant tensorial responses in Remark 10.6: each is the metric variation of an invariant functional, and each therefore satisfies the equivariance law (10.17). Separate *invariance* of the two functionals must not, however, be confused with separate *conservation* of the two tensors, which Theorem 10.5 does *not* assert; the two sectors do each carry an off-shell Noether identity, but they are coupled on shell by an exchange current, as the proof of that theorem makes precise. Define the *effective stress-energy tensor* of the eliminated sector by the same variational rule:

$$T_{\mu\nu}^{\text{eff}} := -\frac{2}{\sqrt{-g}} \frac{\delta \Gamma_{\text{corr}}}{\delta g^{\mu\nu}}. \quad (10.8)$$

The point of (10.8) is that it is a *definition by variation of an explicitly constructed functional* Γ_{corr} , and (10.6) tells us exactly what that functional is. Contrast the common alternative: take a field equation one does not like, move the offending term to the right-hand side of Einstein’s equation, and call the result a stress-energy tensor. That rearrangement is not empty—it does define a symbol, and algebraically the resulting equation is true by construction—but it establishes, by itself, none of the three properties one wants of a stress-energy tensor: it exhibits no variational origin (no functional whose metric derivative the term is), it does not by itself certify that the object transforms as a $(0, 2)$ -tensor, and it yields no conservation law. Definition 10.3 supplies all three: the functional is Γ_{corr} from (10.6), the tensorial transformation follows because T^{eff} is the metric variation of a diffeomorphism-invariant functional (Remark 10.6), and the conservation law is Theorem 10.5. That is the difference between a rearrangement and a derivation.

Proposition 10.4 (Effective Einstein equations). *With the conventions (10.7), stationarity of Γ_{eff} with respect to $g^{\mu\nu}$, i.e. $\delta \Gamma_{\text{eff}} / \delta g^{\mu\nu} = 0$, is equivalent to*

$$G_{\mu\nu} + \Lambda g_{\mu\nu} = 8\pi G_N \left(T_{\mu\nu}^{\text{visible}} + T_{\mu\nu}^{\text{eff}} \right), \quad (10.9)$$

where $G_{\mu\nu} = R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu}$.

Proof. Vary (10.6). The Einstein–Hilbert term in (10.7) has the standard variation

$$\frac{\delta}{\delta g^{\mu\nu}} \frac{1}{16\pi G_N} \int (R - 2\Lambda) \sqrt{-g} d^4x = \frac{\sqrt{-g}}{16\pi G_N} \left(R_{\mu\nu} - \frac{1}{2} R g_{\mu\nu} + \Lambda g_{\mu\nu} \right) = \frac{\sqrt{-g}}{16\pi G_N} (G_{\mu\nu} + \Lambda g_{\mu\nu}),$$

using $\delta \sqrt{-g} = -\frac{1}{2} \sqrt{-g} g_{\mu\nu} \delta g^{\mu\nu}$ and discarding the total-derivative term by the boundary hypothesis [16, §4.3]. By (10.7) and (10.8) the matter and correction terms contribute $-\frac{\sqrt{-g}}{2} T_{\mu\nu}^{\text{visible}}$ and $-\frac{\sqrt{-g}}{2} T_{\mu\nu}^{\text{eff}}$ respectively. Setting the sum to zero and multiplying by $16\pi G_N / \sqrt{-g}$ gives (10.9). $\square$

Theorem 10.5 (Total covariant conservation). *Suppose in addition to Assumption 10.1 that the branch χ_* is diffeomorphism-equivariant, so that Γ_{eff} is a diffeomorphism-invariant functional of (g, φ) . Then, on solutions of the effective visible-matter field equation $\delta\Gamma_{\text{eff}}/\delta\varphi = 0$,*

$$\nabla^\mu (T_{\mu\nu}^{\text{visible}} + T_{\mu\nu}^{\text{eff}}) = 0. \quad (10.10)$$

In general, neither $\nabla^\mu T_{\mu\nu}^{\text{visible}} = 0$ nor $\nabla^\mu T_{\mu\nu}^{\text{eff}} = 0$ holds separately.

Proof. Let ξ be a compactly supported smooth vector field on M and let ϕ_ε be its flow. Diffeomorphism invariance of Γ_{eff} gives $\Gamma_{\text{eff}}[\phi_\varepsilon^* g, \phi_\varepsilon^* \varphi] = \Gamma_{\text{eff}}[g, \varphi]$ for all ε ; differentiating at $\varepsilon = 0$ and using $\delta_\xi g^{\mu\nu} = \mathcal{L}_\xi g^{\mu\nu} = -(\nabla^\mu \xi^\nu + \nabla^\nu \xi^\mu)$ and $\delta_\xi \varphi = \mathcal{L}_\xi \varphi$, we obtain the Noether identity

$$0 = \int_M \left(\frac{\delta\Gamma_{\text{eff}}}{\delta g^{\mu\nu}} \mathcal{L}_\xi g^{\mu\nu} + \frac{\delta\Gamma_{\text{eff}}}{\delta \varphi} \mathcal{L}_\xi \varphi \right) d^4x. \quad (10.11)$$

Restrict now to field configurations satisfying the visible-matter equation of the *reduced* theory, $\delta\Gamma_{\text{eff}}/\delta\varphi = 0$; the second term of (10.11) then drops. Split $\delta\Gamma_{\text{eff}}/\delta g^{\mu\nu}$ using (10.6), (10.7), (10.8): writing

$$\frac{\delta\Gamma_{\text{eff}}}{\delta g^{\mu\nu}} = \frac{\sqrt{-g}}{16\pi G_N} (G_{\mu\nu} + \Lambda g_{\mu\nu}) - \frac{\sqrt{-g}}{2} (T_{\mu\nu}^{\text{visible}} + T_{\mu\nu}^{\text{eff}}) =: \frac{\sqrt{-g}}{16\pi G_N} \mathcal{E}_{\mu\nu},$$

and substituting into (10.11) with $\mathcal{L}_\xi g^{\mu\nu} = -2\nabla^{(\mu} \xi^{\nu)}$ gives

$$0 = -\frac{2}{16\pi G_N} \int_M \sqrt{-g} \mathcal{E}_{\mu\nu} \nabla^\mu \xi^\nu d^4x = \frac{2}{16\pi G_N} \int_M \sqrt{-g} (\nabla^\mu \mathcal{E}_{\mu\nu}) \xi^\nu d^4x,$$

after one covariant integration by parts, the boundary term vanishing because ξ has compact support. As ξ is arbitrary, $\nabla^\mu \mathcal{E}_{\mu\nu} = 0$. Finally, the contracted Bianchi identity gives $\nabla^\mu (G_{\mu\nu} + \Lambda g_{\mu\nu}) = 0$ identically (metric-compatibility making $\nabla^\mu g_{\mu\nu} = 0$), so $\nabla^\mu (T_{\mu\nu}^{\text{visible}} + T_{\mu\nu}^{\text{eff}}) = 0$, which is (10.10).

For the final sentence, the situation must be described with care, because S_{visible} and Γ_{corr} are separately diffeomorphism invariant (Definition 10.3), and each therefore *does* possess its own off-shell Noether identity, obtained by running the argument above on that functional alone. To state these without committing to the tensor type of φ , introduce the *Ward operator* $\mathcal{W}_\nu$: for any diffeomorphism-invariant functional $F[g, \varphi]$ with associated stress tensor $T_{\mu\nu}^F := -\frac{2}{\sqrt{-g}} \delta F / \delta g^{\mu\nu}$ and Euler operator $E_\varphi^F := \delta F / \delta \varphi$, the same computation as above—with $\delta_\xi \varphi = \mathcal{L}_\xi \varphi$ and one integration by parts—yields the off-shell identity

$$\nabla^\mu T_{\mu\nu}^F = \mathcal{W}_\nu (E_\varphi^F), \quad (10.12)$$

where $\mathcal{W}_\nu$ is the formal adjoint of the map $\xi^\nu \mapsto \mathcal{L}_\xi \varphi$, contracted appropriately; it is *linear* in its argument, and its precise form depends on the tensor type of φ but *not* on F . For a scalar field φ it reduces to

$$\mathcal{W}_\nu (E_\varphi) = \frac{1}{\sqrt{-g}} E_\varphi \nabla_\nu \varphi, \quad \text{so that} \quad \nabla^\mu T_{\mu\nu} = \frac{1}{\sqrt{-g}} \frac{\delta F}{\delta \varphi} \nabla_\nu \varphi, \quad (10.13)$$

in the conventions $T_{\mu\nu} = -\frac{2}{\sqrt{-g}} \delta F / \delta g^{\mu\nu}$ and $\delta_\xi \varphi = \mathcal{L}_\xi \varphi$ fixed in (10.7) and above. We stress that (10.13) is the *scalar* case; (10.12) is the statement for general φ , and the scalar formula must not be read as an identity for tensor-valued φ .

Applying (10.12) to S_{visible} and to Γ_{corr} —each invariant, hence each entitled to its own identity—and using the contracted Bianchi identity on the gravitational part of S_{visible} , we

obtain, as identities holding for *all* field configurations,

$$\nabla^\mu T_{\mu\nu}^{\text{visible}} = \mathcal{W}_\nu \left(\frac{\delta S_{\text{visible}}}{\delta \varphi} \right) =: J_\nu^{\text{vis}}, \quad (10.14)$$

$$\nabla^\mu T_{\mu\nu}^{\text{eff}} = \mathcal{W}_\nu \left(\frac{\delta \Gamma_{\text{corr}}}{\delta \varphi} \right) =: J_\nu^{\text{corr}}. \quad (10.15)$$

So the two sectors do have Ward identities, and the reason separate conservation nevertheless fails is not an absence of identities but the *on-shell* condition. The physical configuration satisfies the field equation of the *combined* reduced theory,

$$\frac{\delta \Gamma_{\text{eff}}}{\delta \varphi} = \frac{\delta S_{\text{visible}}}{\delta \varphi} + \frac{\delta \Gamma_{\text{corr}}}{\delta \varphi} = 0, \quad (10.16)$$

which does *not* imply $\delta S_{\text{visible}}/\delta \varphi = 0$ or $\delta \Gamma_{\text{corr}}/\delta \varphi = 0$ separately—it implies only that the two Euler operators are equal and opposite. Since $\mathcal{W}_\nu$ is linear, applying it to (10.16) gives

$$J_\nu^{\text{vis}} + J_\nu^{\text{corr}} = \mathcal{W}_\nu \left(\frac{\delta S_{\text{visible}}}{\delta \varphi} + \frac{\delta \Gamma_{\text{corr}}}{\delta \varphi} \right) = \mathcal{W}_\nu(0) = 0,$$

so the two sector Ward terms are equal and opposite and cancel in the sum of (10.14) and (10.15)—which is exactly (10.10)—while neither need vanish on its own. The common quantity $J_\nu^{\text{vis}} = -J_\nu^{\text{corr}}$ is the *exchange current* between the visible and eliminated sectors: it is the rate at which the eliminated sector feeds energy–momentum to the visible one.

The logical structure is therefore this. Separate diffeomorphism invariance supplies separate *off-shell* identities, (10.14) and (10.15), and nothing more. Separate *conservation* of a sector’s stress tensor requires, in addition, that that sector’s Ward term vanish *on the combined solution*—i.e. that $\mathcal{W}_\nu(\delta S_{\text{visible}}/\delta \varphi) = 0$ there, equivalently (by (10.16) and linearity) that $\mathcal{W}_\nu(\delta \Gamma_{\text{corr}}/\delta \varphi) = 0$ there. This holds, for instance, when Γ_{corr} does not depend on φ at all, or when the visible field happens to satisfy its own equation $\delta S_{\text{visible}}/\delta \varphi = 0$ on the combined solution, or when $\mathcal{W}_\nu$ annihilates the relevant Euler operator. It is *not* supplied by invariance: diffeomorphism invariance of a sector gives that sector its Ward identity, but it does not put the sector on shell, and it therefore cannot by itself force the exchange current to vanish. Absent such a structure, the eliminated sector exchanges energy–momentum with the visible one, and only the sum is conserved. $\square$

Remark 10.6 (Invariance versus covariance). Three statements must be kept apart, and they are frequently run together.

- (1) Γ_{eff} is an *invariant scalar functional*: under the hypotheses of Theorem 10.5 we have $\Gamma_{\text{eff}}[\phi^* g, \phi^* \varphi] = \Gamma_{\text{eff}}[g, \varphi]$. This is a statement about a single number assigned to a whole field configuration.
- (2) $T_{\mu\nu}^{\text{eff}}$ is a *covariant tensorial response*: it transforms equivariantly as a $(0, 2)$ -tensor field under diffeomorphisms, because it is the metric variation of an invariant functional. Geometrically it is a section of the associated tensor bundle $\text{Sym}^2 T^* M$, evaluated on the given configuration, and the diffeomorphism group acts on it by pullback rather than fixing it:

$$T^{\text{eff}}[\phi^* g, \phi^* \varphi] = \phi^* T^{\text{eff}}[g, \varphi], \quad (10.17)$$

in the same pullback convention used for the metric and the fields throughout. Consequently it is *not* an ordinary pointwise scalar function on the naive quotient of field space by diffeomorphisms: the value “ $T_{\mu\nu}^{\text{eff}}(x)$ ” is not by itself a diffeomorphism-invariant number, because the label x is not. Three levels must be kept apart here, and the second is routinely mistaken for the third.

- *Tensorial response.* $T_{\mu\nu}^{\text{eff}}$ itself: an equivariant section of $\text{Sym}^2 T^*M$, transforming by pullback as in (10.17).
- *Covariant scalar fields obtained by contraction.* Objects such as the trace $T^{\text{eff}} := g^{\mu\nu} T_{\mu\nu}^{\text{eff}}$, or $T_{\mu\nu}^{\text{eff}} T^{\text{eff}\mu\nu}$. These are scalar *fields*: they carry no free indices and transform as scalars. But a scalar field is still a function on M , and it is *still not* a diffeomorphism-invariant observable when evaluated at a coordinate-labelled point: under ϕ it goes to $\phi^* T^{\text{eff}}$, so the number $T^{\text{eff}}(x)$ at a fixed coordinate label x is not invariant. Contraction removes the indices; it does not remove the dependence on the labelling.
- *Genuine diffeomorphism-invariant observables.* These require a further step: invariant integration (e.g. $\int_M T^{\text{eff}} \sqrt{-g} d^4x$ over an invariantly specified region, or an integral against an invariantly constructed density), smearing against equivariantly chosen test tensors, or relational evaluation—evaluating one field at a point located by the values of other fields, so that the point is specified by the configuration rather than by a coordinate.

The equivariant response is the primary object; covariant scalar fields are contractions of it; invariant observables are constructed from those by one of the three routes just named.

- (3) The effective field equations (10.9) are *geometrically covariant*: both sides are $(0, 2)$ -tensors and the equation is preserved under diffeomorphisms. Covariance of the equations is a consequence of invariance of the functional; it is not the same statement, and it does not license treating T^{eff} as an invariant.

11. EXACT VERSUS APPROXIMATE REDUCTION

It is worth collecting, in one place, what has been proved exactly and where an approximation would have to enter.

Exact. Theorem 2.3 (the quotient and the closed-range criterion), Proposition 3.2 and Corollary 3.3 (functional descent), Theorem 3.4 (stable observability of a linear functional, i.e. boundedness of the descended functional in the inherited norm iff $a \in \text{Ran } K^*$), Theorem 4.4 (the truncation error, which is an *exact evaluation* of an approximation's cost), Theorem 5.2 (descent iff invariance), Theorem 6.2 (the projected identity), Theorem 6.4 (closure iff $PLQ = 0$), Theorems 7.4 and 7.6 (minimal completion and Volterra closure), Proposition 7.7 (finite computation of N_Q), Theorem 8.2 (nonlinear projectability), Theorem 9.2 (Schur complement), Proposition 10.2 and Theorem 10.5 (envelope identity and total conservation). Apart from the explicitly identified rank- r approximant in the Eckart–Young–Mirsky result, these identities introduce no approximation, limiting procedure, ensemble, or small parameter; the theorem itself evaluates the truncation error exactly.

Where approximation would enter. Four places, and it is important that they are four and not one.

- (1) *Rank truncation.* Replacing K by K_r costs exactly σ_{r+1} in operator norm (Theorem 4.4). The cost is quantified, the reduction is not exact, and it is completely unrelated to $\ker K$.
- (2) *Regularization of unstable inversion.* When $\text{Ran } K$ is not closed, recovering the source from the readout requires a regularization scheme (Tikhonov, spectral cutoff), with a bias–variance trade-off governed by the spectral behaviour of K^*K near 0 [4]. When K is compact, that behaviour is encoded in the discrete sequence $\sigma_n \rightarrow 0$ and the trade-off is classically described by the rate of decay of the σ_n ; for a general bounded

operator with nonclosed range there need be no discrete singular sequence at all— K^*K may have continuous spectrum accumulating at 0—and it is the spectral measure of K^*K near the origin, not a sequence, that governs the scheme. This is in any case an approximation of the *inverse*, not of K .

- (3) *Truncating the memory.* Approximating (6.8) by an autonomous equation—dropping η_{z_0} and $\mathcal{K}$, or replacing $\mathcal{K}$ by a delta function (the Markov approximation)—is legitimate only under an additional hypothesis. The property actually needed is *exponential stability of the unresolved propagator*: constants $M \geq 1$, $\alpha > 0$ with $\|e^{\tau QLQ}\|_{\text{op}} \leq Me^{-\alpha\tau}$, so that $\|\mathcal{K}(\tau)\| \leq M\|PLQ\|\|QLP\|e^{-\alpha\tau}$ and the convolution has an effective memory time of order $1/\alpha$; this is then normally accompanied by a quantitative separation of timescales, α^{-1} small compared with the timescale on which x varies (governed by $\|PLP\|$ and the coupling). These are prerequisites, not a proof. Exponential decay of $e^{\tau QLQ}$ alone does *not* establish that replacing $\mathcal{K}$ by a delta function is controlled: that step additionally requires quantitative control of the kernel’s moments—at minimum $\int_0^\infty \|\mathcal{K}(\tau)\| d\tau$ and $\int_0^\infty \tau \|\mathcal{K}(\tau)\| d\tau$, the second being what multiplies the error made in replacing $x(t-\tau)$ by $x(t)$ —together with a bound on the variation of x over the memory timescale, and then an explicit error estimate or a limiting argument ($\alpha \rightarrow \infty$ at fixed $\int \mathcal{K}$, say) carried out in the concrete model at hand. We supply none of this here, and it should not be assumed to be automatic. Note also that a large norm $\|QLQ\|$ implies nothing of the sort by itself—the norm bounds $\|e^{\tau QLQ}\| \leq e^{\tau\|QLQ\|}$ from above only, and a large-norm generator may have purely oscillatory or even growing propagator. What matters is the location of the spectrum (or, in infinite dimensions, the growth bound of the semigroup), not the size of the generator. We have made no such hypothesis, and Theorem 6.4 shows that without $PLQ = 0$ such a truncation is never exact.
- (4) *Loop corrections.* (9.4) and (10.3) are classical stationary eliminations (Remark 9.5). A quantum treatment produces determinant and loop terms.

The physical reading. What Sections 5–7 establish is this. A model may have a completely coherent observable theory—an honest Hilbert space of observable states, an honest algebra of observable functionals, an honest exact quotient—and still fail to be *autonomous* on that observable state. When it does fail, the exact projected law records precisely three things beyond the naive autonomous term $PLP x$: dependence on unresolved initial data (the term η_{z_0}), history dependence through the memory kernel $\mathcal{K}$, and—if one insists on restoring a genuine initial-value problem—the additional state information E_{dyn} needed to close it. Eliminating a field sector instead of a state sector produces the same structure in field-theoretic guise: an effective operator $A - BC^{-1}B^*$, whose correction may be nonlocal when the composed operator $BC^{-1}B^*$ retains the nonlocal Green-operator directions, and a covariant effective stress-energy T^{eff} defined by variation.

None of this is evidence that our universe has hidden source variables, memory, or any particular effective correction. The theorems are conditional: *if* a model declares a source/readout architecture, *then* these are the necessary and sufficient conditions for descent, and these are the exact reduction identities when descent fails. A memory kernel is not a physical discovery; it is a formula. It becomes physics only when the model that produced L and P is itself physics.

12. CONSEQUENCES FOR SHADOW THEORY AND CONCLUSION

The observable architecture is now precise:

$$\mathcal{H}_{\text{src}} \xrightarrow{K} \text{Ran } K \cong \mathcal{H}_{\text{src}} / \ker K, \quad (12.1)$$

the isomorphism being that of Theorem 2.3, and topologically faithful in the inherited norm exactly when $\text{Ran } K$ is closed (Remark 2.4). Static source quantities descend exactly when they are constant on cosets of $\ker K$ (Proposition 3.2). Source dynamics descend exactly when the source semigroup preserves $\ker K$ (Theorem 5.2). When it does not, exact projection produces dependence on unresolved initial data and a memory kernel (Theorem 6.2), and autonomous closure for all source states is precisely the invariance condition (Theorem 6.4). The minimal dynamical completion retains exactly those unresolved distinctions that generate different observable response functions (Theorem 7.4), and no more (Remark 7.5); in finite dimensions it is computed by a rank condition (Proposition 7.7). The nonlinear counterpart is fiberwise constancy of the projected velocity (Theorem 8.2). Eliminating a source field sector produces Schur-complement effective operators (Theorem 9.2) and, in the presence of gravity, a covariant effective stress-energy that is a genuine variational response with total covariant conservation (Definition 10.3, Theorem 10.5).

This establishes how a coherent observable field theory can be an exact quotient theory without being an autonomous presentation of every source degree of freedom—and it supplies the exact price of that fact, in three separable currencies: initial-state dependence, memory, and additional retained state.

Observable reduction therefore has three mathematically distinct static forms: exact identification by a kernel (D1), unstable inversion along small singular directions or a nonclosed range (D2), and controlled approximation by singular-value truncation (D3). Projected dynamics adds a fourth, which is independent of (D2) and (D3) and, though implied by neither, strictly *stronger* than (D1) rather than independent of it: a static quotient with a nontrivial kernel may fail to support autonomous evolution because its equivalence classes are not preserved by the source dynamics (D4). A nontrivial kernel is necessary for this failure, by (1.1), and never sufficient—whether the kernel is invariant is a separate, dynamical fact, decided by Theorem 5.2 and not by any property of K alone. The exact projection identity, the minimal response completion, the Schur-complement reduction, and the variational effective stress-energy are the field-theoretic machinery connecting source structure to observable laws. They are supplied here as mathematics, with their hypotheses stated and their limits marked, as the dynamical layer of the sequence [22, 23, 24, 25].

REFERENCES

- [1] H. Brezis, *Functional Analysis, Sobolev Spaces and Partial Differential Equations*, Universitext, Springer, New York, 2011.
- [2] C. Eckart and G. Young, *The approximation of one matrix by another of lower rank*, Psychometrika **1** (1936), no. 3, 211–218.
- [3] K.-J. Engel and R. Nagel, *One-Parameter Semigroups for Linear Evolution Equations*, Graduate Texts in Mathematics, vol. 194, Springer, New York, 2000.
- [4] H. W. Engl, M. Hanke, and A. Neubauer, *Regularization of Inverse Problems*, Mathematics and its Applications, vol. 375, Kluwer, Dordrecht, 1996.
- [5] G. Gripenberg, S.-O. Londen, and O. Staffans, *Volterra Integral and Functional Equations*, Encyclopedia of Mathematics and its Applications, vol. 34, Cambridge University Press, Cambridge, 1990.
- [6] R. E. Kalman, *On the general theory of control systems*, in: Proceedings of the First International Congress of Automatic Control (Moscow, 1960), Butterworths, London, 1961, pp. 481–492.
- [7] T. Kato, *Perturbation Theory for Linear Operators*, 2nd ed., Grundlehren der mathematischen Wissenschaften, vol. 132, Springer, Berlin, 1980.
- [8] J. M. Lee, *Introduction to Smooth Manifolds*, 2nd ed., Graduate Texts in Mathematics, vol. 218, Springer, New York, 2013.
- [9] L. Mirsky, *Symmetric gauge functions and unitarily invariant norms*, Quarterly Journal of Mathematics (Oxford), Ser. 2, **11** (1960), 50–59.

- [10] H. Mori, *Transport, collective motion, and Brownian motion*, Progress of Theoretical Physics **33** (1965), no. 3, 423–455.
- [11] A. Pazy, *Semigroups of Linear Operators and Applications to Partial Differential Equations*, Applied Mathematical Sciences, vol. 44, Springer, New York, 1983.
- [12] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*, Addison-Wesley, Reading, MA, 1995.
- [13] M. Reed and B. Simon, *Methods of Modern Mathematical Physics I: Functional Analysis*, rev. and enl. ed., Academic Press, New York, 1980.
- [14] E. Schmidt, *Zur Theorie der linearen und nichtlinearen Integralgleichungen. I. Teil: Entwicklung willkürlicher Funktionen nach Systemen vorgeschriebener*, Mathematische Annalen **63** (1907), no. 4, 433–476.
- [15] E. D. Sontag, *Mathematical Control Theory: Deterministic Finite-Dimensional Systems*, 2nd ed., Texts in Applied Mathematics, vol. 6, Springer, New York, 1998.
- [16] R. M. Wald, *General Relativity*, University of Chicago Press, Chicago, 1984.
- [17] S. Weinberg, *The Quantum Theory of Fields, Volume II: Modern Applications*, Cambridge University Press, Cambridge, 1996.
- [18] E. Zeidler, *Applied Functional Analysis: Main Principles and Their Applications*, Applied Mathematical Sciences, vol. 109, Springer, New York, 1995.
- [19] F. Zhang (ed.), *The Schur Complement and Its Applications*, Numerical Methods and Algorithms, vol. 4, Springer, New York, 2005.
- [20] R. Zwanzig, *Ensemble method in the theory of irreversibility*, Journal of Chemical Physics **33** (1960), no. 5, 1338–1341.
- [21] R. Zwanzig, *Nonequilibrium Statistical Mechanics*, Oxford University Press, Oxford, 2001.
- [22] J. Rodgers, *Source–Readout Non-Equivalence: Descent and Equivariant Reconstruction Obstructions*, Zenodo, 2026. [doi:10.5281/zenodo.21369967](https://doi.org/10.5281/zenodo.21369967).
- [23] J. Rodgers, *Target-Relative Necessity of Completion: When Readout Loss Obstructs, and What a Sufficient Extension Must Retain*, Zenodo, 2026. [doi:10.5281/zenodo.21370501](https://doi.org/10.5281/zenodo.21370501).
- [24] J. Rodgers, *Canonical Minimal Source Completion: The Coarsest Readout Extension on which a Nominated Family of Source Relations Becomes Well Defined*, Zenodo, 2026. [doi:10.5281/zenodo.21370763](https://doi.org/10.5281/zenodo.21370763).
- [25] 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](https://doi.org/10.5281/zenodo.21370985).

INDEPENDENT RESEARCHER