

Observation-Determined Normal Forms: Stability, Obstructions, and Refinement in Constraint and Rewrite Systems

Bernhard Mueller^{1*}

David Matscheko

Jonathan Hill

July 12, 2026

Paper release: r1533 **Released:** July 12, 2026

Author affiliations: ¹Pragma Research Inc.

Abstract

Confluence compares reductions from one initial state. Many distributed and reconstruction problems instead compare states or executions that expose the same observable boundary. We study this distinction first for finite rewrite and constraint systems after quotienting representation changes that are declared unobservable, and then for compact metric systems and refinement towers.

The exact layer separates unrealizable, uniquely reconstructing, and ambiguous observable fibers. It also characterizes identification modulo an arbitrary silent equivalence as endpoint agreement across all sources with the same observation; this is distinct from both same-source confluence and all-schedule termination. The associated unique-normal-form statements are classical; we record them only to identify the correct hypotheses and to distinguish failed normalization from an impossible boundary. The principal contribution is a robust, refinement-compatible calculus. A residual error-bound modulus controls distance to the consistent set, while an inverse-observation modulus controls distance between consistent states with close observations. Their composition yields a sharp two-output estimate and, conditional on settling and observation-drift receipts, bounds disagreement between schedule endpoints without further critical-pair, commutation, or contraction analysis. We prove that the inverse-observation term is exact at zero residual and that the universal coefficient of the residual term cannot be improved.

For inverse systems, one-step normalization defects telescope across arbitrary refinement depth. At one fixed level, approximate solvers admit a receipt-only comparison, with no refinement tail. Across different levels, an anchor-optimized theorem adds the amplified normalizer defects between each solver level and a common comparison level, together with protected-observation mismatch there. For a compatible admissible input, nonexpansive restrictions, summable pathwise defects, and vanishing receipt errors make the solver families induce the same coordinatewise projective limit, unchanged on cofinal subsequences; zero defects recover the exact inverse-limit normalizer. Stagewise uniqueness alone is insufficient: uniform inverse conditioning is necessary for uniform zero-residual stability, while a uniform residual error bound is separately required away from the consistent set. We also give the finite conditional-expectation projector associated with observation fibers and a directly checkable matrix receipt, isolate supporting repair obstructions, prove finite update bounds for ranked functional systems, derive singular-value certificates for linear quotient systems, and record coNP- and Π_2^P -completeness

*Corresponding author: Bernhard Mueller, bernhard@floatingpragma.ai.

classifications for succinct boundary determination and strong local repair. A small Rule-90 carrier is retained only as a formal regression test.

Keywords: abstract rewriting, observation-determined normal forms, observational confluence, error bounds, metric subregularity, boundary reconstruction, refinement, local repair, Lean 4.

Mathematics Subject Classification (2020): Primary 68Q42; Secondary 49J53, 68Q25, 68Q85, 94B05.

1 Introduction

A confluent rewrite system has a familiar geometry: paths may split, but paths from the same source eventually rejoin. In many applications the source itself is hidden. What is fixed is an interface, input, conserved sector, sensor record, or boundary condition. The relevant question is then not

$$x \rightarrow^* y, \quad x \rightarrow^* z \quad \Longrightarrow \quad y \downarrow z,$$

but rather

$$B(x) = B(x') \quad \Longrightarrow \quad \text{the settled observable outputs of } x \text{ and } x' \text{ agree.}$$

The two statements have different proof obligations. Confluence is a property of paths. Observable reconstruction is a property of the consistent set and of the observation map on that set.

We use *observation-determined normal form* for a settled state determined through an observation map on the consistent set. This is distinct from the control-theory use of “observable normal form” for a state-space form obtained by coordinate transformation, including in Boolean control networks [12].

Let Q be a state space after quotienting representation changes that are regarded as silent, let $C \subseteq Q$ be the consistent states, and let

$$B : Q \longrightarrow \mathcal{B}$$

be protected observable data. For $b \in \mathcal{B}$, set

$$C_b := C \cap B^{-1}(b).$$

The three possibilities

$$|C_b| = 0, \quad |C_b| = 1, \quad |C_b| > 1$$

mean, respectively, that the observation is unrealizable, reconstructing, or ambiguous. Termination and confluence do not alter this trichotomy. They control how a computation reaches C ; they do not determine how many points of C share one observation.

The distinction becomes more consequential under perturbation. Exact injectivity of $B|_C$ says nothing about conditioning. A sequence of finite systems may have injective observation maps at every stage and nevertheless become arbitrarily unstable under refinement. To capture the missing information, we use two moduli:

- (i) an error-bound modulus carrying a constraint residual to distance from C ;
- (ii) an inverse-observation modulus carrying observation discrepancy within C to state discrepancy.

The first is an error-bound or metric-subregularity object in the sense of variational analysis [14, 16, 20, 26]. The second is the modulus of continuity of the inverse of $B|_C$. We do not claim either notion as new. The object studied here is their use as an *observation-determined normal-form certificate* for constraint and rewrite systems and, especially, the uniform and multi-stage refinement calculus that results.

1.1 Main contributions

The manuscript has six technical contributions.

C1. Correct exact interface. We state a canonical partial normalizer and a boundary-indexed unique-normal-form theorem with the empty-fiber case separated from weak normalization. An audited-terminal theorem further separates computational quiescence from semantic consistency. Identification modulo a silent equivalence is characterized exactly by cross-source endpoint agreement. The exact confluence implications are explicitly treated as classical.

C2. Sharp observational stability. For a residual Φ and an observation map B , we define intrinsic moduli η_Φ and ω_B . For arbitrary residual levels δ_x, δ_y ,

$$\begin{aligned} d_Q(x, y) \leq & \eta_\Phi(\delta_x) + \eta_\Phi(\delta_y) \\ & + \omega_B(d_B(Bx, By) + L_B[\eta_\Phi(\delta_x) + \eta_\Phi(\delta_y)]). \end{aligned} \tag{1.1}$$

At zero residual the bound is exact by definition. Across all finite metric systems, the coefficient 2 in the equal-residual specialization cannot be lowered.

C3. Uniform refinement towers. One-step approximate naturality follows from consistency and observation defects. We distinguish same-level receipt-conditioned agreement from genuinely cross-level comparison. An anchor-optimized theorem compares independently computed outputs at arbitrary refinement levels after restriction to a common level. Compatible nested inputs leave only the intervening defect segment, and summable pathwise defects yield an implementation-independent projective limit. A counterexample shows that stagewise injectivity is not enough without a uniform inverse modulus.

C4. Repair and selection obstructions. A total exact local repair preserving an exterior collar exists if and only if the local consistency relation projects surjectively onto that collar. On a finite full-support probability space, observation-fiber resampling is the weighted orthogonal conditional-expectation projector and has a finite matrix recognition receipt. A stabilizer obstruction prevents an equivariant selector on ambiguous fibers. These results are elementary, but their integration prevents a normalizer from silently becoming a sector-selection axiom.

C5. Constructive and linear classes. Ranked functional-audit systems have singleton-or-empty fibers, terminate under every effective asynchronous schedule with an explicit update bound, and settle synchronously in dependency depth; acyclic Boolean circuits compile directly into this class. Linear quotient systems admit rank tests and singular-value bounds that instantiate (1.1); finite linear dynamics yield the usual observability matrix as a special case.

C6. Receipts, complexity, and formal audit. A standard finite Markov drift lemma turns a row-wise kernel inequality into endpoint, stationary, and expected-occupation receipts, without asserting infinite-horizon confinement under persistent noise. For succinct Boolean systems, we record tight coNP- and Π_2^P -completeness classifications in this formulation; these supporting results are not part of the principal novelty claim. We audit the accompanying Lean sources theorem by theorem and state exactly which results remain to be formalized.

1.2 Novelty boundary and intended venue

The exact theory overlaps several established areas. Abstract rewriting supplies Newman’s lemma, Church–Rosser properties, unique normal forms, and normalization theory [3, 8, 21, 32, 36]. Rewriting modulo equivalence and confluence under invariants already address quotient and relevant-state restrictions [11, 13, 18, 27]. Quantitative rewriting enriches rewrite relations themselves and proves metric analogues of classical confluence theorems [19]. Error bounds and inverse stability are mature subjects in variational analysis [14, 26]. Functoriality of rewriting is also established in categorical graph transformation [17].

Accordingly, the paper makes no claim that a triangle inequality, an error bound, inverse continuity, a section criterion, or boundary-indexed unique normal forms are individually new. The contribution advanced for review is the integrated certificate architecture, particularly the explicit anchor-optimized cross-level receipt comparison and its cofinal-limit consequence, rather than the triangle inequality or Lipschitz telescope in isolation. Approximate inverse-system theory is important neighboring prior art, and the displayed estimate itself is elementary; we therefore claim an organizing synthesis, not priority for approximate commutativity or inverse limits. External specialist review remains necessary. The intended contribution is in mathematical theoretical computer science, with rewriting and semantics as the primary audience.

2 Relation to existing theories

2.1 Rewriting, invariants, and quotients

The classical properties of an abstract reduction system are standard [8, 21, 36]. Confluence restricted to an invariant or taken modulo an equivalence has been developed in Constraint Handling Rules and graph transformation [11, 13, 18]. Those frameworks are close to the first arrow in

$$\text{presentations} \longrightarrow Q \longrightarrow \mathcal{B} :$$

they identify states that should count as the same or restrict attention to relevant states. Our observation map is the second arrow. It may be strictly coarser than equality in Q , and it is required to be injective only on the consistent subset when exact reconstruction is desired.

Unique-normal-form properties weaker than confluence are well known and have been automated for term rewriting [3]. The fiber trichotomy and partial-normalizer universal property here are elementary set theory; only the reachable-normal-form and fiberwise-confluence corollaries specialize classical rewriting theory.

2.2 Quantitative rewriting versus output stability

Gavazzo and Di Florio enrich rewriting relations over quantales and derive quantitative Newman, Church–Rosser, and critical-pair theorems [19]. Here the rewrite relation remains ordinary. Distances are attached to two semantic interfaces: residual-to-solution error and inversion of observable

data on the solution set. Thus the theory does not estimate a quantitative rewrite path; it estimates disagreement between settled outputs of possibly unrelated schedules or initial interiors.

2.3 Error bounds and inverse mappings

The inequality

$$\text{dist}(x, C) \leq \eta(\Phi(x))$$

is an error-bound property. In variational analysis, local linear versions are equivalent to metric subregularity or calmness of suitable set-valued mappings [14, 16, 26]. The inverse-observation modulus is a global finite or compact analogue of inverse Lipschitz/calmness information. Our use of these notions is intentionally transparent: the novelty claim concerns their operational consequences for schedules and refinement towers, not their analytic origin.

2.4 Approximate inverse systems

Approximate inverse systems and almost-commuting diagrams have a substantial topological literature, including stability and resolution theorems for compacta [28–30]. Those works perturb or relax compatibility of bonding data in topological or cover terms. Our setting keeps the ambient inverse system exact and instead gives pointwise metric bounds for normalizer defects, solver receipts, and protected-observation mismatch. The connection limits the novelty claim: the contribution is an explicit certificate adapted to observation-determined normalizers, not the invention of approximate inverse systems.

2.5 Static reconstruction and functional observability

Control-theoretic observability reconstructs a state, or a prescribed function of state, from an output history. Linear observability and functional observability supply rank and conditioning tests; recent nonlinear work characterizes when a nonlinear functional can be reconstructed from measurements and studies the quality of the inverse map [24, 25, 31]. This is the closest neighboring language for the map $B|_C$. Our setting is nevertheless static at the theorem interface: C is the already-settled consistent set, B may be any protected quotient observation, and the new questions are how residual error and inverse conditioning control schedule disagreement and how these certificates propagate through refinement towers.

2.6 Constraints, databases, and global sections

Acyclic database schemes and positive-semiring relations provide local-to-global consistency results [7, 9]. Constraint satisfaction studies local consistency as an algorithmic reduction [15]. Database repair semantics distinguishes the existence of many consistent repairs from the answers invariant across them [4]. Sheaf-theoretic contextuality similarly turns failure of a global section into an obstruction [2]. Our empty-fiber certificate belongs to this broad lineage. The additional layer is the interaction with normal-form stability and refinement.

2.7 Asynchronous functional networks

Kahn process networks provide timing-independent determinacy for networks of functional components, while generalized Kahn principles treat broader abstract asynchronous, including nondeterministic, models under their stated hypotheses [1, 23]. Ranked functional systems below are a finite terminating special case with residual audits. Their purpose is not to reprove those principles, but

to give a class in which the exact boundary hypotheses, update bound, and obstruction output are explicit.

2.8 Formalized rewriting

Formalized rewriting has established libraries such as CoLoR in Coq [10]. Recent developments include a constructive Agda library for abstract rewriting [5], a Lean 4 framework for confluence and normalization of lambda calculi [34], and a Rocq library for adhesive-category graph rewriting [6]. The accompanying standalone Lean artifact formalizes selected exact, stability, refinement, repair, ranked, and stochastic results from this paper; its coverage boundary is stated in [Section 11](#).

3 Exact observable fibers

Definition 3.1 (Observable quotient system). An *observable quotient system* is a tuple

$$\mathfrak{S} = (Q, C, \mathcal{B}, B),$$

where Q is a set, $C \subseteq Q$ is the consistent subset, $\mathcal{B}$ is an observation set, and $B : Q \rightarrow \mathcal{B}$ is the protected observation map. In applications, Q may itself be a quotient Σ/Γ of presentations by gauge, implementation, or representation changes.

Definition 3.2 (Fiber trichotomy). For $b \in \mathcal{B}$, define $C_b = C \cap B^{-1}(b)$. The observation b is *unrealizable*, *reconstructing*, or *ambiguous* according as $|C_b|$ is 0, 1, or greater than 1. The system is *observable-determined* if $B|_C$ is injective.

Definition 3.3 (Canonical partial normalizer). If $B|_C$ is injective, define

$$N_B : Q \longrightarrow C \sqcup \{\perp\}$$

by returning the unique member of $C_{B(x)}$ when the fiber is nonempty and $\perp$ otherwise. The symbol $\perp$ denotes a mathematically proved no-extension outcome, not a failure of search.

Remark 3.4 (Proof-carrying interpretation). For finite data with decidable membership and equality, exhaustive search returns the dependent alternative

$$\left(\sum_{c \in C} [B(c) = B(x)] \right) \sqcup [C_{B(x)} = \emptyset].$$

Thus “obstructed” carries a proof of empty fiber. In the ordinary classical setting the same alternative is defined by excluded middle and choice, but injectivity alone is not an algorithm for deciding which side holds. Even for finite succinct inputs, such a proof need not be computationally cheap; [Section 9](#) shows that realizability is NP-complete.

Theorem 3.5 (Universal property). *The following are equivalent.*

- (a) $B|_C$ is injective.
- (b) There is a unique map $N : Q \rightarrow C \sqcup \{\perp\}$ such that:
 - (i) $N(x) = c \in C$ implies $B(c) = B(x)$;
 - (ii) $N(x) = \perp$ if and only if $C_{B(x)} = \emptyset$;

- (iii) $N(c) = c$ for all $c \in C$;
- (iv) $B(x) = B(y)$ implies $N(x) = N(y)$.

(c) $B|_C : C \rightarrow B(C)$ is a bijection.

If $B(Q) \subseteq B(C)$, then N is a total idempotent retraction $Q \rightarrow C$ preserving B .

Proof. If $B|_C$ is injective, every nonempty fiber contains exactly one point, so N_B has all listed properties and any competing map is forced to agree with it fiber by fiber. Conversely, if $c, c' \in C$ have $B(c) = B(c')$, then boundary extensionality and the fixed-point property give $c = N(c) = N(c') = c'$. The equivalence with (c) is immediate. $\square$

Let $\rightarrow$ be a relation on Q , write $\rightarrow^*$ for its reflexive transitive closure, and let $\text{NF}(\rightarrow)$ be the normal forms.

Definition 3.6. The relation $\rightarrow$ is *observation preserving* if $x \rightarrow y$ implies $B(x) = B(y)$, and *complete for C* if $\text{NF}(\rightarrow) = C$.

Proposition 3.7 (Reachable normal forms live in the observable fiber). *If $\rightarrow$ preserves B and is complete for C , then every normal form reachable from x belongs to $C_{B(x)}$.*

Proof. Completeness puts the reachable normal form in C , and observation preservation puts it in the fiber of $B(x)$. $\square$

Theorem 3.8 (Cross-source endpoint identification modulo equivalence). *Assume that $\rightarrow$ preserves B and that $\text{NF}(\rightarrow) = C$. Let $\approx$ be an equivalence relation on C . The following are equivalent:*

- (i) for all $c, d \in C$, $B(c) = B(d)$ implies $c \approx d$;
- (ii) whenever $B(x) = B(y)$, $x \rightarrow^* n_x$, $y \rightarrow^* n_y$, and $n_x, n_y \in \text{NF}(\rightarrow)$, one has $n_x \approx n_y$.

Thus observation identifiability modulo $\approx$ is exactly agreement modulo $\approx$ of reachable endpoints from possibly different sources with the same observation.

Proof. Under (i), Proposition 3.7 puts n_x, n_y in C and observation preservation gives $B(n_x) = B(x) = B(y) = B(n_y)$, hence $n_x \approx n_y$. Conversely, take $c, d \in C$ with $B(c) = B(d)$. Since $C = \text{NF}(\rightarrow)$ and $\rightarrow^*$ is reflexive, apply (ii) to $c \rightarrow^* c$ and $d \rightarrow^* d$. $\square$

Corollary 3.9 (Silent-quotient criterion). *In addition to the hypotheses of Theorem 3.8, suppose $c \approx d$ implies $B(c) = B(d)$. Then $B|_C$ induces a surjection*

$$\overline{B} : C/\approx \longrightarrow B(C), \quad [c] \longmapsto B(c),$$

and $\overline{B}$ is injective, hence bijective, if and only if either condition of Theorem 3.8 holds. This includes quotients by silent representation or gauge changes; equality recovers ordinary endpoint uniqueness.

Proof. Silence makes $\overline{B}$ well defined and it is surjective by its codomain. Its injectivity says exactly that $B(c) = B(d)$ implies $[c] = [d]$, equivalently $c \approx d$. $\square$

Remark 3.10 (Three separate obligations). Same-source confluence compares arms of a peak from one state; Theorem 3.8 compares endpoints from all sources in one observation fiber. Weak normalization is a separate existence hypothesis ensuring that every source has at least one endpoint. Strong normalization, fairness, or another liveness argument is still needed to conclude that every allowed schedule settles. None of these three obligations may be substituted for either of the others.

Corollary 3.11 (Correct empty/singleton alternative). *Under the hypotheses of Proposition 3.7:*

- (i) *if $C_b = \emptyset$, no state of $B^{-1}(b)$ reaches a normal form;*
- (ii) *if $C_b = \{c_b\}$, every reachable normal form from $B^{-1}(b)$ equals c_b ;*
- (iii) *if, in addition, every state of $B^{-1}(b)$ is weakly normalizing, then every such state reduces to c_b , and $\rightarrow$ is confluent on that fiber.*

Proof. The first two claims follow directly from Proposition 3.7. For (iii), normalize both arms of any peak. Their normal forms are both c_b , providing a join. $\square$

The completeness hypothesis $\text{NF}(\rightarrow) = C$ is sometimes too strong: a computation may halt uniquely and then fail a separate semantic audit. The next theorem separates those two layers.

Theorem 3.12 (Audited terminal alternative). *Let $T = \text{NF}(\rightarrow)$ be the computational terminal set and let $C \subseteq T$ be the semantically consistent terminals. Assume that $\rightarrow$ preserves B , every state is weakly normalizing, and $B|_T$ is injective. Then every $x \in Q$ reaches a unique terminal $t_x \in T$, and*

$$B(t_x) = B(x), \quad t_x \in C \iff C_{B(x)} \neq \emptyset.$$

Thus a realizable observation has a unique reachable terminal at its consistent extension, whereas an unrealizable observation can still have a unique reachable terminal, but only outside C . The theorem does not assert strong normalization of every execution.

Proof. Weak normalization supplies a reachable $t_x \in T$, and observation preservation gives $B(t_x) = B(x)$. Any other reachable terminal has the same observation, so injectivity of $B|_T$ makes it equal to t_x . If $t_x \in C$, it witnesses a nonempty consistent fiber. Conversely, if $c \in C_{B(x)}$, then $c \in T$ and $B(c) = B(t_x)$; injectivity on T gives $c = t_x$, hence $t_x \in C$. $\square$

When $T = C$, Theorem 3.12 reduces to the realizable part of Corollary 3.11. When $T \supsetneq C$, it prevents an audit failure from being misreported as nontermination or as a consistent normal form.

Remark 3.13 (Classical status). Part (iii) is the invariant-fiber form of the classical implication from weak normalization and uniqueness of normal forms with respect to reduction to confluence. It is included to identify the exact hypotheses needed for observable reconstruction; it is not part of the novelty claim.

4 Observational stability moduli

Let (Q, d_Q) be a finite metric space, let $(\mathcal{B}, d_{\mathcal{B}})$ be any metric space, let $\emptyset \neq C \subseteq Q$, and let $\Phi : Q \rightarrow \mathbb{R}_{\geq 0}$ satisfy

$$C = \Phi^{-1}(0).$$

All maxima below therefore exist. The same formulas extend to compact spaces by replacing maxima with suprema and imposing the continuity hypotheses stated in Proposition 4.13; the displayed Lipschitz estimates continue to assume a finite Lipschitz constant for B .

Definition 4.1 (Residual error-bound modulus). Define

$$\eta_{\Phi}(t) := \max\{\text{dist}_Q(x, C) : \Phi(x) \leq t\}.$$

An *error-bound modulus* is any nondecreasing η satisfying

$$\text{dist}_Q(x, C) \leq \eta(\Phi(x)).$$

Definition 4.2 (Inverse-observation modulus). Define

$$\omega_B(r) := \max\{d_Q(c, c') : c, c' \in C, d_B(Bc, Bc') \leq r\}.$$

The system is *stably observable-determined* if $\omega_B(r) \rightarrow 0$ as $r \downarrow 0$.

Proposition 4.3 (Minimality of the intrinsic moduli). *The functions η_Φ and ω_B are nondecreasing. Moreover:*

- (i) $\eta_\Phi(0) = 0$;
- (ii) η_Φ is the least nondecreasing error-bound modulus on the attained residual values;
- (iii) $\omega_B(0) = 0$ if and only if $B|_C$ is injective;
- (iv) ω_B is the least nondecreasing modulus satisfying

$$d_Q(c, c') \leq \omega(d_B(Bc, Bc')) \quad (c, c' \in C).$$

For finite C , injectivity of $B|_C$ implies that there is an $r_0 > 0$ such that $\omega_B(r) = 0$ for $0 \leq r < r_0$. When $|C| > 1$, one may take r_0 to be the minimum distance between distinct points of $B(C)$; when $|C| = 1$, any positive r_0 works.

Proof. Monotonicity follows from inclusion of the defining sets. Since $C = \Phi^{-1}(0)$, residual zero has distance zero from C . If η is any nondecreasing error-bound modulus and $\Phi(x) \leq t$, then $\text{dist}(x, C) \leq \eta(\Phi(x)) \leq \eta(t)$; taking the maximum gives $\eta_\Phi(t) \leq \eta(t)$. The proof for ω_B is identical. The statement at zero is exactly injectivity, and finiteness gives a positive minimum separation among distinct boundary images. $\square$

Let L_B be any Lipschitz constant for B on Q .

Theorem 4.4 (Heterogeneous two-output estimate). *If $x, y \in Q$ satisfy*

$$\Phi(x) \leq \delta_x, \quad \Phi(y) \leq \delta_y, \quad d_B(Bx, By) \leq \varepsilon,$$

then

$$\begin{aligned} d_Q(x, y) &\leq \eta_\Phi(\delta_x) + \eta_\Phi(\delta_y) \\ &\quad + \omega_B(\varepsilon + L_B[\eta_\Phi(\delta_x) + \eta_\Phi(\delta_y)]). \end{aligned} \tag{4.1}$$

Proof. Choose nearest points $c_x, c_y \in C$. Then

$$d_Q(x, c_x) \leq \eta_\Phi(\delta_x), \quad d_Q(y, c_y) \leq \eta_\Phi(\delta_y).$$

Lipschitz continuity and the triangle inequality give

$$\begin{aligned} d_B(Bc_x, Bc_y) &\leq d_B(Bc_x, Bx) + d_B(Bx, By) + d_B(By, Bc_y) \\ &\leq \varepsilon + L_B[\eta_\Phi(\delta_x) + \eta_\Phi(\delta_y)]. \end{aligned}$$

The definition of ω_B bounds $d_Q(c_x, c_y)$. A final triangle inequality yields (4.1). $\square$

The symmetric form used most often is

$$d_Q(x, y) \leq 2\eta_\Phi(\delta) + \omega_B(\varepsilon + 2L_B\eta_\Phi(\delta)). \tag{4.2}$$

Corollary 4.5 (Rate transfer). *Suppose, on the residual and observation ranges under consideration,*

$$\eta_\Phi(t) \leq at^p, \quad \omega_B(r) \leq br^q$$

for constants $a, b \geq 0$ and exponents $p, q > 0$. Then

$$d_Q(x, y) \leq a(\delta_x^p + \delta_y^p) + b\left(\varepsilon + L_B a(\delta_x^p + \delta_y^p)\right)^q$$

under the hypotheses of Theorem 4.4. In particular, linear error bounds and a Lipschitz inverse give a linear output certificate.

Proof. Substitute the two rate bounds into (4.1). □

4.1 Optimality and the exact observational modulus

Definition 4.6 (Settled-output modulus). For $\delta, r \geq 0$, define

$$\Theta_{\Phi, B}(\delta, r) := \max \{d_Q(x, y) : \Phi(x), \Phi(y) \leq \delta, d_B(Bx, By) \leq r\}.$$

Theorem 4.7 (Sharpness). *The following statements hold.*

(i) *At zero residual,*

$$\Theta_{\Phi, B}(0, r) = \omega_B(r).$$

Thus the inverse-observation term in (4.2) is exact.

(ii) *No universal estimate of the form*

$$\Theta_{\Phi, B}(\delta, 0) \leq a \eta_\Phi(\delta) + \omega_B(0)$$

can hold for all finite metric systems with a constant $a < 2$.

Proof. When $\delta = 0$, the admissible points are precisely C , so the two defining maxima coincide.

For (ii), fix $s > 0$ and let $Q = \{c, x, y\}$ with distances

$$d(c, x) = d(c, y) = s, \quad d(x, y) = 2s.$$

Let $C = \{c\}$, let B be constant, and set $\Phi(c) = 0$, $\Phi(x) = \Phi(y) = 1$. Then $\eta_\Phi(1) = s$, $\omega_B(0) = 0$, and $\Theta_{\Phi, B}(1, 0) = 2s$. □

Remark 4.8. The theorem does not assert that every term in (4.1) can be attained simultaneously in every structured class. It states the correct universal boundary: the observation term is forced already on C , and the coefficient 2 on two independent residual errors is globally sharp.

4.2 Approximate schedule independence

A finite execution is a sequence $x = x_0 \rightarrow x_1 \rightarrow \dots \rightarrow x_n$. It is (δ, β) -settled if its endpoint z satisfies

$$\Phi(z) \leq \delta, \quad d_B(Bz, Bx) \leq \beta.$$

No exact observation-preservation assumption is made.

Corollary 4.9 (Approximate schedule independence). *Suppose two executions start at the same state x and end at z_1, z_2 , where the i th execution is (δ_i, β_i) -settled. Then*

$$d_Q(z_1, z_2) \leq \eta_\Phi(\delta_1) + \eta_\Phi(\delta_2) + \omega_B(\beta_1 + \beta_2 + L_B[\eta_\Phi(\delta_1) + \eta_\Phi(\delta_2)]). \quad (4.3)$$

The same bound holds for executions starting from two different states with equal initial observations.

Proof. The endpoint observation discrepancy is at most $\beta_1 + \beta_2$. Apply [Theorem 4.4](#). $\square$

Corollary 4.10 (High-probability schedules). *Let Z_1, Z_2 be random endpoints of two schedules from a common initial state in finite Q . Suppose schedule i is (δ_i, β_i) -settled with probability at least $1 - p_i$. Let D denote the right-hand side of (4.3). Then*

$$\mathbb{P}[d_Q(Z_1, Z_2) \leq D] \geq 1 - p_1 - p_2$$

and

$$\mathbb{E}[d_Q(Z_1, Z_2)] \leq D + (p_1 + p_2) \text{diam}(Q).$$

Proof. Use the union bound for the two settling events. On their intersection apply [Corollary 4.9](#); on its complement use the diameter bound. $\square$

The preceding result consumes a probabilistic settling receipt. The following standard drift calculation is one finite mechanism for producing such a receipt.

Proposition 4.11 (Finite Markov drift receipt). *Let K be a Markov kernel on finite Q , write*

$$(Kf)(x) = \sum_{y \in Q} K(x, y)f(y),$$

and let $V : Q \rightarrow \mathbb{R}_{\geq 0}$. Suppose $0 \leq \kappa < 1$, $\xi \geq 0$, and

$$(KV)(x) \leq \kappa V(x) + \xi \quad (x \in Q).$$

For the chain (X_n) started at x ,

$$\mathbb{E}_x[V(X_n)] \leq \kappa^n V(x) + \xi \frac{1 - \kappa^n}{1 - \kappa}.$$

Consequently, for every $\delta > 0$,

$$\mathbb{P}_x[V(X_n) > \delta] \leq \frac{\kappa^n V(x) + \xi(1 - \kappa^n)/(1 - \kappa)}{\delta}.$$

If $V = \Phi$ and every transition with positive probability preserves B , then the complementary event supplies a $(\delta, 0)$ -settling receipt for [Corollary 4.10](#). When the kernel is meant to sample rewrite schedules, its support must additionally consist of permitted rewrite steps or stuttering moves. This is a one-time endpoint bound; pathwise confinement over a time interval requires an additional excursion or concentration hypothesis.

Proof. The tower property and the drift inequality give $a_{n+1} \leq \kappa a_n + \xi$ for $a_n = \mathbb{E}_x[V(X_n)]$. Induction solves this scalar recurrence and gives the displayed geometric sum. The probability bound is Markov's inequality. Positive-probability transitions preserve the initial observation along every sampled path. $\square$

Corollary 4.12 (Stationary and occupation receipts). *Under the hypotheses of Proposition 4.11, every stationary distribution π satisfies, for $\delta > 0$,*

$$\mathbb{E}_\pi V \leq \frac{\xi}{1-\kappa}, \quad \pi\{V > \delta\} \leq \frac{\xi}{(1-\kappa)\delta}.$$

For a chain started at x and $N \geq 1$,

$$\mathbb{E}_x \left[\frac{1}{N} \sum_{n=0}^{N-1} \mathbf{1}_{\{V(X_n) > \delta\}} \right] \leq \frac{V(x)}{N\delta(1-\kappa)} + \frac{\xi}{\delta(1-\kappa)}.$$

The same right-hand side bounds the probability of being outside the δ -tube at an independent uniformly sampled time in $\{0, \dots, N-1\}$. When $\xi > 0$, these expectation and occupation bounds do not imply infinite-horizon pathwise confinement; persistent noise can permit repeated excursions.

Proof. Integrate the drift inequality against π and use stationarity. For the occupation bound, use $\mathbf{1}_{\{V > \delta\}} \leq V/\delta$, sum the estimate in Proposition 4.11, and bound $\sum_{n=0}^{N-1} \kappa^n \leq (1-\kappa)^{-1}$. $\square$

4.3 Compact and uniform families

Proposition 4.13 (Compact-space extension). *Let Q be compact metric, let $C \subseteq Q$ be nonempty and closed, let $\Phi : Q \rightarrow \mathbb{R}_{\geq 0}$ be continuous with $C = \Phi^{-1}(0)$, and let $B : Q \rightarrow \mathcal{B}$ be continuous. Then $\eta_\Phi(t) \rightarrow 0$ as $t \downarrow 0$. If $B|_C$ is injective, then its inverse on $B(C)$ is uniformly continuous and $\omega_B(r) \rightarrow 0$ as $r \downarrow 0$.*

Proof. If $\eta_\Phi(t)$ did not tend to zero, there would be $t_n \downarrow 0$ and x_n with $\Phi(x_n) \leq t_n$ but $\text{dist}(x_n, C) \geq a > 0$. A convergent subsequence has limit x with $\Phi(x) = 0$, hence $x \in C$, contradicting the distance bound. The second statement follows because a continuous bijection from compact C to the Hausdorff space $B(C)$ is a homeomorphism, whose inverse is uniformly continuous. $\square$

Now let $\{\mathfrak{S}_i\}_{i \in I}$ be a nonempty family of compact or finite metric systems with a common Lipschitz bound L for the maps B_i . Define

$$H(t) := \sup_i \eta_i(t), \quad \Omega(r) := \sup_i \omega_i(r).$$

We assume these suprema are finite on the ranges under consideration; this uniform boundedness is not automatic when the diameters of the systems are unbounded.

A *uniform settled-output modulus* is a function $\Psi(t, r)$ such that $\Theta_{\Phi_i, B_i}(t, r) \leq \Psi(t, r)$ for every $i \in I$.

Theorem 4.14 (Uniform observational stability). *If $H(t) \rightarrow 0$ and $\Omega(r) \rightarrow 0$ at zero, then every member of the family satisfies the uniform estimate*

$$d_i(x, y) \leq 2H(\delta) + \Omega(r + 2LH(\delta))$$

whenever $\Phi_i(x), \Phi_i(y) \leq \delta$ and $d_{\mathcal{B}_i}(B_i x, B_i y) \leq r$. Conversely, any uniform settled-output modulus restricts at zero residual to a modulus dominating Ω .

Proof. Apply Theorem 4.4 stagewise and take suprema. At zero residual, Theorem 4.7(i) gives the converse statement. $\square$

Example 4.15 (Stagewise uniqueness without uniform stability). For $n \geq 1$, let $C_n = Q_n = \{0, 1\}$ with discrete state distance, and let $B_n(0) = 0$, $B_n(1) = 1/n$ in $\mathbb{R}$. Every B_n is injective, yet for every $r > 0$ one has $\sup_n \omega_{B_n}(r) = 1$. Thus exact uniqueness at every finite stage does not supply a stable refinement limit.

Example 4.16 (Stagewise error bounds without a uniform error bound). For $n \geq 1$, let $Q_n = \{0, 1\}$ with the discrete metric, let $C_n = \{0\}$, let B_n be constant, and put

$$\Phi_n(0) = 0, \quad \Phi_n(1) = 1/n.$$

For every fixed n , $\eta_n(t) \rightarrow 0$ as $t \downarrow 0$, and the unique consistent fiber is perfectly conditioned. Nevertheless, for every $t > 0$ one has $\sup_n \eta_n(t) = 1$. Thus uniform residual control and uniform inverse-observation control are independent requirements.

4.4 Parallel composition and observation design

Theorem 4.17 (Exact product calculus). *Let I be nonempty and finite. For each $i \in I$, let*

$$\mathfrak{S}_i = (Q_i, C_i, \mathcal{B}_i, B_i, \Phi_i)$$

be a finite metric system with $C_i \neq \emptyset$. Form the parallel product

$$Q = \prod_{i \in I} Q_i, \quad C = \prod_{i \in I} C_i, \quad B = (B_i)_{i \in I}, \quad \Phi(x) = \max_i \Phi_i(x_i),$$

with maximum metrics on Q and $\prod_i \mathcal{B}_i$. Then

$$\eta_\Phi(t) = \max_i \eta_{\Phi_i}(t), \quad \omega_B(r) = \max_i \omega_{B_i}(r).$$

Consequently, parallel composition preserves a uniform stability certificate without a dimension-dependent loss in the maximum metric.

Proof. For the maximum product metric,

$$\text{dist}_Q(x, C) = \max_i \text{dist}_{Q_i}(x_i, C_i).$$

The constraint $\Phi(x) \leq t$ is equivalent to $\Phi_i(x_i) \leq t$ for every i . Taking maxima gives the upper bound for η_Φ ; equality is attained by choosing a maximizing component and arbitrary consistent points in all other components. The proof for ω_B is identical: the product observation discrepancy is at most r exactly when every component discrepancy is at most r . $\square$

Corollary 4.18 (Adding observations cannot worsen conditioning). *Let $D : C \rightarrow \mathcal{D}$ be an additional observation and equip $\mathcal{B} \times \mathcal{D}$ with the maximum metric. For the enriched map $B^+ = (B, D)$,*

$$\omega_{B^+}(r) \leq \omega_B(r).$$

Thus extra quotient-visible sensors can only improve the intrinsic inverse-observation certificate; they may, of course, carry their own implementation error outside C .

Proof. The pairs of consistent states satisfying $d_{\mathcal{B} \times \mathcal{D}}(B^+c, B^+c') \leq r$ form a subset of those satisfying $d_{\mathcal{B}}(Bc, Bc') \leq r$. $\square$

5 Naturality and refinement towers

For an observable system define its reconstructing, realizable domain by

$$Q^{\text{adm}} := \{q \in Q : \text{there exists a unique } c \in C \text{ with } B(c) = B(q)\}.$$

Every normalizer in this section is understood to act on a chosen metric subspace of Q^{adm} , with the consistent set restricted to the corresponding normalizer image; every coarse map is assumed to preserve the chosen admissible subspaces. To avoid carrying restriction notation through every formula, we relabel those subspaces and restricted consistent sets as Q_s, Q_r, Q_n and C_s, C_r, C_n . Thus $B_n|_{C_n}$ is injective on the declared domain. Obstructed and ambiguous fibers remain governed by [Section 3](#); they are not silently assigned a finite normalizer value in the metric estimates below.

Let

$$\mathfrak{S}_s = (Q_s, C_s, \mathcal{B}_s, B_s, \Phi_s), \quad \mathfrak{S}_r = (Q_r, C_r, \mathcal{B}_r, B_r, \Phi_r)$$

be finite or compact observable systems on the admissible domains declared above. Their normalizers N_s and N_r are therefore total on these domains. Let $f : Q_s \rightarrow Q_r$ be a coarse state map and $h : \mathcal{B}_s \rightarrow \mathcal{B}_r$ a coarse observation map.

Definition 5.1 (Morphism defects). The pair (f, h) has observation defect ε_B and consistency defect δ_C if

$$d_{\mathcal{B}_r}(B_r(fq), h(B_sq)) \leq \varepsilon_B \quad (q \in Q_s), \quad (5.1)$$

$$\Phi_r(fc) \leq \delta_C \quad (c \in C_s). \quad (5.2)$$

Theorem 5.2 (One-step approximate naturality). Assume B_r is L_r -Lipschitz. Under (5.1)–(5.2),

$$d_{Q_r}(fN_sq, N_rf q) \leq \eta_r(\delta_C) + \omega_r(2\varepsilon_B + L_r\eta_r(\delta_C)) \quad (5.3)$$

for every relevant $q \in Q_s$.

Proof. Put $y = fN_sq$ and $z = N_rf q$. By (5.2), $\Phi_r(y) \leq \delta_C$. Choose $c \in C_r$ with $d_r(y, c) \leq \eta_r(\delta_C)$. Since both normalizers preserve observations,

$$\begin{aligned} d_{\mathcal{B}_r}(B_rc, B_rz) &\leq d_{\mathcal{B}_r}(B_rc, B_ry) + d_{\mathcal{B}_r}(B_rfN_sq, hB_sN_sq) \\ &\quad + d_{\mathcal{B}_r}(hB_sq, B_rf q) \\ &\leq L_r\eta_r(\delta_C) + 2\varepsilon_B. \end{aligned}$$

Both c and z lie in C_r , so the inverse-observation modulus bounds $d_r(c, z)$. Add $d_r(y, c)$. □

Corollary 5.3 (Exact naturality from uniqueness). If $f(C_s) \subseteq C_r$ and $B_r \circ f = h \circ B_s$, then

$$f \circ N_s = N_r \circ f.$$

Proof. Both sides are consistent coarse states with the same coarse observation. Reconstructing fibers make them equal. Equivalently, set both defects to zero in [Theorem 5.2](#). □

Corollary 5.4 (Stability under perturbation of the consistency model). Let (Q, d_Q) carry two nonempty compact consistent sets C and $\tilde{C}$, two observation maps $B, \tilde{B} : Q \rightarrow \mathcal{B}$, and total maps $N, \tilde{N} : Q \rightarrow Q$ satisfying

$$N(Q) \subseteq C, \quad \tilde{N}(Q) \subseteq \tilde{C}, \quad B \circ N = B, \quad \tilde{B} \circ \tilde{N} = \tilde{B}.$$

Suppose

$$d_H(C, \tilde{C}) \leq \alpha, \quad \sup_{q \in Q} d_{\mathcal{B}}(Bq, \tilde{B}q) \leq \beta,$$

where d_H is Hausdorff distance. If B is L -Lipschitz and $B|_C$ has inverse modulus ω_B , then

$$d_Q(Nx, \tilde{N}x) \leq \alpha + \omega_B(2\beta + L\alpha) \quad (x \in Q).$$

Proof. Choose $c \in C$ with $d_Q(c, \tilde{N}x) \leq \alpha$. Since N and $\tilde{N}$ preserve their respective observations,

$$\begin{aligned} d_{\mathcal{B}}(Bc, BNx) &\leq d_{\mathcal{B}}(Bc, B\tilde{N}x) + d_{\mathcal{B}}(B\tilde{N}x, \tilde{B}\tilde{N}x) \\ &\quad + d_{\mathcal{B}}(\tilde{B}x, Bx) \\ &\leq L\alpha + 2\beta. \end{aligned}$$

Both c and Nx lie in C , so $d_Q(c, Nx) \leq \omega_B(L\alpha + 2\beta)$. Add the Hausdorff error. $\square$

Remark 5.5. Categorical functoriality of rewrite steps is an established subject [17]. [Corollary 5.3](#) is a different and simpler statement: it concerns the terminal normalizer, and uniqueness allows it to be proved without simulating each fine rewrite. The nontrivial content for applications is the defect estimate, its accumulation across a tower, and the perturbation certificate of [Corollary 5.4](#).

5.1 Accumulation through arbitrary depth

Let

$$Q_0 \xleftarrow{\rho_{1,0}} Q_1 \xleftarrow{\rho_{2,1}} Q_2 \leftarrow \dots$$

be an inverse sequence of nonempty finite or compact metric spaces with Lipschitz restriction maps, with $\rho_{m,n} = \rho_{n+1,n} \circ \dots \circ \rho_{m,m-1}$. Let $N_n : Q_n \rightarrow Q_n$ be normalizers. The boundedness of each Q_n makes the following pathwise and uniform one-step normalization defects finite:

$$\alpha_n(q) := d_n(\rho_{n+1,n}N_{n+1}q, N_n\rho_{n+1,n}q), \quad a_n := \sup_{q \in Q_{n+1}} \alpha_n(q).$$

For $j \geq n$, write

$$K_{j,n} := \text{Lip}(\rho_{j,n}), \quad K_{n,n} = 1.$$

Theorem 5.6 (Telescoping refinement bound). *For every $m > n$ and every $q \in Q_m$,*

$$d_n(\rho_{m,n}N_mq, N_n\rho_{m,n}q) \leq \sum_{j=n}^{m-1} K_{j,n}a_j. \quad (5.4)$$

Proof. For $j = n, \dots, m$, define

$$p_j := \rho_{j,n}N_j\rho_{m,j}q.$$

Then $p_m = \rho_{m,n}N_mq$ and $p_n = N_n\rho_{m,n}q$. For $j < m$, put $r = \rho_{m,j+1}q$. The difference between consecutive terms is

$$\begin{aligned} d_n(p_{j+1}, p_j) &= d_n(\rho_{j,n}\rho_{j+1,j}N_{j+1}r, \rho_{j,n}N_j\rho_{j+1,j}r) \\ &\leq K_{j,n}a_j. \end{aligned}$$

Summing the telescoping chain proves (5.4). $\square$

Corollary 5.7 (Summable-defect asymptotic naturality). *If all restriction maps are nonexpansive, then*

$$d_n(\rho_{m,n}N_m q, N_n \rho_{m,n} q) \leq \sum_{j=n}^{m-1} a_j.$$

If $\sum_j a_j < \infty$, the defect is uniformly small when the comparison level tends outward along the refinement index. More generally, the sufficient condition is

$$\sup_{m>n} \sum_{j=n}^{m-1} K_{j,n} a_j \longrightarrow 0 \quad (n \rightarrow \infty).$$

Combining [Theorem 5.2](#) with [Theorem 5.6](#) gives an explicit tower certificate in terms of residual, observation, and Lipschitz data at each stage.

Corollary 5.8 (Modulus-generated tower certificate). *Suppose the n th restriction map has consistency defect δ_n and observation defect ε_n , and let*

$$a_n^* := \eta_n(\delta_n) + \omega_n(2\varepsilon_n + L_n \eta_n(\delta_n)).$$

Then [\(5.4\)](#) holds with a_j replaced by a_j^ .*

Definition 5.9 (Run-specific approximate solver receipt). *At level m , an output $\hat{x}_m \in Q_m$ produced from $q_m \in Q_m$ has receipt (δ_m, β_m) when*

$$\Phi_m(\hat{x}_m) \leq \delta_m, \quad d_{B_m}(B_m \hat{x}_m, B_m q_m) \leq \beta_m.$$

Set

$$e_m(\delta, \beta) := \eta_m(\delta) + \omega_m(\beta + L_m \eta_m(\delta)).$$

Uniform receipts for a solver map mean that the two inequalities hold for every admissible input, but the comparison results below require only the displayed run-specific evidence.

Lemma 5.10 (Receipt-to-exact comparison). *If $\hat{x}_m$ has receipt (δ_m, β_m) relative to q_m , then*

$$d_m(\hat{x}_m, N_m q_m) \leq e_m(\delta_m, \beta_m). \tag{5.5}$$

Proof. The exact normal form has residual zero and preserves the observation of q_m . Apply the heterogeneous two-output estimate to $\hat{x}_m$ and $N_m q_m$, using $\eta_m(0) = 0$. $\square$

Proposition 5.11 (Same-level implementation agreement). *Let $n \leq m$. If $\hat{x}_m$ and $\hat{y}_m$ are outputs on the same input q_m with receipts (δ_m, β_m) and (δ'_m, β'_m) , then*

$$d_n(\rho_{m,n} \hat{x}_m, \rho_{m,n} \hat{y}_m) \leq K_{m,n}(e_m(\delta_m, \beta_m) + e_m(\delta'_m, \beta'_m)).$$

No refinement-defect term occurs in this same-level comparison.

Proof. Compare both outputs with the same exact state $N_m q_m$ using [Lemma 5.10](#), then apply the $K_{m,n}$ -Lipschitz restriction map. $\square$

For $n \leq h \leq r$ and $q_r \in Q_r$, define the amplified pathwise defect segment

$$A_{r,h}^{(n)}(q_r) := \sum_{j=h}^{r-1} K_{j,n} \alpha_j(\rho_{r,j+1} q_r),$$

with an empty sum equal to zero. By definition of a_j ,

$$A_{r,h}^{(n)}(q_r) \leq \sum_{j=h}^{r-1} K_{j,n} a_j.$$

Lemma 5.12 (Pathwise telescoping to an anchor). *For $n \leq h \leq r$ and $q_r \in Q_r$,*

$$d_n(\rho_{r,n}N_r q_r, \rho_{h,n}N_h \rho_{r,h} q_r) \leq \mathbf{A}_{r,h}^{(n)}(q_r).$$

Proof. For $j = h, \dots, r$, put $p_j = \rho_{j,n}N_j \rho_{r,j} q_r$. Consecutive terms satisfy

$$d_n(p_{j+1}, p_j) \leq K_{j,n} \alpha_j(\rho_{r,j+1} q_r).$$

Summing these inequalities proves the claim. $\square$

Corollary 5.13 (Fine-to-coarse solver certificate). *For $m \geq n$ and an output $\hat{x}_m$ with receipt (δ_m, β_m) relative to q_m ,*

$$d_n(\rho_{m,n}\hat{x}_m, N_n \rho_{m,n} q_m) \leq K_{m,n} e_m(\delta_m, \beta_m) + \mathbf{A}_{m,n}^{(n)}(q_m).$$

Proof. Use Lemma 5.10, apply $\rho_{m,n}$, and then use Lemma 5.12 with $h = n$. $\square$

Theorem 5.14 (Anchored cross-level solver comparison). *Let $n \leq h \leq \min\{m, \ell\}$. Let $q_m \in Q_m$ and $q_\ell \in Q_\ell$, and suppose $\hat{x}_m$ and $\hat{y}_\ell$ have receipts (δ_m, β_m) and $(\delta'_\ell, \beta'_\ell)$, respectively. Define the protected observation mismatch at the anchor by*

$$\gamma_h := d_{\mathcal{B}_h}(B_h \rho_{m,h} q_m, B_h \rho_{\ell,h} q_\ell).$$

Then

$$\begin{aligned} d_n(\rho_{m,n}\hat{x}_m, \rho_{\ell,n}\hat{y}_\ell) \\ \leq K_{m,n} e_m(\delta_m, \beta_m) + \mathbf{A}_{m,h}^{(n)}(q_m) + K_{h,n} \omega_h(\gamma_h) + \mathbf{A}_{\ell,h}^{(n)}(q_\ell) + K_{\ell,n} e_\ell(\delta'_\ell, \beta'_\ell). \end{aligned} \quad (5.6)$$

The two pathwise segments may be replaced by their uniform a_j bounds. Because the estimate holds for every admissible h , the right-hand side may be minimized over the possible anchors.

Proof. Compare each approximate output with its exact normal form using Lemma 5.10, and apply the corresponding restriction to level n . Use Lemma 5.12 to move both exact branches to level h . The two anchor normal forms are consistent, preserve the observations of their anchor inputs, and hence differ in Q_h by at most $\omega_h(\gamma_h)$. Applying $\rho_{h,n}$ contributes $K_{h,n}$. The triangle inequality gives (5.6). $\square$

Corollary 5.15 (Nested compatible solver levels). *Under the output and receipt hypotheses of Theorem 5.14, let $n \leq m \leq \ell$ and suppose*

$$B_m q_m = B_m \rho_{\ell,m} q_\ell.$$

Then

$$\begin{aligned} d_n(\rho_{m,n}\hat{x}_m, \rho_{\ell,n}\hat{y}_\ell) &\leq K_{m,n} e_m(\delta_m, \beta_m) + \mathbf{A}_{\ell,m}^{(n)}(q_\ell) \\ &\quad + K_{\ell,n} e_\ell(\delta'_\ell, \beta'_\ell). \end{aligned}$$

In particular, the middle term is at most $\sum_{j=m}^{\ell-1} K_{j,n} a_j$. The compatibility hypothesis holds when $q_m = \rho_{\ell,m} q_\ell$.

Proof. Take $h = m$ in Theorem 5.14. The first pathwise segment is empty and $\omega_m(0) = 0$ on the observation-determined consistent domain. $\square$

Corollary 5.16 (Cofinal projective implementation independence). *Assume every Q_n is complete and every restriction map is nonexpansive. Let $q = (q_n) \in \varprojlim Q_n$ be a compatible admissible input and suppose its pathwise exact-normalizer defects are summable:*

$$\sum_{j=0}^{\infty} \alpha_j(q_{j+1}) < \infty.$$

For $i \in \{1, 2\}$, let $\hat{x}_m^{(i)}$ be outputs at q_m whose receipt errors

$$e_m^{(i)} := e_m(\delta_m^{(i)}, \beta_m^{(i)}) \longrightarrow 0.$$

Then, for every fixed n , the limits

$$x_n = \lim_{m \rightarrow \infty} \rho_{m,n} \hat{x}_m^{(i)}$$

exist, are independent of i , and are unchanged by passage to a cofinal subsequence. The family $x = (x_n)$ belongs to $\varprojlim Q_n$, and

$$d_n(x_n, N_n q_n) \leq \sum_{j=n}^{\infty} \alpha_j(q_{j+1}).$$

If all pathwise normalizer defects vanish, then $x_n = N_n q_n$ for every n .

Proof. For $s \geq r \geq n$, compatibility gives $q_r = \rho_{s,r} q_s$. [Corollary 5.15](#) and nonexpansiveness give, for either pair of implementations,

$$d_n(\rho_{r,n} \hat{x}_r^{(i)}, \rho_{s,n} \hat{x}_s^{(i')}) \leq e_r^{(i)} + e_s^{(i')} + \sum_{j=r}^{s-1} \alpha_j(q_{j+1}).$$

The right-hand side tends to zero as $r, s \rightarrow \infty$. Completeness gives the coordinate limits, and the same estimate makes them independent of the implementation and of a cofinal subsequence. Continuity of the restriction maps passes compatibility to the limits. Finally, [Corollary 5.13](#) gives

$$d_n(\rho_{m,n} \hat{x}_m^{(i)}, N_n q_n) \leq e_m^{(i)} + \sum_{j=n}^{m-1} \alpha_j(q_{j+1});$$

let $m \rightarrow \infty$. □

The uniform hypotheses used earlier imply $e_m^{(i)} \rightarrow 0$ whenever $\delta_m^{(i)}, \beta_m^{(i)} \rightarrow 0$, the observation maps have a common Lipschitz bound, and $H(t), \Omega(t) \rightarrow 0$ at zero. Uniform summability of $\sum_j \alpha_j$ implies the pathwise summability condition. Thus raw receipt parameters alone are not enough: uniform conditioning and the amplified restriction factors remain load-bearing.

5.2 Exact inverse limits

Assume now that all one-step defects vanish and that $\rho_{n+1,n}(C_{n+1}) \subseteq C_n$. Let

$$Q_\infty := \varprojlim Q_n = \{(q_n) : \rho_{n+1,n} q_{n+1} = q_n\},$$

and similarly let $C_\infty = \varprojlim C_n$. Here Q_∞ is the inverse limit of the admissible domains fixed at the start of the section. It is the full ambient inverse limit only when every observation value is reconstructing and realizable.

Theorem 5.17 (Inverse-limit normalizer). *For each n , let $N_n : Q_n \rightarrow C_n \subseteq Q_n$ be an idempotent observation-preserving normalizer that fixes every point of C_n . Assume the restriction maps form an inverse system, preserve the consistent subsets, and satisfy*

$$\rho_{n+1,n} N_{n+1} = N_n \rho_{n+1,n} \quad (n \geq 0).$$

Assume also that the observation spaces carry compatible restriction maps $\beta_{n+1,n} : \mathcal{B}_{n+1} \rightarrow \mathcal{B}_n$ with

$$B_n \rho_{n+1,n} = \beta_{n+1,n} B_{n+1}.$$

Then

$$N_\infty((q_n)) = (N_n q_n)$$

defines an idempotent map $Q_\infty \rightarrow C_\infty$ preserving the induced inverse-limit observation. If, for every n and every $b_n \in B_n(C_n)$, the fiber $C_n \cap B_n^{-1}(b_n)$ is a singleton, then every realizable inverse-limit consistent observation fiber is a singleton.

Proof. Naturality makes $(N_n q_n)$ a compatible family. Because each coordinate lies in C_n and each N_n is idempotent, the resulting map lands in C_∞ and is idempotent. Compatibility of the B_n makes observation preservation coordinatewise. If two inverse-limit consistent families have the same inverse-limit observation, their n th coordinates lie in the same finite consistent observation fiber for every n , hence agree at every coordinate. $\square$

Remark 5.18 (What the theorem does not prove). The theorem supplies a normalizer on an already defined inverse system. It does not prove compactness, nonemptiness, a continuum dynamics, or adequacy of the chosen refinement maps. Those are separate mathematical or model-specific obligations.

6 Repairability and selection obstructions

The existence of a unique consistent state with a given observation does not imply that a prescribed local write support can construct it. This section separates extension, uniqueness, and repair support.

Let X_W be an interior write space, let X_D be a protected collar space, and let

$$R \subseteq X_W \times X_D$$

be the locally consistent relation.

Definition 6.1 (Strong local repair). A *strong local repair* is a map

$$R : X_W \times X_D \longrightarrow R$$

such that

- (i) $\pi_D R(x, d) = d$ for every (x, d) ;
- (ii) $R(r) = r$ for every $r \in R$.

Thus every local state is repaired, the collar is preserved, and already consistent states are fixed.

Theorem 6.2 (Collar-section criterion). *For nonempty finite X_W and finite X_D , a strong local repair exists if and only if the collar projection*

$$\pi_D : R \longrightarrow X_D$$

is surjective.

Proof. A repair sends every (x, d) to a point of R with collar d , so surjectivity is necessary. Conversely, choose a section $s : X_D \rightarrow R$ of π_D and define

$$R(z) = \begin{cases} z, & z \in R, \\ s(\pi_D z), & z \notin R. \end{cases}$$

This is a strong repair. □

If $R = \emptyset$ while $X_W \times X_D \neq \emptyset$, nonexistence of a repair is immediate; when the domain is empty, the unique empty map is a vacuous repair. The next proposition assumes $R \neq \emptyset$ because ordinary real-valued distance to the empty set is convention-dependent; an extended-distance formulation may instead assign that case margin $+\infty$.

Corollary 6.3 (Finite no-repair certificate). *For nonempty X_W , a single collar value $d \notin \pi_D(R)$ certifies that no total exact repair preserving that collar exists.*

Proposition 6.4 (Robust no-repair margin). *Suppose X_W and X_D are metric spaces, $\emptyset \neq R \subseteq X_W \times X_D$, and the product metric dominates collar distance:*

$$d_{X_W \times X_D}((w, d), (w', d')) \geq d_{X_D}(d, d').$$

For $d \in X_D$, set

$$\mu_R(d) := \text{dist}_{X_D}(d, \pi_D(R)).$$

Then for every $w \in X_W$,

$$\text{dist}_{X_W \times X_D}((w, d), R) \geq \mu_R(d).$$

If R is compact and $d \notin \pi_D(R)$, then $\mu_R(d) > 0$. Hence no procedure whose output preserves the collar value d exactly can approximate consistency to error below $\mu_R(d)$.

Proof. For every $(w', d') \in R$, the product-metric assumption gives

$$d((w, d), (w', d')) \geq d_{X_D}(d, d') \geq \mu_R(d).$$

Take the infimum. Compactness makes $\pi_D(R)$ compact and therefore closed, so a point outside it has positive distance. □

Definition 6.5 (Repair width). Let $\mathcal{W}$ be a finite family of admissible write supports with a cost function $w : \mathcal{W} \rightarrow \mathbb{R}_{\geq 0}$. For each $W \in \mathcal{W}$, let $R_W \subseteq X_W \times X_{D(W)}$ be the induced local consistency relation and let $D_{\text{adm}}(W) \subseteq X_{D(W)}$ be the declared admissible collar domain. The *repair width* is

$$\min\{w(W) : D_{\text{adm}}(W) \subseteq \pi_{D(W)}(R_W)\},$$

with value $+\infty$ when no admissible support has the required projection property.

The definition prevents a one-site no-go from being misreported as a global inconsistency: increasing the transaction support may restore surjectivity.

6.1 Finite observation-conditioned resampling

There is nevertheless a canonical *averaging* operation once a reference measure is supplied; unlike a section, it does not select one state from an ambiguous fiber.

Theorem 6.6 (Fiber averaging is conditional expectation). *Let Q be finite, let $\mu(x) > 0$ with $\sum_x \mu(x) = 1$, and let $B : Q \rightarrow \mathcal{B}$. For $b \in B(Q)$ put $F_b = B^{-1}(b)$ and $\mu(F_b) = \sum_{y \in F_b} \mu(y)$. Define*

$$(P_B f)(x) = \frac{1}{\mu(F_{B(x)})} \sum_{y \in F_{B(x)}} \mu(y) f(y).$$

Then $P_B = \mathbb{E}_\mu[\cdot \mid \sigma(B)]$. In the weighted inner product $\langle f, g \rangle_\mu = \sum_x \mu(x) f(x) g(x)$ it satisfies

$$P_B^2 = P_B, \quad P_B^* = P_B, \quad \|P_B f\|_{2,\mu} \leq \|f\|_{2,\mu},$$

and its fixed space is exactly

$$\mathcal{M}_B = \{f \in \mathbb{R}^Q : B(x) = B(y) \Rightarrow f(x) = f(y)\},$$

the observation-measurable functions.

Moreover, a Markov matrix K on Q equals the matrix of P_B if and only if the following finite receipt holds:

$$K(x, y) = 0 \quad \text{when } B(x) \neq B(y), \quad (\text{R1})$$

$$K(x, \cdot) = K(x', \cdot) \quad \text{when } B(x) = B(x'), \quad (\text{R2})$$

$$\mu(x) K(x, y) = \mu(y) K(y, x) \quad \text{for all } x, y. \quad (\text{R3})$$

Equivalently, its entries are

$$K(x, y) = \mathbf{1}_{\{B(x)=B(y)\}} \frac{\mu(y)}{\mu(F_{B(x)})}.$$

Proof. The displayed average is constant on each fiber and fixes every fiber-constant function, which proves the conditional-expectation and fixed space statements. Grouping both variables by fibers gives

$$\langle P_B f, g \rangle_\mu = \sum_{b \in B(Q)} \frac{(\sum_{x \in F_b} \mu(x) f(x)) (\sum_{y \in F_b} \mu(y) g(y))}{\mu(F_b)},$$

which is symmetric in f, g . Hence P_B is an orthogonal projection and is L^2 -contractive. Its matrix plainly satisfies (R1)–(R3). Conversely, (R1)–(R2) give one probability row k_b on each fiber. By (R3), $\mu(x) k_b(y) = \mu(y) k_b(x)$ there, so normalization forces $k_b(y) = \mu(y)/\mu(F_b)$. $\square$

Remark 6.7. [Theorem 6.6](#) recognizes the exact finite resampling projector from checkable matrix identities. It neither chooses a representative of a fiber nor supplies a spectral gap or a convergence rate for a different repair dynamics; those require additional structure. A noncircular receipt enumerates the states and fibers, derives K independently from the declared transition table or transition counts, and then verifies (R1)–(R3); constructing K from the target entry formula is not a verification.

6.2 Symmetry and hidden tie-breaking

Let a group G act on C and $\mathcal{B}$, and assume $B|_C$ is equivariant.

Theorem 6.8 (Equivariant-section criterion). *Let a finite group G act on C and $B(C)$, and let $B|_C : C \rightarrow B(C)$ be an equivariant surjection. A G -equivariant section $s : B(C) \rightarrow C$ exists if and only if, for one (equivalently every) representative b of each G -orbit in $B(C)$, the fiber C_b contains a point fixed by the stabilizer*

$$G_b = \{g \in G : gb = b\}.$$

Proof. Necessity follows from $gs(b) = s(gb) = s(b)$ for $g \in G_b$. Conversely, choose an orbit representative b and a G_b -fixed point $c_b \in C_b$. Define $s(gb) = gc_b$. If $gb = g'b$, then $g^{-1}g' \in G_b$, so $gc_b = g'c_b$; hence the definition is well posed. It is visibly a section and equivariant. Repeat independently on each orbit. $\square$

Corollary 6.9 (Stabilizer obstruction). *If some fiber C_b has no G_b -fixed point, there is no equivariant boundary-only selector on that orbit. In particular, a transitive ambiguous fiber with no fixed point cannot be canonically collapsed without adding symmetry-breaking data.*

Remark 6.10 (No selection by normalization). Lexicographic order, an energy, a prior, maximum entropy, or an external sector label can select an ambiguous fiber. None is entailed by consistency and the observation alone. The theorem is elementary, but it is a useful firewall: a normal-form map cannot silently become a physical or probabilistic selector without additional structure.

7 Ranked functional-audit systems

We now give a broad constructive class in which the partial normalizer, asynchronous dynamics, and obstruction output are all explicit.

Definition 7.1 (Ranked functional-audit system). Let

$$V = V_0 \sqcup V_1 \sqcup \cdots \sqcup V_m$$

be a finite ranked site set with $V_0 \neq \emptyset$. Every site $v \in V$ has a nonempty finite state set X_v , and sites in V_0 are boundary sites. Every generated site $v \in V_k$, $k > 0$, has a parent set

$$P(v) \subseteq V_0 \cup \cdots \cup V_{k-1},$$

and a function

$$F_v : \prod_{u \in P(v)} X_u \longrightarrow X_v.$$

A state satisfies the generating constraints when

$$x_v = F_v(x|_{P(v)}) \quad (v \notin V_0).$$

A finite family of audits $A_j(x) \in \{0, 1\}$ imposes residual global constraints. A state is globally consistent when it satisfies all generating constraints and every audit returns 1.

Write

$$Q = \prod_{v \in V} X_v, \quad B : Q \longrightarrow \prod_{v \in V_0} X_v$$

for the full state space and the boundary-coordinate projection. For a boundary value $b \in \prod_{v \in V_0} X_v$, define $E(b)$ recursively by rank.

Theorem 7.2 (Singleton-or-empty fiber). *For every boundary b , exactly one of the following holds:*

- (i) *every audit passes on $E(b)$ and the consistent fiber is $\{E(b)\}$;*
- (ii) *some audit fails on $E(b)$ and the consistent fiber is empty.*

Proof. Rank recursion constructs one state satisfying all generating constraints. Any other such state with boundary b agrees at rank 0 and therefore agrees inductively at every rank. The audits either accept this sole candidate or eliminate it. $\square$

Definition 7.3 (Asynchronous functional update). For $v \notin V_0$, let T_v replace x_v by $F_v(x|_{P(v)})$ and leave every other coordinate unchanged. An update is *effective* when it changes x_v .

Define recursively

$$M(v) = 0 \quad (v \in V_0), \quad M(v) = 1 + \sum_{u \in P(v)} M(u) \quad (v \notin V_0).$$

Theorem 7.4 (Finite asynchronous normalization with an update bound). *In every sequence of effective asynchronous updates, site v is updated at most $M(v)$ times. Consequently every such sequence has length at most*

$$\sum_{v \in V} M(v),$$

and every maximal sequence terminates at $E(Bx)$ independently of update order.

Proof. Proceed by induction on rank. Boundary sites never change. Assume the bound holds below rank k . The parent tuple of a rank- k site can change at most $\sum_{u \in P(v)} M(u)$ times. Before the first parent change, between consecutive parent changes, and after the last parent change, at most one effective update at v is possible on each fixed-parent interval: after that update, v equals the value prescribed by the fixed parent tuple. Hence v changes at most $M(v)$ times. Summing over sites proves termination.

At a maximal endpoint no functional update is enabled, so all generating equations hold. The endpoint has the original boundary and therefore equals the unique generated extension $E(Bx)$. $\square$

Corollary 7.5 (Synchronous settling by dependency depth). *Let $x^{(0)} = x$. In one synchronous round, keep the boundary coordinates fixed and replace every generated coordinate using the values from the previous round:*

$$x_v^{(k+1)} = \begin{cases} x_v^{(k)}, & v \in V_0, \\ F_v(x^{(k)}|_{P(v)}), & v \notin V_0. \end{cases}$$

For $0 \leq k \leq m$, after round k every site in $V_0 \cup \dots \cup V_k$ agrees with $E(Bx)$. Consequently $x^{(m)} = E(Bx)$. This is a depth bound for the explicit ranked system; correspondence with a separately specified compiler or update rule requires its own theorem.

Proof. Induct on k . Boundary values are fixed. If all sites of rank at most $k-1$ are correct after round $k-1$, then every parent of a rank- k site is correct, so its round- k update evaluates the defining function at the same parent tuple as $E(Bx)$. Sites of lower rank remain correct because their parents were already correct in the preceding round. $\square$

Corollary 7.6 (Acyclic Boolean circuit compiler). *Every nonempty finite acyclic Boolean circuit compiles to a ranked functional system with one $\{0,1\}$ -valued site per input or gate. Inputs and constant gates form V_0 , with the constant coordinates restricted to their prescribed values; every other gate is ranked by its longest source-to-gate path, has its incoming wires as parents, and uses its truth function as F_v . The extension $E(b)$ is ordinary circuit evaluation on the admissible input b . Updating all nonconstant gate sites in parallel from the preceding round reaches $E(b)$ in at most the circuit depth, independently of the initial gate-register values.*

Proof. The topological ranking makes every nonconstant gate depend only on lower ranks, so the compilation satisfies the definition on the boundary domain where constants have their prescribed values. Recursive evaluation is E , and the depth bound is [Corollary 7.5](#). $\square$

Corollary 7.7 (Certified solver). *The procedure “settle the generating equations asynchronously, then evaluate the audits” returns either the unique consistent extension or a named failed audit. It cannot confuse an empty fiber with nontermination.*

Theorem 7.8 (Propagation of boundary perturbations). *Equip every X_v with a metric d_v . Suppose that for each generated site v there are coefficients $L_{vu} \geq 0$ such that*

$$d_v(F_v(p), F_v(p')) \leq \sum_{u \in P(v)} L_{vu} d_u(p_u, p'_u).$$

For a boundary site $r \in V_0$, define coefficients recursively by

$$K_{rr} = 1, \quad K_{vr} = 0 \ (v \in V_0, v \neq r), \quad K_{vr} = \sum_{u \in P(v)} L_{vu} K_{ur} \ (v \notin V_0).$$

Then for all boundary values b, b' and every site v ,

$$d_v(E(b)_v, E(b')_v) \leq \sum_{r \in V_0} K_{vr} d_r(b_r, b'_r).$$

With the ℓ^1 product metrics,

$$d_Q(E(b), E(b')) \leq \Lambda_E d_B(b, b'), \quad \Lambda_E := \max_{r \in V_0} \sum_{v \in V} K_{vr}.$$

Thus the generated extension supplies an explicit inverse-observation modulus $\omega_B(t) \leq \Lambda_E t$ on every accepted audit branch.

Proof. The sitewise estimate follows by induction on rank. Boundary sites satisfy it by definition. At a generated site, apply the local Lipschitz inequality and substitute the induction hypotheses; regrouping coefficients gives the recursion for K_{vr} . Summing over v and then bounding each boundary coefficient by Λ_E proves the ℓ^1 estimate. $\square$

Remark 7.9. The local updates need not commute when one site depends on another. Order independence follows from well-founded dependency and unique extension, not from a global pairwise-commutation hypothesis. [Theorem 7.8](#) additionally exposes the amplification carried by every dependency path.

8 Linear quotient systems and observability

Let V, W, Z be finite-dimensional real inner-product spaces. Let

$$A : V \rightarrow Z, \quad C = \ker A, \quad G \subseteq C$$

be a gauge subspace. Let $B : V \rightarrow W$ vanish on G , so it descends to V/G . Equip V/G with the quotient norm.

Proposition 8.1 (Linear information-boundary criterion). *The induced map $C/G \rightarrow W$ is injective if and only if*

$$\ker(B|_C) = G.$$

Equivalently,

$$\text{rank}(B|_C) = \dim C - \dim G.$$

The same statement holds over any field.

Proof. The kernel of the induced map is $\ker(B|_C)/G$. □

Let

$$H = C \cap G^\perp.$$

Under the injectivity hypothesis, $B|_H$ is injective. Let

$$\kappa_A := \|A^\dagger\|, \quad \kappa_B := \|(B|_H)^\dagger\|,$$

where $\dagger$ denotes the Moore–Penrose pseudoinverse. Thus, when $H \neq \{0\}$, κ_B is the reciprocal of the smallest singular value of $B|_H$; on the zero space the pseudoinverse convention gives $\kappa_B = 0$. The constant κ_A is the global linear error-bound constant for $C = \ker A$.

Theorem 8.2 (Linear stability certificate). *For all $x, y \in V$,*

$$\begin{aligned} d_{V/G}([x], [y]) &\leq \kappa_A(\|Ax\| + \|Ay\|) \\ &\quad + \kappa_B(\|B(x - y)\| + \|B\|\kappa_A(\|Ax\| + \|Ay\|)). \end{aligned} \tag{8.1}$$

Proof. Let $c_x = P_C x$ and $c_y = P_C y$ be the orthogonal projections onto C . The pseudoinverse estimate gives

$$\|x - c_x\| \leq \kappa_A \|Ax\|, \quad \|y - c_y\| \leq \kappa_A \|Ay\|.$$

Let $h_x = P_H c_x$ and $h_y = P_H c_y$. Since $C = G \oplus H$, the vectors h_x, h_y are the minimal-norm representatives of the quotient classes $[c_x], [c_y]$. Because B vanishes on G ,

$$B(h_x - h_y) = B(c_x - c_y).$$

Hence

$$\|B(h_x - h_y)\| \leq \|B(x - y)\| + \|B\|(\|x - c_x\| + \|y - c_y\|),$$

and injectivity of $B|_H$ gives

$$d_{C/G}([c_x], [c_y]) = \|h_x - h_y\| \leq \kappa_B \|B(c_x - c_y)\|.$$

Finally,

$$d_{V/G}([x], [c_x]) \leq \|x - c_x\|, \quad d_{V/G}([y], [c_y]) \leq \|y - c_y\|,$$

so the quotient triangle inequality yields (8.1). □

The finite-field version retains the rank criterion but not singular values. It is the quotient form of an information-set condition from coding theory [22].

8.1 Finite linear dynamics

Let $T : U \rightarrow U$ be a linear update, let $S : U \rightarrow Y$ be a sensor, and let $G \subseteq U$ satisfy $T(G) \subseteq G$ and $S(G) = 0$. Equip Y^{k+1} with its product inner product. The length- k observation map is

$$\mathcal{O}_k x = (Sx, STx, \dots, ST^k x).$$

Corollary 8.3 (Observability modulo gauge). *The time tube $\mathcal{O}_k$ reconstructs U/G if and only if*

$$\ker \mathcal{O}_k = G.$$

Over $\mathbb{R}$ or $\mathbb{C}$ and when $G^\perp \neq \{0\}$, the smallest singular value of $\mathcal{O}_k$ on $G^\perp$ is the separation margin; its reciprocal (equivalently, the norm of the restricted pseudoinverse) is the inverse-stability constant. On the zero space only the pseudoinverse convention is used. The static Rule-90 example in [Section 10](#) instead instantiates the earlier finite-field quotient rank criterion; it is not asserted to instantiate this time-tube corollary.

This is the classical observability matrix viewed through the quotient-boundary language; no novelty is claimed for the control-theoretic rank test [\[24\]](#).

9 Complexity of succinct boundary problems

The finite theory is elementary when Q and C are listed explicitly. It becomes computationally nontrivial when consistency is given by a Boolean circuit or formula.

Definition 9.1 (Succinct Boolean boundary instance). A succinct instance consists of a Boolean circuit $F : \{0, 1\}^n \rightarrow \{0, 1\}$ and a coordinate set $S \subseteq \{1, \dots, n\}$. The consistent set is $C = F^{-1}(1)$ and the observation is the projection $B_S(x) = x|_S$.

Theorem 9.2 (Boundary realizability is NP-complete). *Given (F, S, b) , deciding whether*

$$\{x : F(x) = 1, B_S(x) = b\}$$

is nonempty is NP-complete.

Proof. A satisfying extension is a polynomial witness. SAT reduces to the problem by taking $S = \emptyset$. $\square$

Theorem 9.3 (Boundary determination is coNP-complete). *Given (F, S) , deciding whether $B_S|_C$ is injective is coNP-complete, even when S contains all but one variable.*

Proof. Non-injectivity has an NP certificate: two distinct satisfying assignments agreeing on S . For coNP-hardness, reduce UNSAT while keeping the consistent set nonempty. Given a formula $\varphi(x_1, \dots, x_m)$, introduce bits z, t and define

$$F(x, z, t) = (\neg t \wedge [x = 0^m] \wedge \neg z) \vee (t \wedge \varphi(x)), \quad S = \{x_1, \dots, x_m, t\}.$$

The baseline assignment $(0^m, 0, 0)$ is always consistent. If φ is satisfiable, any satisfying x gives the two distinct consistent assignments $(x, 0, 1)$ and $(x, 1, 1)$ with the same boundary. If φ is unsatisfiable, the baseline is the only consistent assignment, so the boundary restriction is injective. Thus boundary determination is equivalent to unsatisfiability even for nonempty C , and S observes all but one variable. $\square$

Theorem 9.4 (Strong repair existence is Π_2^P -complete). *Let a Boolean circuit $R(w, d)$ represent a local consistency relation on write bits w and protected collar bits d . Deciding whether an unrestricted total strong collar-preserving repair (equivalently, a set-theoretic section of the collar projection) exists is Π_2^P -complete.*

Proof. By [Theorem 6.2](#), a strong repair exists exactly when

$$\forall d \exists w \quad R(w, d) = 1.$$

This is in Π_2^P and is Π_2^P -hard by the canonical quantified-Boolean-formula problem with one universal and one existential block [\[33, 35\]](#). $\square$

Remark 9.5. These results explain why an abstract finite normalizer can be mathematically canonical yet computationally difficult. The ranked functional class avoids the hardness by supplying a directed generative presentation; the linear class avoids it through rank computation.

Proposition 9.6 (Explicit-table algorithms). *If Q , C , B , d_Q , d_B , and Φ are given by explicit finite tables, then:*

- (i) *the fiber trichotomy and partial normalizer are computable in time $O(|Q| + |C|)$ after hashing observations;*
- (ii) *ω_B is computable from $O(|C|^2)$ pairs;*
- (iii) *η_Φ is computable from all-pairs distances to C and a sort by residual value.*

10 Examples and separation results

Example 10.1 (Confluence does not imply observable reconstruction). Let $Q = C = \{a, b\}$ with the empty rewrite relation and a constant observation map. Every state is a normal form and the relation is confluent, but the sole observation fiber contains two consistent states.

Example 10.2 (A two-bit positive witness). Let $Q = \mathbb{F}_2^2$, let $B(u, v) = u$, and put $C = \{(0, 0), (1, 0)\}$. For each $u \in \mathbb{F}_2$, take the step $(u, 1) \rightarrow (u, 0)$. It preserves B , the normal forms are exactly C , and every source (u, v) terminates at $(u, 0)$. Thus sources with the same observation have the same endpoint, jointly witnessing normalization, same-source confluence, and the stronger cross-source property of [Theorem 3.8](#) for equality.

Example 10.3 (Observable determination does not imply realizability). Let $C = \{c\} \subset Q$ and let $B = \{0, 1\}$ with $B(c) = 0$. The restriction $B|_C$ is injective, but the fiber over 1 is empty.

Example 10.4 (Termination does not imply completeness). Let $Q = \{x, n\}$, let $x \rightarrow n$, and declare $C = \{x\}$. The relation terminates, but its only normal form is not consistent.

Example 10.5 (A singleton fiber does not replace observation preservation). Let $Q = \{x, c, d\}$, let $C = \{c, d\}$, set $B(x) = B(c) = 0$ and $B(d) = 1$, and take $x \rightarrow d$ as the sole rewrite step. Then the normal forms are exactly C and the consistent fiber over 0 is the singleton $\{c\}$. Nevertheless, the boundary-0 state x normalizes to d , not to c , because the rewrite changes the observation. Thus the observation-preservation hypothesis in [Corollary 3.11](#) is load-bearing.

10.1 Rule 90 as a regression test

No cellular automaton is used in the general theory. Rule 90 is retained because a three-cell instance simultaneously exhibits a hidden kernel, a proper information boundary, a deficient boundary, and a reverse-repair obstruction.

Over $\mathbb{F}_2$, define

$$F(a, b, c) = (b, a + c, b).$$

A one-step spacetime record is a pair (s, t) satisfying $t = F(s)$.

Proposition 10.6 (Width-three Rule 90). *The kernel and image are*

$$\ker F = \text{span}\{(1, 0, 1)\}, \quad \text{im } F = \{(u, v, u) : u, v \in \mathbb{F}_2\}.$$

The coordinate observation $B_{01}(u, v, w) = (u, v)$ is injective on $\text{im } F$, while $B_{02}(u, v, w) = (u, w)$ is not. No total exact repair can preserve an arbitrary target row and alter only the seed.

Proof. The kernel and image follow immediately from the formula for F . On the image, the first two coordinates determine (u, v, u) , whereas the first and third coordinates repeat the same bit. A reverse repair preserving every target would require F to be surjective by [Theorem 6.2](#); for example $(0, 0, 1)$ is not in the image. $\square$

The accompanying standalone Lean artifact proves these exact width-three statements and includes theorem-level axiom-audit commands. Rule 90 is a regression test, not a rule selected by the theory or a load-bearing assumption.

11 Formal artifact and coverage

The accompanying `ObservableNormalForms` directory is a standalone Lean 4 project. It pins Lean and Mathlib at version 4.29.1, includes a pinned, checksummed Lake manifest, and builds both on its own and as a library of the parent formalization repository. Its theorem-bearing files contain no `sorry` or `admit`. The recorded `#print axioms` reports contain only standard Mathlib axioms—`propext`, `Classical.choice`, and `Quot.sound` where required—and no `sorryAx`.

Module	Machine-checked coverage
Exact	Proof-carrying partial normalizers, the $(a) \Leftrightarrow (b)$ partial-normalizer equivalence, reachable observable fibers, the empty/singleton alternatives, audited terminals, and the observation-leak counterexample.
ObserverConfluence	The equivalence in Theorem 3.8 , weak-normalization endpoint existence, the matching two-bit positive system, and a coarse-sensor counterexample separating same-source confluence from cross-source endpoint uniqueness.
Stability	The heterogeneous and symmetric output estimates, approximate schedule independence, and a sensor-enrichment certificate.
Refinement	One-step and exact naturality, arbitrary-depth metric telescoping, same-level agreement without a tower term, fine-to-coarse comparison, and the anchored and nested metric cores.
Repair	The collar-section criterion, no-repair and robust-margin certificates, and counterexamples showing why the nonempty-domain hypotheses are necessary.
ConditionalResampling	The finite algebraic characterization of the weighted fiber kernel in Theorem 6.6 , observation-measurable fixed space, idempotence, weighted self-adjointness, a Pythagorean identity, and weighted L^2 contraction, together with the exact R1–R3 matrix-recognition equivalence.
Functional	Synchronous settling through dependency rank, fixed-point behavior, and generated-extension uniqueness.
Stochastic	Finite Markov affine-drift iteration and the corresponding one-time finite-state tail estimate.
Examples/Rule90	Exact width-three kernel, image, information-boundary, deficient-boundary, and reverse-repair obstruction results.

The artifact’s proof index gives theorem-level names and labels; its build receipt records the parent and standalone commands, tool versions, and axiom audit. This is substantial coverage of the finite core, but it is not a complete mechanization of the manuscript. In particular, the intrinsic-modulus minimality and sharpness results, compact and uniform extensions, probability-event and occupation wrappers, the Hausdorff perturbation wrapper, dependent cofinal and inverse-limit constructions, equivariant sections, the circuit-specialization wrapper, asynchronous path-gain estimate, linear-algebraic certificates, and complexity classifications remain paper proofs. The status distinction is deliberate: only declarations listed in the proof index are claimed as machine checked.

12 Discussion

12.1 Why the factorization is useful

The estimate

$$\text{residual} \xrightarrow{\eta} \text{distance to consistency} \xrightarrow{\omega} \text{observable uniqueness}$$

separates two burdens that are often mixed. A rewrite algorithm, numerical solver, or physical implementation controls η by showing that a small residual means closeness to C . The observation architecture controls ω by showing that $B|_C$ is not merely injective but well conditioned. Critical-pair arguments may help establish exact completion or small residuals, but they are not required in the final output comparison.

12.2 Why finite uniqueness is not enough

In a refinement program, proving $B_r|_{C_r}$ injective for every r does not prevent the inverse condition number from diverging. [Example 4.15](#) is deliberately minimal. The correct limit-facing statement is uniform stable separation, not a list of stagewise singleton claims. This is the main stability requirement emphasized for refinement-sensitive applications.

12.3 Limitations

The theory assumes that a meaningful quotient, consistent set, residual, observation metric, and refinement map have already been declared. It does not select those objects. In particular:

- a normalizer is not a probability measure;
- an inverse-limit normalizer is not a continuum field theory;
- a finite update bound is not a wall-clock liveness theorem;
- a boundary pseudometric does not acquire additional semantics without a separate interpretation theorem;
- a small Rule-90 example does not establish cellular-automaton universality or generality.

These exclusions are necessary for both mathematical correctness and journal review.

13 Artifact and data availability

The repository artifact contains the PDF, LaTeX source, bibliography, and the standalone `ObservableNormalForms` Lean project. The formal directory includes a pinned toolchain and dependency manifest, theorem-to-source index, axiom audit, reproducible build receipt, archive manifest, and SHA-256 checksums. The mathematical proofs in this manuscript are self-contained; the Lean artifact checks the subset stated in [Section 11](#). No empirical dataset is used.

14 Conclusion

Confluence is about paths from one state. Observable normalization is about what partial data determine after consistency is enforced. Their exact intersection is classical unique-normal-form theory: under observation preservation and $\text{NF}(\rightarrow) = C$, identification modulo any silent equivalence is equivalent to agreement of normal endpoints across all sources with the same observation. This does not by itself provide normalization or all-schedule settlement. Their robust intersection is governed by two analytic objects: an error bound to the consistent set and an inverse modulus for the observation map on that set.

The resulting estimate is universal and sharp in its two unavoidable components. More importantly, it survives refinement: one-step defects telescope, same-level solver agreement needs no tower term, and compatible cross-level computations differ by their receipt errors and the intervening normalizer-defect tail. For a compatible admissible input in complete spaces, nonexpansive restrictions, vanishing receipt errors, and summable pathwise tails give a common projective limit; exact compatibility recovers the inverse-limit normalizer. On finite probability spaces, fiber averaging supplies the canonical conditional-expectation projector, while its matrix receipt keeps that

averaging distinct from representative selection or a spectral-gap claim. Empty fibers and ambiguous fibers remain visible rather than being hidden inside a repair operator. The theory therefore supplies a precise answer to three different questions:

Does a consistent extension exist?
 Is it observably unique and stable?
 Can the declared local support construct it?

No one of these questions answers the others. Keeping them separate is the central mathematical lesson and the reason the framework applies across rewrite, reconstruction, and distributed-inference models.

Statements and Declarations

Funding declaration: The authors declare that no external grants or dedicated third-party research funding were received for the preparation of this manuscript.

Author affiliations declaration: Bernhard Mueller is affiliated with Pragma Research Inc. No institutional affiliations were declared for David Matscheko or Jonathan Hill in this submission. The corresponding-author email address is provided in the correspondence declaration.

Consent to Participate declaration: not applicable.

Consent to Publish declaration: not applicable.

Author Contribution declaration: Bernhard Mueller developed the observation-determined normal-form framework, the stability, refinement, repair, and complexity results, the manuscript text, and the accompanying Lean artifact. David Matscheko and Jonathan Hill contributed proof review, checking of the exact and stability results, and consistency review of the formal artifact and manuscript. All authors reviewed and approved the manuscript.

Correspondence declaration: Correspondence should be addressed to Bernhard Mueller, bernhard@floatingpragma.ai.

Data Availability declaration: No empirical dataset is used in this work. The LaTeX source, bibliography, and the standalone `ObservableNormalForms` Lean 4 project—with a pinned toolchain and dependency manifest, a theorem-to-source index, an axiom audit, a reproducible build receipt, and SHA-256 checksums—are available through the project repository.

Code Availability declaration: The standalone `ObservableNormalForms` Lean 4 project and its build sources are maintained in the project repository, together with a source bundle so that the manuscript and the formal artifact can be rebuilt without relying on private local paths.

Ethics declaration: not applicable.

Competing Interests declaration: The authors declare no competing financial or non-financial interests for this manuscript.

Use of AI-assisted tools declaration: AI-assisted tools were used for language editing, drafting assistance, formatting support, and submission-package preparation. The authors reviewed and remain responsible for all scientific content, claims, references, and declarations.

References

- [1] Samson Abramsky. A generalized kahn principle for abstract asynchronous networks. In *Mathematical Foundations of Programming Semantics*, volume 442 of *Lecture Notes in Computer Science*, pages 1–21, 1990. doi: 10.1007/BFb0040252.

- [2] Samson Abramsky and Adam Brandenburger. The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13:113036, 2011. doi: 10.1088/1367-2630/13/11/113036.
- [3] Takahito Aoto and Yoshihito Toyama. Automated proofs of unique normal forms w.r.t. conversion for term rewriting systems. In *Frontiers of Combining Systems (FroCoS 2019)*, volume 11715 of *Lecture Notes in Computer Science*, pages 330–347, 2019. doi: 10.1007/978-3-030-29007-8_19.
- [4] Marcelo Arenas, Leopoldo Bertossi, and Jan Chomicki. Consistent query answers in inconsistent databases. In *Proceedings of the Eighteenth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 68–79, 1999. doi: 10.1145/303976.303983.
- [5] Sam Arkle and Andrew Polonsky. A formalization of abstract rewriting in agda, 2026.
- [6] Samuel Arsac, Russ Harmer, and Damien Pous. Adhesive category theory for graph rewriting in rocq. In *Proceedings of the 15th ACM SIGPLAN International Conference on Certified Programs and Proofs*, pages 59–74, 2026. doi: 10.1145/3779031.3779105.
- [7] Albert Atserias and Phokion G. Kolaitis. Consistency, acyclicity, and positive semirings. In *Samson Abramsky on Logic and Structure in Computer Science and Beyond*, pages 623–668. Springer, 2023. doi: 10.1007/978-3-031-24117-8_17.
- [8] Franz Baader and Tobias Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998. doi: 10.1017/CBO9781139172752.
- [9] Catriel Beeri, Ronald Fagin, David Maier, and Mihalis Yannakakis. On the desirability of acyclic database schemes. *Journal of the ACM*, 30(3):479–513, 1983. doi: 10.1145/2402.322389.
- [10] Frédéric Blanqui and Adam Koprowski. CoLoR: A Coq library on well-founded rewrite relations and its application to the automated verification of termination certificates. *Mathematical Structures in Computer Science*, 21(4):827–859, 2011. doi: 10.1017/S0960129511000120.
- [11] Graham Campbell and Detlef Plump. Confluence up to garbage in graph transformation. *Theoretical Computer Science*, 884:1–22, 2021. doi: 10.1016/j.tcs.2021.06.010.
- [12] Daizhan Cheng, Zhiqiang Li, and Hongsheng Qi. Realization of boolean control networks. *Automatica*, 46(1):62–69, 2010. doi: 10.1016/j.automatica.2009.10.036.
- [13] Henning Christiansen and Maja H. Kirkeby. On proving confluence modulo equivalence for constraint handling rules. *Formal Aspects of Computing*, 29:57–95, 2017. doi: 10.1007/s00165-016-0396-9.
- [14] Radek Cibulka, Asen L. Dontchev, and Alexander Y. Kruger. Strong metric subregularity of mappings in variational analysis and optimization. *Journal of Mathematical Analysis and Applications*, 457:1247–1282, 2018. doi: 10.1016/j.jmaa.2016.11.045.
- [15] Victor Dalmau and Jakub Opršal. Local consistency as a reduction between constraint satisfaction problems. In *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 1–15, 2024. doi: 10.1145/3661814.3662068.
- [16] Asen L. Dontchev and R. Tyrrell Rockafellar. *Implicit Functions and Solution Mappings: A View from Variational Analysis*. Springer, 2 edition, 2014. doi: 10.1007/978-1-4939-1037-3.

- [17] Dominique Duval, Rachid Echahed, and Frédéric Prost. Categorical abstract rewriting systems and functoriality of graph transformation. *Electronic Communications of the EASST*, 41, 2011. doi: 10.14279/tuj.eceasst.41.612.
- [18] Daniel Gall and Thom Frühwirth. Confluence modulo equivalence with invariants in constraint handling rules. In *Functional and Logic Programming (FLOPS 2018)*, volume 10818 of *Lecture Notes in Computer Science*, pages 116–131, 2018. doi: 10.1007/978-3-319-90686-7_8.
- [19] Francesco Gavazzo and Cecilia Di Florio. Quantitative and metric rewriting: Abstract, non-expansive, and graded systems. *arXiv preprint arXiv:2206.13610*, 2022. doi: 10.48550/arXiv.2206.13610.
- [20] Alan J. Hoffman. On approximate solutions of systems of linear inequalities. *Journal of Research of the National Bureau of Standards*, 49(4):263–265, 1952. doi: 10.6028/jres.049.027.
- [21] Gérard Huet. Confluent reductions: Abstract properties and applications to term rewriting systems. *Journal of the ACM*, 27(4):797–821, 1980. doi: 10.1145/322217.322230.
- [22] W. Cary Huffman and Vera Pless. *Fundamentals of Error-Correcting Codes*. Cambridge University Press, 2003. doi: 10.1017/CBO9780511807077.
- [23] Gilles Kahn. The semantics of a simple language for parallel programming. In Jack L. Rosenfeld, editor, *Information Processing 74*, pages 471–475, Amsterdam, 1974. North-Holland.
- [24] Thomas Kailath. *Linear Systems*. Prentice-Hall, 1980.
- [25] Costas Kravaris. On functional observability of nonlinear systems and the design of functional observers with assignable error dynamics. *IFAC-PapersOnLine*, 59(19):668–673, 2025. doi: 10.1016/j.ifacol.2025.11.112.
- [26] Alexander Y. Kruger. Error bounds and metric subregularity. *Optimization*, 64(1):49–79, 2015. doi: 10.1080/02331934.2014.938074.
- [27] Salvador Lucas. Confluence of conditional rewriting modulo. *arXiv preprint arXiv:2504.01847*, 2025. doi: 10.48550/arXiv.2504.01847.
- [28] Sibe Mardešić. On approximate inverse systems and resolutions. *Fundamenta Mathematicae*, 142(3):241–255, 1993. doi: 10.4064/fm-142-3-241-255.
- [29] Sibe Mardešić and Leonard R. Rubin. Approximate inverse systems of compacta and covering dimension. *Pacific Journal of Mathematics*, 138(1):129–144, 1989. doi: 10.2140/pjm.1989.138.129.
- [30] Sibe Mardešić and Jack Segal. Stability of almost commutative inverse systems of compacta. *Topology and its Applications*, 31(3):285–299, 1989. doi: 10.1016/0166-8641(89)90025-4.
- [31] Arthur N. Montanari, Leandro Freitas, Daniele Proverbio, and Jorge Gonçalves. Functional observability and subspace reconstruction in nonlinear systems. *Physical Review Research*, 4(4):043195, 2022. doi: 10.1103/PhysRevResearch.4.043195.
- [32] M. H. A. Newman. On theories with a combinatorial definition of equivalence. *Annals of Mathematics*, 43(2):223–243, 1942. doi: 10.2307/1968867.
- [33] Christos H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

- [34] Arthur Ramos, Anjolina Oliveira, Ruy de Queiroz, and Tiago de Veras. A modular lean 4 framework for confluence and strong normalization of lambda calculi with products and sums, 2025.
- [35] Larry J. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3(1): 1–22, 1976. doi: 10.1016/0304-3975(76)90061-X.
- [36] Terese. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2003. doi: 10.1017/CBO9780511541792.