Irreversible Compression Commitment: A Typed Axiomatic Framework for Evolution, Extraction, and Commitment

Claudio Irrgang

May 3, 2026

Abstract

This paper introduces Irreversible Compression Commitment (ICC), a typed axiomatic framework for sequential irreversible commitment processes. Each commitment step is decomposed into a three-stage pipeline with explicit domain and codomain types: deterministic flow on density matrices, deterministic many-to-one extraction of outcome probabilities, and non-deterministic sampling to a single committed outcome. Three axioms constrain the pipeline (finiteness, boundary compression, and commitment decomposition), with compression stated at the finite boundary level of distinguishable inputs and outputs rather than on the continuous density-matrix interior.

Under these axioms, the main theorem is structural. It establishes strict boundary compression, non-recoverability of prior inputs from committed outcomes, and non-deterministic irreversibility; a corollary captures the discreteness, irreversibility, and sequentiality of the chain index. The primary nontrivial content of the framework, however, is not this single-step structural theorem but the chain accumulation theorem and the typed collision-geometry proposition: the exponential growth of compatible histories under iterated compression and the typed constraint on where collisions occur. These results go beyond the pigeonhole-derived non-recoverability by quantifying and structuring the ambiguity induced by the pipeline. The contribution of the framework is not the discovery of a new dynamical law, but rather the formal isolation of a compact typed package sufficient for these consequences and the explicit tracking of which axioms support which results. Three targeted weakenings produce three distinct loss patterns, demonstrating that the structural ingredients are independent.

The framework does not resolve the physical mechanism of single-outcome selection, and it does not establish that any physical system instantiates ICC. Rather, it provides an abstract target for comparing candidate realizations, weakened variants, and neighboring formal frameworks.

Contents

1 Introduction **3**

2 Formal Framework **4**

2.1 The generator 4
2.2 The typed pipeline 5
2.3 Axioms 6
2.4 The commitment chain 9
2.5 Component summary 9
2.6 Quantum realization 9

2.7 Classical realizations . . . 10
2.7.1 Coarse-grained Markov chain . . . 10
2.7.2 Analog-to-digital converter . . . 10

3 Structural Consequences: Compression, Accumulation, and Collision Geometry 11
3.1 Dependency localization . . . 19
3.2 Scope, dependencies, and interpretation . . . 22

4 Open Conjecture 23

5 Scope and Limitations 26
5.1 Strength of axioms and structural scope . . . 26
5.2 The measurement problem . . . 26
5.3 Basis and pointer-observable assumptions . . . 26
5.4 Markov and multi-channel extensions . . . 27
5.5 Categorical placement . . . 27
5.6 Comparison to Bennett–Landauer . . . 27
5.7 Regime classification and verification scope . . . 28

6 Conclusion 29

A Propositional Dependency Check 30

B Load-Bearing Dependency Map 31

C Worked Example: Two-State System 33

D Constructive Effective Realization in a Spin-Boson Measurement Model 33
D.1 Model setup . . . 34
D.2 Single-step checklist verification . . . 35
D.3 Basis-rigidity resolution . . . 35
D.4 Peakedness . . . 36
D.5 Chain extension . . . 36
D.6 What is not claimed . . . 37

1 Introduction

This paper studies a recurrent formal pattern: a system receives an environment input, evolves under a dynamical law, and yields a committed outcome from which the input cannot in general be recovered. This pattern (evolution followed by extraction followed by irreversible commitment) has analogues in quantum mechanics, thermodynamics, and information theory. The ingredients are individually familiar. The question addressed here is what becomes formally visible when they are packaged together as explicit axioms with typed signatures.

This paper introduces Irreversible Compression Commitment (ICC), a typed axiomatic framework for sequential compression-commitment processes. The framework decomposes each commitment step into three stages with explicit domain and codomain types: a deterministic flow map on density matrices, a deterministic many-to-one extraction of outcome probabilities, and a non-deterministic sampling step that produces a single committed outcome. Three axioms constrain the pipeline: finiteness, boundary compression, and commitment decomposition. Compression claims are placed at the finite boundary level (distinguishable inputs in, distinguishable outcomes out), not on the continuous density-matrix interior. Under these axioms, the main theorem establishes strict boundary compression, non-recoverability of prior inputs from committed outcomes, and non-deterministic irreversibility as structural consequences of the typed pipeline; a corollary captures the discreteness, irreversibility, and sequentiality of the chain index.

The contribution of this work is not the discovery of a new dynamical law, nor the derivation of irreversibility from weak premises. We use "irreversibility" strictly in the sense of non-invertibility of committed histories; the framework makes no claim about thermodynamic arrows or fundamental time asymmetry. ICC's value is in factoring this notion into separate, traceable structural ingredients, not in deriving it from weaker assumptions.

The methodological core is *dependency localization*: the explicit tracking of which axiom subparts are load-bearing for which structural consequences. The typed decomposition of the pipeline is the vehicle that makes this localization precise. ICC isolates irreversible commitment into three separable contributors (boundary compression, deterministic non-invertibility, and non-deterministic single-outcome sampling). Boundary compression alone is sufficient for non-recoverability via pigeonhole; the remaining axioms localize where in the pipeline information loss occurs.

The single-step theorem is therefore structural: once these load-bearing assumptions are fixed, its consequences follow directly. The nontrivial content lies not in any single conclusion but in the demonstration that these ingredients localize to different axiom subparts, with different consequences lost under different weakenings. ICC functions as a diagnostic for which structural ingredients drive irreversibility, a comparison lens for related formal frameworks, and a classification scheme for candidate physical realizations. Whether any particular physical system satisfies the axioms is a separate question, treated here as an open conjecture rather than an established result.

The framework does not solve the measurement problem; it isolates the stage at which the unresolved step enters. The non-deterministic sampling stage is explicitly left unspecified: its mechanism is not given, and the formal results do not depend on its resolution. Interpretive discussion is separated from the formal results and labeled throughout. This includes any correspondence between the chain index and physical time, and any correspondence between structural non-recoverability and thermodynamic irreversibility. The framework applies to systems admitting a fixed pointer observable [16] across all steps; systems with basis drift across environments fall outside the present formulation and require the per-step relaxation discussed in Section 5. The paper does not claim empirical confirmation, cosmological applicability, or priority over existing treatments of irreversibility.

The ingredients of this framework have antecedents in adjacent literature. Decoherence theory explains the emergence of effectively diagonal density matrices but does not by itself specify a

single-outcome selection mechanism [15, 16]. Work on retrodiction studies conditions under which past states can be inferred given additional inferential structure [14]. Categorical treatments of quantum channels provide formal language for composing completely positive maps [9]. The theory of quantum instruments [2] formalizes measurement with post-measurement states; ICC's Activation–Sample decomposition can be viewed as a typed refinement of an instrument into a deterministic extraction stage and a non-deterministic commitment stage. Consistent histories [4, 6, 13] studies sequences of projector-defined events on closed quantum systems, with consistency conditions selecting which families of histories admit classical probability assignment. ICC's commitment chain may be viewed as fixing one such family, with the typed pipeline making the projection structure explicit at each step. The frameworks differ in that ICC axiomatizes single-outcome commitment as a primitive rather than handling outcome selection within the calculus of histories itself. Axiomatic information thermodynamics derives thermodynamic-style constraints from information-theoretic primitives [8]; the compression conditions there are entropy-based, whereas ICC's boundary compression is a finite-cardinality inequality on distinguishable inputs and outputs. ICC occupies a narrower structural niche: it axiomatizes a typed pipeline from evolution through probability extraction to single-outcome commitment, localizes the unresolved step, and derives which irreversible consequences depend on which structural ingredient. The contrasts with these literatures are narrow and specific.

ICC does not assert that physical dynamics are fundamentally compressive; rather, it identifies compression of distinguishable configurations as the structural condition under which an operational arrow emerges.

The primary nontrivial content of the framework is not the single-step structural theorem (Theorem 3.1) but the chain accumulation theorem (Theorem 3.3) and the typed collision-geometry proposition (Proposition 3.11): the exponential growth of compatible histories under iterated compression and the typed constraint on where collisions occur in the $\Delta_{k-1} \times E_k$ factorization. These results go beyond the pigeonhole-derived non-recoverability by quantifying and structuring the ambiguity induced by the pipeline; the single-step theorem is the entry point, not the endpoint, of the structural account.

The remainder of the paper is organized as follows. Section 2 defines the typed pipeline, states the three axioms, and establishes notation. Section 3 states and proves the main structural theorem and its corollaries. Section 4 formulates the open conjectures connecting the formal framework to quantum decoherence. Section 5 discusses limitations, scope boundaries, and open questions. Section 6 summarizes the results.

2 Formal Framework

We define a sufficient axiomatic package for irreversible commitment structure. Picture each step as a state of affairs that gets nudged by an environmental input, evolves briefly, and snaps to one committed outcome that becomes the input to the next step. The framework applies to a single system undergoing a sequential chain of such compression-commitment events: at each step k, the system occupies a committed state $A_{k-1} \in \Delta_{k-1}$ (a finite set), receives an environment input $e_k \in E_k$ (a finite set), and one application of the commitment operator M produces the next committed state A_k. Formal definitions and axioms are stated first; interpretive remarks are labeled and separated.

2.1 The generator

Definition 2.1 (Generator)**.** For each environment input $e_k \in E_k$, the generator $C(e_k)$ is a vector field on the space of $N \times N$ density matrices $\mathrm{DM}(N)$, where N is a positive integer satisfying $N \geq |\Delta_k|$ for the given step k (or, for a finite chain of length n, a single $N \geq \max_k |\Delta_k|$ suffices). Given $\rho \in \mathrm{DM}(N)$, it returns $d\rho/dt = C(e_k)[\rho]$. A generator is *admissible* if the initial

value problem $d\rho/dt = C(e_k)[\rho]$ has a unique solution on $[0, \Delta t]$ for every $\rho_0 \in \mathrm{DM}(N)$. All generators in this framework are assumed admissible.

The generator defines the equation of motion but does not solve it. C is not a stage in the pipeline; it parameterizes the flow map Φ defined below. In the quantum setting, the Lindblad equation $d\rho/dt = -i[H, \rho] + \mathcal{D}[\rho]$ [5, 12] is an instance of C.

2.2 The typed pipeline

The commitment process σ decomposes into three stages with explicit types.

Definition 2.2 (Flow map)**.** $\Phi^{\Delta t}_{C(e_k)}\colon \mathrm{DM}(N) \to \mathrm{DM}(N)$. Given an initial density matrix ρ_0 and an admissible generator $C(e_k)$, the flow map $\Phi^{\Delta t}_{C(e_k)}(\rho_0)$ is the unique solution at time Δt to the initial value problem

$$\frac{d\rho}{dt} = C(e_k)[\rho], \qquad \rho(0) = \rho_0.$$

Equivalently, $\rho(\Delta t)$ satisfies the integral equation $\rho(\Delta t) = \rho_0 + \int_0^{\Delta t} C(e_k)[\rho(t')]\, dt'$. Φ is deterministic: the same ρ_0 and C produce the same $\rho(\Delta t)$.

Remark 2.3. Admissibility (Definition 2.1) is a regularity hypothesis on C, not a derived consequence of the axioms. In the Lindblad case the generator is bounded and linear on a finite-dimensional space, so admissibility is automatic. For the abstract framework, admissibility restricts the class of generators to those for which the flow map Φ is well-defined.

Remark 2.4 (Step duration). Throughout this paper, $\Delta t > 0$ denotes a single fixed step duration, common to every step k of the chain. That is, each application of the flow map $\Phi^{\Delta t}_{C(e_k)}$ integrates the generator over the same interval $[0, \Delta t]$. This convention simplifies notation and is sufficient for all results in Section 3: the theorems hold at the level of the pipeline's boundary cardinalities and injectivity properties, neither of which depends on the value of Δt or its constancy across steps. The extension to step-dependent durations $\Delta t_k > 0$ is straightforward: each step's flow map $\Phi^{\Delta t_k}_{C(e_k)}$ remains well-defined provided $C(e_k)$ is admissible (Definition 2.1), and all per-step structural results of Theorem 3.1 carry over unchanged.

Remark 2.5 (Nominal state space). The space $\mathrm{DM}(N)$ is the framework's nominal state space: it carries the type signatures of C, Φ, and Activation, and it anchors the quantum realization of Section 2.6. The structural results of Section 3 depend on cardinalities and typing rather than on the full quantum state-space structure, so classical realizations (such as the coarse-grained Markov chain and the analog-to-digital converter of Section 2.6) instantiate the framework by embedding their state spaces into $\mathrm{DM}(N)$ as diagonal density matrices over a chosen basis. The choice of $\mathrm{DM}(N)$ as the nominal carrier is therefore a structural commitment with downstream scope consequences (Section 5), not a hidden universality claim.

Definition 2.6 (Activation)**.** Activation$\colon \mathrm{DM}(N) \to (P_k, \mu_k)$ is a deterministic, many-to-one map (not necessarily basis-diagonal) from density matrices to probability distributions over a finite outcome set P_k. The framework's structural results depend only on this abstract typing, not on the specific extraction rule. The formal constraints on Activation are asserted in Axiom ICC-3(b). A concrete realization, used in the quantum interpretation throughout this paper, is diagonal extraction in a fixed orthonormal basis $\{|i\rangle\}_{i \in P_k}$: $\mu_k(i) = \langle i|\rho(\Delta t)|i\rangle$ for each $i \in P_k$. The notation (P_k, μ_k) denotes a distribution μ_k on the finite set P_k and is used throughout. Other deterministic many-to-one maps satisfying Axiom ICC-3(b) yield variant realizations.

Remark 2.7. The bra-ket notation $\mu_k(i) = \langle i|\rho(\Delta t)|i\rangle$ presupposes a choice of orthonormal basis $\{|i\rangle\}_{i \in P_k}$ in which the diagonal is extracted. This basis dependence is a parameter of the framework, not derived from it. In the quantum setting, the basis may be selected

by environment-induced decoherence (Section 2.6); in the abstract framework, it is fixed by assumption.

Definition 2.8 (Sample)**.** Sample: $(P_k, \mu_k) \to P_k$. Sample takes a probability distribution over a finite outcome set and returns a single committed outcome $A_k \in P_k$. The formal constraints on Sample are asserted in Axiom ICC-3(c).

Remark 2.9. The framework does not specify how a single outcome is selected from μ_k. In a quantum setting, this is the measurement problem: how one definite outcome arises from a distribution over alternatives. The results of Section 3 do not depend on any particular resolution of this issue.

Definition 2.10 (Commitment process)**.** The commitment process is the composition

$$\sigma_{C(e_k)} \coloneqq \text{Sample} \circ \text{Activation} \circ \Phi^{\Delta t}_{C(e_k)}.$$

The pipeline composes as:

$$\text{DM}(N) \xrightarrow{\Phi} \text{DM}(N) \xrightarrow{\text{Activation}} (P_k, \mu_k) \xrightarrow{\text{Sample}} P_k.$$

Remark 2.11. The decomposition of σ into a deterministic extraction stage and a non-deterministic commitment stage may be viewed as a typed unpacking of the quantum instrument formalism [2], which packages measurement and post-measurement state into a single map. The present framework makes these internal stages and their types explicit. The chain construction (Section 2.4) extends this typing to a sequence of commitment steps, paralleling the history-sequence structure of consistent histories [4, 6].

Definition 2.12 (Commitment operator)**.** The commitment operator M takes a committed state $A_{k-1} \in \Delta_{k-1}$ and an environment input e_k, and returns a committed state A_k:

$$M(A_{k-1}, e_k) = \sigma_{C(e_k)}(\rho_0), \quad \rho_0 = |A_{k-1}\rangle\langle A_{k-1}|.$$

The notation $M = \sigma \circ C$ is shorthand for "σ applied using C as generator." C parameterizes σ; it is not a stage in the pipeline.

Remark 2.13 (Pure-state re-initialization). The convention $\rho_0 = |A_{k-1}\rangle\langle A_{k-1}|$ in Definition 2.12 treats each committed state as re-preparing a pure projector on the fixed basis. This aligns the post-Sample state with a von Neumann projection postulate and is an idealization: realistic open-system commitment need not produce a pure state at the start of the next step. The framework adopts the pure-state form for cleanliness. Relaxing to a state-preparation map $S\colon P_k \to \text{DM}(N)$ at each step is straightforward and preserves the per-step structural results of Theorem 3.1, since the proofs depend only on the typed signatures and non-invertibility properties, not on the rank of ρ_0. The pure-state reinitialization additionally enforces a Markov structure on the chain: each step starts from a pure projector on the previous committed outcome, with no cross-step correlations retained. Irreversibility claims in this framework are therefore local to the commitment boundary at each step, not accumulated through retained state. A relaxation to non-Markov chains is a separate extension, distinct from the rank relaxation noted above.

2.3 Axioms

Axiom ICC-1 (Finite)**.** $|\Delta_k| < \infty$ *and* $|E_k| < \infty$ *for all* k.

Definition 2.14 (Distinguishable elements)**.** Let $k \geq 1$ be a chain step.

(a) **Committed states and outcomes.** Two elements of Δ_{k-1} (or of P_k) are *distinguishable* if and only if they carry different labels under the fixed orthonormal basis $\{|i\rangle\}_{i=1}^{N}$ of Definition 2.23(b). Because $\Delta_{k-1} \cong P_{k-1}$ via the canonical identification of Definition 2.23(c), the same label-equality relation governs both sets. $|\Delta_{k-1}|$ and $|P_k|$ therefore count distinct basis labels.

(b) **Environment inputs.** Two elements $e, e' \in E_k$ are *distinguishable* if and only if they parameterize different admissible generators: $C(e) \neq C(e')$ as vector fields on $\mathrm{DM}(N)$ (Definition 2.1). $|E_k|$ counts the resulting equivalence classes.

The cardinalities $|\Delta_{k-1}|$, $|P_k|$, and $|E_k|$ in Axiom ICC-2(a) are understood throughout in this sense.

Remark 2.15. Part (a) makes the compression claim basis-relative in the same way that Activation is (Remark 2.7). If the basis is refined, some previously collapsed outcomes split into distinct labels; the inequality $|P_k| < |\Delta_{k-1}| \cdot |E_k|$ is then stated in terms of the refined counts. Different choices of distinguishability structure may also alter whether the inequality holds at all: a sufficiently coarse partition could collapse outcomes so that the inequality fails, and a sufficiently fine one could violate finiteness (Axiom ICC-1). The basis is therefore not a notational convenience but a parameter on which the framework's applicability conditions depend. Part (b) identifies environment inputs up to their effect on the generator; two inputs that induce the same dynamics are treated as a single input for the purpose of the compression count. Both choices are parameters of the framework, not derived from it.

Axiom ICC-2 (Compression)**.** *At each step:*

(a) ***Boundary compression.*** *The number of distinguishable committed outcomes is strictly less than the number of distinguishable input configurations (distinguishability as in Definition 2.14):*

$$|P_k| < |\Delta_{k-1}| \cdot |E_k|.$$

(b) ***Pipeline non-invertibility.*** *The deterministic pipeline* $\mathrm{Activation} \circ \Phi_{C(e_k)}^{\Delta t}$ *from input states to outcome distributions is non-invertible: distinct inputs (in particular, distinct boundary inputs* $|A_{k-1}\rangle\langle A_{k-1}|$*) may produce identical outcome distributions.*

Remark 2.16. The density matrix $\rho \in \mathrm{DM}(N)$ is a continuous intermediate representation. The compression claim in Axiom ICC-2 is about the finite boundaries (distinguishable inputs in, distinguishable outcomes out; see Definition 2.14), not about the continuous interior of the pipeline.

Remark 2.17. Axiom ICC-2(a) does not exclude the degenerate cases $|P_k| = 1$ (a single committed outcome regardless of input) or $|E_k| = 1$ (trivial environment). The formal results hold in these regimes but are uninteresting: a single-outcome system trivially satisfies all compression and non-recoverability claims, and a trivial environment reduces the compression inequality to $|P_k| < |\Delta_{k-1}|$. Similarly, $|\Delta_{k-1}| = 1$ (a single prior state) permits $|P_k| = 1$ without violating the axiom in an interesting way. The intended scope of the framework is the non-degenerate regime $|\Delta_{k-1}| \geq 2$, $|P_k| \geq 2$, and $|E_k| \geq 2$, where compression is substantive.

Remark 2.18. Axiom ICC-2 is the primary compression axiom. Non-recoverability is a consequence (Theorem 3.1(ii)), not the axiom's direct content. The framework does not derive non-recoverability from weak premises; it isolates a sufficient axiomatic package and identifies the structural consequences.

Remark 2.19 (Semantic independence of ICC-2(b)). Axiom ICC-2(b) is not implied by Axiom ICC-2(a) together with Axiom ICC-3(b). We exhibit an explicit counterexample.

Set $|\Delta_0| = 2$, $|E_1| = 2$, $|P_1| = 2$, so Axiom ICC-2(a) holds: $|P_1| = 2 < 4 = |\Delta_0|\,|E_1|$. Take $N = 2$ (a single qubit). Define two unitary generators via the Hamiltonians

$$H(e_0) \coloneqq \tfrac{\pi}{6}\,\sigma_y, \qquad H(e_1) \coloneqq \tfrac{\pi}{5}\,\sigma_y,$$

with no dissipator, so $\Phi^1_{C(e_j)} = \exp(-\mathrm{i}\,H(e_j))(\,\cdot\,)\exp(\mathrm{i}\,H(e_j))$. Label the two boundary inputs $\rho_0 = |0\rangle\langle 0|$ and $\rho_1 = |1\rangle\langle 1|$, and let Activation extract the diagonal of the evolved state as a distribution on $P_1 = \{0, 1\}$. The $(0, 0)$ entry of each evolved state determines the full distribution; the four values are

$$\cos^2\tfrac{\pi}{6} = \tfrac{3}{4}, \quad \sin^2\tfrac{\pi}{6} = \tfrac{1}{4}, \quad \cos^2\tfrac{\pi}{5} = \tfrac{3+\sqrt{5}}{8}, \quad \sin^2\tfrac{\pi}{5} = \tfrac{5-\sqrt{5}}{8}.$$

These four values are pairwise distinct (all six pairwise differences are non-zero, verified independently by exact symbolic computation and by a machine-checked Lean 4 / `mathlib4` formalization), so Activation $\circ\,\Phi$ is injective on these four boundary inputs and Axiom ICC-2(b) fails. Activation remains globally many-to-one on $\mathrm{DM}(N)$ (as required by Axiom ICC-3(b)): the injectivity holds only on these four boundary inputs, not as a property of Activation on the full space of density matrices. The distinction is that Axiom ICC-2(a) constrains only finite boundary cardinalities, whereas Axiom ICC-2(b) requires injectivity into a continuous simplex of distributions. The downstream consequence is Proposition 3.22: dropping Axiom ICC-2(b) removes one recovery obstacle and weakens Theorem 3.1(iii).

Axiom ICC-3 (Commitment)**.** *The commitment process σ satisfies the following constraints.*

(a) ***Decomposition.*** $\sigma = \mathrm{Sample} \circ \mathrm{Activation} \circ \Phi^{\Delta t}_{C(e_k)}$.

(b) ***Activation is deterministic and many-to-one.*** Activation: $\mathrm{DM}(N) \to (P_k, \mu_k)$ *is deterministic: the same $\rho(\Delta t)$ always produces the same (P_k, μ_k). It is many-to-one: distinct density matrices may produce the same μ_k, with the information distinguishing them not preserved in the output. In the diagonal-extraction realization of Definition 2.6, the discarded information consists of the off-diagonal entries of $\rho(\Delta t)$.*

(c) ***Sample is non-deterministic.*** Sample: $(P_k, \mu_k) \to P_k$ *returns a committed outcome $A_k \in P_k$ drawn according to μ_k. The individual outcome is non-deterministic: the same μ_k may yield different A_k on different applications. (When μ_k is a point mass, Sample is trivially deterministic for that step; Theorem 3.1(iii) applies non-vacuously only when μ_k assigns positive probability to at least two outcomes.)*

Remark 2.20 (Structural vs representational loss). Boundary compression (Axiom ICC-2(a)) is representation-independent: it constrains finite-cardinality counts without reference to a basis or to the specific Activation map. Activation-induced loss (Axiom ICC-3(b)) is representation-dependent: which density matrices collapse to the same μ_k depends on the chosen extraction rule. The two contributions are therefore distinct in kind, not just in stage along the pipeline. Activation need not be lossy on the boundary in principle; loss arises only for specific extraction choices satisfying Axiom ICC-3(b). An informationally complete extraction would not satisfy Axiom ICC-3(b) and would lie outside the framework's scope.

Remark 2.21 (Probabilistic interpretation). Axiom ICC-3(c) asserts that an individual Sample outcome is non-deterministic and drawn according to μ_k. The operational gloss (that statistics over many independent applications to the same μ_k converge to μ_k) is a probabilistic interpretation of the axiom over an ensemble of trials. The chain itself is a single sequence; the convergence claim is properly about a meta-level ensemble of such sequences and does not enter the formal results of Section 3.

Remark 2.22. The generator C is not axiomatized. C is the equation that Φ integrates. Axiom ICC-2 constrains the pipeline (what happens when C is solved). Different generators (different physics) may produce the same pipeline constraint.

2.4 The commitment chain

The chain is the sequential application of M:

$$A_1 = M(A_0, e_1),$$
$$A_2 = M(A_1, e_2),$$
$$A_k = M(A_{k-1}, e_k).$$

Each committed output becomes the state input to the next step. The environment provides fresh input at each step. The identification of committed state spaces with outcome sets is stated formally in Definition 2.23(c).

Definition 2.23 (Initialization and basis consistency)**.** The chain requires:

(a) An initial committed state $A_0 \in \Delta_0$, where Δ_0 indexes a subset of an orthonormal basis of the N-dimensional space, so that $\rho_0 = |A_0\rangle\langle A_0| \in \mathrm{DM}(N)$.

(b) A single fixed orthonormal basis $\{|i\rangle\}_{i=1}^{N}$ used throughout the chain: the outcome labels in P_k at each step correspond to the basis elements used by Activation (Remark 2.7).

(c) At each step, the committed state space Δ_k is identified with the outcome set P_k via the canonical identification $P_k \cong \Delta_k$. This identification carries labels forward: $A_k \in P_k$ becomes $A_k \in \Delta_k$, and $\rho_0 = |A_k\rangle\langle A_k|$ is formed using the fixed basis of (b) at step $k{+}1$.

The axioms do not explain A_0; they specify the structure of each step given a chain starting at A_0.

Remark 2.24. The system's distinguishable state count $|\Delta_k|$ may grow, shrink, or remain the same at any step, depending on the environment input e_k. What is always true: the combined input $|\Delta_{k-1}| \cdot |E_k|$ is strictly compressed to $|P_k|$ (Axiom ICC-2).

2.5 Component summary

The following table summarizes the typed components.

Component	Type	Deterministic	Status
$C(e_k)$	vector field on $\mathrm{DM}(N)$	yes	known
$\Phi_C^{\Delta t}$	$\mathrm{DM}(N) \to \mathrm{DM}(N)$	yes	known (admissible C)
Activation	$\mathrm{DM}(N) \to (P_k, \mu_k)$	yes	known
Sample	$(P_k, \mu_k) \to P_k$	no	open

Three components have known instances. Sample does not: its mechanism is the measurement problem. The axioms constrain M without requiring the Sample mechanism to be identified.

2.6 Quantum realization

In the quantum setting, the components have standard identifications:

- $C(e_k)[\rho] = -i[H(e_k), \rho] + \mathcal{D}(e_k)[\rho]$ (Lindblad generator).
- $\Phi_{C(e_k)}^{\Delta t}$ is the corresponding CPTP map.
- Activation extracts the diagonal of $\rho(\Delta t)$ in a chosen or decoherence-preferred basis (Born rule).

- Sample is the unresolved step: the mechanism by which one outcome is selected from μ_k.

In the CPTP reading, Φ may already encode a first information-loss layer: unitary evolution of a larger system–environment state followed by tracing out the environment. Activation then supplies a second, representation-dependent layer by extracting outcome probabilities, for example through diagonal extraction in a chosen basis. This two-layer reading belongs to the quantum interpretation of the framework, not to the abstract ICC axioms themselves.

The decoherence-preferred basis cited above is environment-dependent; the framework's single-fixed-basis requirement of Definition 2.23(b) is therefore in tension with the candidate selection mechanism. Section 5 discusses this tension and the per-step relaxation.

These identifications are not required by the axioms. The axioms specify structural constraints on the pipeline; the quantum setting is one candidate realization.

2.7 Classical realizations

The framework admits classical instances that embed into DM(N) via diagonal density matrices over the basis of distinguishable states (Remark 2.5). Two are given here.

2.7.1 Coarse-grained Markov chain

Take a finite state space of microstates partitioned into coarse classes, with a per-step transition rule selected by the environment input. In the coarse-grained Markov chain classical instance, the components have standard identifications:

- $C(e_k)$: the parameter selecting the per-step transition kernel for environment input e_k.
- $\Phi^{\Delta t}_{C(e_k)}$: applies the selected transition kernel, evolving the diagonal density matrix over microstates.
- Activation: the deterministic many-to-one aggregation from microstate distribution to a distribution over coarse classes.
- Sample: commits one coarse class per step.

The structural results of Theorem 3.1 apply unchanged in this representation; it instantiates the framework without invoking decoherence or a quantum-mechanical substrate.

2.7.2 Analog-to-digital converter

Take $K = 1024$ distinguishable input voltage levels (the boundary Δ_{k-1}), $M = 4$ noise environments (the parameters E_k, indexing distinct generators $C(e_k)$ as additive Gaussian noise of differing amplitude), and an 8-bit quantizer producing $P = 256$ output bins. In the analog-to-digital converter classical instance, the components have standard identifications:

- $C(e_k)$: additive Gaussian noise of environment-selected amplitude.
- $\Phi^{\Delta t}_{C(e_k)}$: adds the environment-selected noise to the input voltage.
- Activation: the quantizer $Q(\cdot)$, deterministic and many-to-one.
- Sample: the comparator latch, committing one of 256 bins per cycle.

Axiom ICC-2(a) holds with a factor-of-16 gap ($256 < 4096$). Axiom ICC-2(b) holds non-trivially at every noise scale: at zero noise, exactly four distinct sub-bin voltages collapse onto each output bin, so the deterministic pipeline is many-to-one with multiplicity four; at positive noise, voltage intervals produce overlapping μ_k distributions, extending the non-invertibility from discrete

sub-bin levels to continuous regions. The non-determinism of Axiom ICC-3(c) is the axiom that requires positive noise scale to be non-vacuous: at zero noise, μ_k is a point mass and Sample is trivially deterministic; at positive noise, near-threshold inputs spread μ_k across adjacent bins and Sample is genuinely non-deterministic.

3 Structural Consequences: Compression, Accumulation, and Collision Geometry

This section states the structural consequences of Axioms ICC-1–ICC-3 and the typed recursive construction of Section 2.4. Theorem 3.1 establishes the minimal structural consequence of boundary compression (non-injectivity and non-recoverability via pigeonhole). The stronger content of the section appears in Theorem 3.3 and Proposition 3.11, which respectively quantify the resulting ambiguity (chain accumulation $\geq \prod_k \lceil 1/\alpha_k \rceil$) and constrain its typed structure (collisions are forced to be cross-axis under per-axis injectivity). Each result is proved from explicitly cited premises.

Theorem 3.1 (Main Structural Theorem)**.** *Under Axioms ICC-1–ICC-3, admissibility of C (Definition 2.1), the initialization of Definition 2.23, and the commitment chain of Section 2.4 (all relative to the fixed basis of Definition 2.23(b)), the following hold at each step k:*

(i) ***Strict boundary compression.*** *$|P_k| < |\Delta_{k-1}| \cdot |E_k|$. The map from distinguishable input configurations (Definition 2.14) to committed outcomes is therefore non-injective.*

(ii) ***Non-recoverable antecedent.*** *The input pair (A_{k-1}, e_k) is not recoverable from the committed outcome A_k.*

(iii) ***Non-deterministic irreversibility.*** *Applying M to the same input (A_{k-1}, e_k) twice may produce different committed outcomes. Neither application can be inverted.*

Proof. We prove each part separately.

Part (i): Strict boundary compression.

The inequality $|P_k| < |\Delta_{k-1}| \cdot |E_k|$ is direct from Axiom ICC-2(a). Both sides are finite by Axiom ICC-1. By Axiom ICC-3(c), each input pair (A_{k-1}, e_k) has a non-empty set of possible committed outcomes (the support of μ_k in P_k). If the supports for distinct input pairs were pairwise disjoint, the union of supports would contain at least $|\Delta_{k-1}| \cdot |E_k|$ distinct elements, contradicting $|P_k| < |\Delta_{k-1}| \cdot |E_k|$. Hence at least two distinct input pairs are compatible with the same committed outcome. This holds independently of which outcome Sample selects on any particular application. A machine-checked formalization of this part is available in Lean 4 against `mathlib4` (see Appendix A).

Part (ii): Non-recoverable antecedent.

Recovery of (A_{k-1}, e_k) from A_k requires inverting the pipeline in reverse. This inversion is blocked by three independent obstructions, any one of which suffices to prevent recovery:

(a) The deterministic pipeline Activation $\circ\, \Phi^{\Delta t}_{C(e_k)}$ is non-invertible by Axiom ICC-2(b): distinct inputs may produce the same outcome distribution.

(b) Activation is many-to-one by Axiom ICC-3(b): distinct evolved density matrices may produce the same μ_k, so the information distinguishing them is destroyed at this stage. In the diagonal-extraction realization, this corresponds to off-diagonal entries being discarded.

(c) Sample is non-deterministic by Axiom ICC-3(c): from the committed outcome A_k alone, one cannot reconstruct the full distribution μ_k from which it was drawn.

Part (ii) is best read as a layered analysis of recovery failure across the reverse pipeline, not as three logically independent barriers. Obstacle (c) blocks reconstruction of μ_k from A_k; this is the first step any recovery attempt must solve. Obstacle (b) blocks recovery of $\rho(\Delta t)$ from μ_k at the Activation-inversion stage. Obstacle (a) blocks recovery of the input from the full deterministic pipeline, which includes both Activation and Φ; it therefore subsumes the scope of (b) but is attributed to a different and independently load-bearing axiom (see Remark 2.19). Full recovery requires solving all three stages in reverse; failure at any stage suffices to prevent it. (The extension from single-step non-recoverability to the full chain is the content of Corollary 3.19.)

Note that part (i) operates at a different structural level than the three pipeline obstacles above: Axiom ICC-2(a) ensures collisions exist at the input-to-output boundary by pigeonhole, while obstacles (a), (b), (c) identify *where* along the reverse path recovery fails. Boundary compression is the dominant mechanism at the global level; the pipeline obstacles localize the realization within the architecture. All non-recoverability results in this paper ultimately reduce to Axiom ICC-2(a); the remaining axioms do not generate non-recoverability but localize its realization within the pipeline. Axiom ICC-3(c) in particular is responsible for trajectory variability under identical inputs (Theorem 3.1(iii)) but is not the source of non-recoverability, which follows from boundary compression alone.

Part (iii): Non-deterministic irreversibility.

Given the same input (A_{k-1}, e_k), the generator $C(e_k)$ and initial state ρ_0 are identical. Φ is deterministic (Definition 2.2), so $\rho(\Delta t)$ is the same. Activation is deterministic (Axiom ICC-3(b)), so μ_k is the same. But Sample is non-deterministic (Axiom ICC-3(c)): the same μ_k may yield different A_k. Independently, the deterministic pipeline is non-invertible (Axiom ICC-2(b)): even if A_k were known, the input cannot be recovered. □

Remark 3.2 (Role of non-deterministic sampling beyond variability). While non-determinism is not required for non-recoverability (Theorem 3.1(ii) survives the removal of Axiom ICC-3(c) via the remaining obstacles), it is required for trajectory branching under identical inputs. In its absence, the chain collapses to a deterministic map on boundary states, and the multiplicity described in Theorem 3.3 reflects only static ambiguity (many input histories per output) without dynamic branching (multiple realizations per input). Axiom ICC-3(c) thus separates ambiguity of histories from variability of realizations: the former is a property of the boundary map, the latter is a property of how the same input is re-instantiated across runs.

Theorem 3.3 (Multiplicity bound and chain accumulation)**.** *Under Axioms ICC-1 and ICC-2(a):*

(i) Per-step bound. *At each step k, fix any deterministic realization of* Sample*, so that the committed-outcome map $f_k\colon \Delta_{k-1} \times E_k \to P_k$ is a function. Then*

$$\max_{p \in P_k} \left| f_k^{-1}(p) \right| \;\geq\; \left\lceil \frac{|\Delta_{k-1}| \cdot |E_k|}{|P_k|} \right\rceil, \tag{1}$$

and this bound is tight.

(ii) Chain accumulation. *Let $f_1, \ldots, f_n$ be fixed realizations of the committed-outcome maps at steps $1, \ldots, n$, and let $g_n := f_n \circ \cdots \circ f_1$ denote the composed chain map (using the identification $P_k \cong \Delta_k$ of Definition 2.23(c)). Then*

$$\max_{p_n \in P_n} \left| g_n^{-1}(p_n) \right| \;\geq\; \prod_{k=1}^{n} \left\lceil \frac{|\Delta_{k-1}| \cdot |E_k|}{|P_k|} \right\rceil. \tag{2}$$

In particular, if the per-step compression ratio $\alpha_k := |P_k|/(|\Delta_{k-1}| \cdot |E_k|)$ is bounded by $\alpha < 1$ uniformly, the right-hand side is at least $\lceil 1/\alpha \rceil^n$, growing without bound as $n \to \infty$.

Proof. **Part (i).** The $|\Delta_{k-1}| \cdot |E_k|$ inputs of f_k distribute among $|P_k|$ outcomes; by integer pigeonhole, at least one outcome has preimage size at least $\lceil |\Delta_{k-1}| \cdot |E_k| / |P_k| \rceil$. The bound is achieved by any function distributing inputs as evenly as possible across outcomes. This part is also machine-checked in Lean 4 (see Appendix A).

Part (ii). By induction on n. The base case is part (i). For the inductive step, step n's map f_n has domain $\Delta_{n-1} \times E_n \cong P_{n-1} \times E_n$ (by Definition 2.23(c)) and codomain P_n; applying part (i) to f_n yields a step-n multiplicity factor of $\lceil |\Delta_{n-1}| \cdot |E_n| / |P_n| \rceil = \lceil 1/\alpha_n \rceil$. Multiplying by the inductive bound through step $n-1$ gives Equation (2). □

Remark 3.4 (Scope and stability). The bounds Equations (1) and (2) are existential on the *maximum*: at least one (terminal) outcome has many consistent input histories. They do not assert that every outcome has a large preimage. Theorem 3.1(i) is recovered as the case where the bound is ≥ 2; Theorem 3.3 sharpens this to a quantitative rate, and complements Corollary 3.19: the corollary states that no recovery procedure can invert the chain; the present bound states how many distinct input histories remain consistent with each present committed state. The bound is robust under perturbation: small changes in the per-step ratios α_k do not affect $\prod_k \lceil 1/\alpha_k \rceil$ until they cross a rounding threshold, and the bound is monotone in compression ($\alpha'_k \leq \alpha_k$ for all k implies the corresponding chain bound is at least as large).

Remark 3.5 (Mean preimage size and average chain accumulation). Theorem 3.3 establishes existential lower bounds: at least one outcome carries a large preimage. Two complementary counting facts sharpen the picture without invoking any probability distribution over inputs. Throughout, $\alpha_k \coloneqq |P_k| / (|\Delta_{k-1}| \cdot |E_k|) \in (0, 1]$ denotes the per-step compression ratio of Theorem 3.3.

Per-step mean. For any deterministic realization $f_k \colon \Delta_{k-1} \times E_k \to P_k$, the mean preimage size averaged over the codomain P_k is exactly

$$\frac{1}{|P_k|} \sum_{p \in P_k} |f_k^{-1}(p)| \;=\; \frac{|\Delta_{k-1}| \cdot |E_k|}{|P_k|} \;=\; \frac{1}{\alpha_k}. \tag{3}$$

This is a counting identity: each input in $\Delta_{k-1} \times E_k$ appears in exactly one preimage, so the preimage sizes sum to $|\Delta_{k-1}| \cdot |E_k|$; dividing by $|P_k|$ gives the mean. The averaging is over the codomain; no measure on inputs is invoked.

Min/max bracket. By the same identity, the preimage sizes cannot all equal $1/\alpha_k$ when $1/\alpha_k \notin \mathbb{Z}$. There exists at least one outcome with preimage size $\leq \lfloor 1/\alpha_k \rfloor$ (otherwise the sum would exceed $|\Delta_{k-1}| \cdot |E_k|$), and by Equation (1) there exists at least one outcome with preimage size $\geq \lceil 1/\alpha_k \rceil$. When $1/\alpha_k \notin \mathbb{Z}$, these are different: the preimage-size distribution over P_k is necessarily inhomogeneous.

Average chain accumulation. The same identity applied to the composed chain map $g_n \coloneqq f_n \circ \cdots \circ f_1$ gives an exact average over P_n:

$$\frac{1}{|P_n|} \sum_{p_n \in P_n} |g_n^{-1}(p_n)| \;=\; \prod_{k=1}^{n} \frac{1}{\alpha_k}. \tag{4}$$

The domain of g_n has size $|\Delta_0| \cdot \prod_{k=1}^{n} |E_k|$; telescoping the identification $|\Delta_k| = |P_k|$ (Definition 2.23(c)) and the compression ratios $\alpha_k = |P_k| / (|\Delta_{k-1}| \cdot |E_k|)$ gives $|P_n| = (\prod_k \alpha_k) \cdot |\Delta_0| \cdot \prod_k |E_k|$; the mean follows by division. Comparing Equation (4) with Equation (2): the maximum bound grows at least as $\prod_k \lceil 1/\alpha_k \rceil$ (with ceiling at each step), while the average is exactly $\prod_k (1/\alpha_k)$ (no ceiling). The gap between them quantifies how unevenly consistent histories distribute across terminal outcomes; a large gap indicates the chain's collision multiplicity is concentrated in a small fraction of outcomes.

Corollary 3.6 (Monotonic ambiguity growth)**.** *Let A_n be the committed state at step n. The size of the consistent-history set $|g_n^{-1}(A_n)|$ is non-decreasing in n and strictly increases whenever $|P_k| < |\Delta_{k-1}| \cdot |E_k|$ at step k.*

Proof. Immediate from Theorem 3.3(ii): each step contributes a multiplicative factor of at least $\lceil |\Delta_{k-1}| \cdot |E_k| / |P_k| \rceil \geq 1$, with strict inequality whenever Axiom ICC-2(a) is non-trivial at that step. □

Proposition 3.7 (Monotone cumulative compression deficit)**.** *Define the* cumulative compression deficit *through step n as*

$$D_n \;:=\; \sum_{k=1}^{n} \log_2\!\left(\frac{|\Delta_{k-1}| \cdot |E_k|}{|P_k|}\right). \tag{5}$$

Under Axioms ICC-1 and ICC-2(a):

(i) *$D_n \geq 0$ for all n, with $D_0 = 0$.*

(ii) *D_n is monotone non-decreasing in n.*

(iii) *$D_n > D_{n-1}$ strictly whenever step n satisfies $|P_n| < |\Delta_{n-1}| \cdot |E_n|$.*

Proof. Axiom ICC-2(a) ensures $|\Delta_{k-1}| \cdot |E_k| \geq |P_k|$ at every step, so each summand $\log_2((|\Delta_{k-1}| \cdot |E_k|)/|P_k|)$ is non-negative; strict inequality at step k gives strict positivity of the corresponding summand. The cumulative sum is therefore non-negative, non-decreasing, and strictly increasing at every compressed step. □

Remark 3.8 (Interpretation of the deficit). D_n counts, in bits, the cumulative reduction in distinguishable input configurations across the chain: the term at step k is the number of bits required to specify which input pair (A_{k-1}, e_k) produced the committed outcome A_k, beyond the bits already specified by A_k itself. Because $\log_2$ is used here as a coding-length count (not as a Shannon entropy with a prior on inputs), D_n is a pure cardinality construction; no probability axiom or thermodynamic primitive enters. The relation to the chain multiplicity bound is $D_n \leq \log_2 \prod_{k=1}^{n} \lceil 1/\alpha_k \rceil$, with equality when each $1/\alpha_k$ is an integer.

Remark 3.9 (Global form of the deficit). Under the chain identification $|\Delta_k| = |P_k|$ (Definition 2.23(c)), the per-step terms in Equation (5) telescope: each $|P_k|$ in step k's denominator cancels with $|\Delta_k| = |P_k|$ in step $(k+1)$'s numerator, so that only the initial input boundary, the environment cardinalities, and the terminal output cardinality remain. The cumulative deficit therefore admits the closed form

$$D_n \;=\; \log_2(|\Delta_0| \cdot \textstyle\prod_{k=1}^{n} |E_k|) \;-\; \log_2 |P_n|. \tag{6}$$

Two chains with the same input-space cardinality, the same total environment cardinality, and the same terminal output cardinality have the same cumulative compression deficit, regardless of the intermediate sizes $|P_1|, \ldots, |P_{n-1}|$. D_n is in this sense a global structural invariant of the chain endpoints: it depends only on the input boundary, the environment volume, and the terminal output, not on the path taken through intermediate cardinalities. In particular, the cumulative deficit depends only on endpoint cardinalities and total environment volume, and is independent of the intermediate pipeline structure.

Remark 3.10 (Per-step and per-bit normalizations of the deficit). Two normalized versions of D_n are convenient when comparing chains of different lengths or different input cardinalities.

Per-step deficit rate. Define

$$\bar{d}_n \;:=\; \frac{D_n}{n} \;=\; \frac{1}{n} \sum_{k=1}^{n} \log_2(1/\alpha_k), \tag{7}$$

the mean number of bits lost per step. Under uniform compression $\alpha_k = \alpha < 1$ for all k, $\bar{d}_n = \log_2(1/\alpha)$ is constant in n and equal to the per-step bit cost of compression. $\bar{d}_n$ is intensive: two chains with the same per-step compression have the same $\bar{d}_n$, regardless of length.

Per-bit fractional loss. For each step, the fractional bit loss is

$$\beta_k \;:=\; \frac{\log_2\big((|\Delta_{k-1}|\cdot|E_k|)/|P_k|\big)}{\log_2(|\Delta_{k-1}|\cdot|E_k|)} \;\in\; [0,1], \tag{8}$$

the fraction of distinguishable input bits at step k not preserved as outcome bits at the boundary. $\beta_k = 0$ corresponds to no compression ($|P_k| = |\Delta_{k-1}|\cdot|E_k|$); $\beta_k = 1$ corresponds to total compression to a single outcome ($|P_k| = 1$). The chain-level fractional loss admits a closed form under the chain identification $|\Delta_k| = |P_k|$ (Definition 2.23(c)):

$$\bar{\beta}_n \;:=\; \frac{D_n}{\log_2(|\Delta_0|\cdot\prod_{k=1}^{n}|E_k|)} \;=\; 1 - \frac{\log_2|P_n|}{\log_2(|\Delta_0|\cdot\prod_{k=1}^{n}|E_k|)} \;\in\; [0,1], \tag{9}$$

the fraction of total distinguishable input bits over the n-step chain that are not preserved as terminal outcome bits. The closed form follows from telescoping $D_n = \log_2((|\Delta_0|\cdot\prod_k|E_k|)/|P_n|)$ under the chain identification. Both normalizations remain pure cardinality constructions; no entropy or probability axiom enters.

Proposition 3.11 (Typed collision geometry)**.** *Let $M\colon \Delta_{k-1}\times E_k \to P_k$ be the boundary map at step k for any fixed deterministic realization of* Sample*.*

(i) *If for every $e\in E_k$ the partial map $M(\cdot\,,e)\colon \Delta_{k-1}\to P_k$ is injective (*per-environment injectivity*), every collision $M(a_1,e_1)=M(a_2,e_2)$ with $(a_1,e_1)\neq(a_2,e_2)$ satisfies $e_1\neq e_2$.*

(ii) *Symmetrically, if for every $a\in\Delta_{k-1}$ the map $M(a,\cdot)\colon E_k\to P_k$ is injective (*per-input injectivity*), every collision satisfies $a_1\neq a_2$.*

(iii) *Under Axioms ICC-1 and ICC-2(a), at least one collision exists by Theorem 3.1(i). If both per-environment and per-input injectivity hold, every such collision is necessarily* cross-axis *in the typed factorization $\Delta_{k-1}\times E_k$: $a_1\neq a_2$ and $e_1\neq e_2$.*

Proof. Parts (i) and (ii) are immediate from the contrapositive of injectivity for the relevant partial maps. Part (iii) combines (i), (ii), and Theorem 3.1(i). □

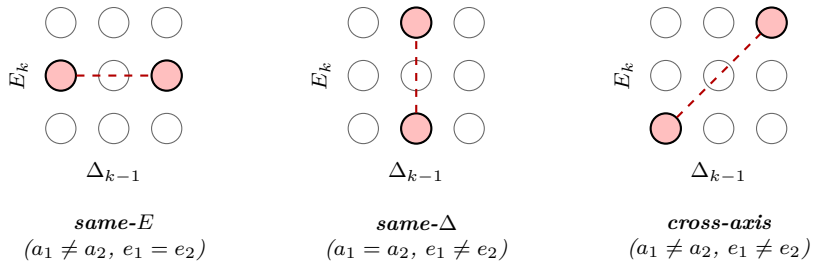

Figure 1: Typed collision geometry on the $\Delta_{k-1}\times E_k$ input grid (Propositions 3.11 and 3.14). Each panel shows two input configurations (filled circles) that map to the same committed outcome under M. **Left:** same-E collision, ruled out by per-input injectivity. **Center:** same-Δ collision, ruled out by per-environment injectivity. **Right:** cross-axis collision, the only kind permitted under both per-axis injectivities (Class A of Proposition 3.14).

Remark 3.12 (Geometric content beyond pigeonhole). Proposition 3.11 sharpens Theorem 3.1(i) from existence of collisions to a constraint on their typed location. Pigeonhole alone does not specify where in $\Delta_{k-1}\times E_k$ collisions live; the typed factorization does. The per-environment-injective regime is natural in near-reversible per-environment dynamics, where the flow followed by Activation is injective for each fixed environment input; the only collisions Axiom ICC-2(a)

can force in this regime are those that mix two environment inputs. This identifies a structural constraint on ambiguity generation: under per-environment injectivity, all compression-induced ambiguity must be mediated through environment variation. The constraint becomes nontrivial when combined with Theorem 3.3, where it determines the axis along which exponential ambiguity growth occurs (Theorem 3.13).

Theorem 3.13 (Structured chain ambiguity under typed injectivity)**.** *Assume Axioms ICC-1 and ICC-2(a), and suppose that for every step k the boundary map $M_k\colon \Delta_{k-1} \times E_k \to P_k$ satisfies per-environment injectivity in the sense of Proposition 3.11(i). Then:*

(i) *All collisions at each step are cross-axis (Proposition 3.11).*

(ii) *For any fixed terminal outcome $p_n \in P_n$, every compatible input history in $g_n^{-1}(p_n)$ must differ from every other compatible history in at least one environment coordinate e_k.*

(iii) *Consequently, the chain multiplicity lower bound of Theorem 3.3 is realized entirely through environment variation: the number of distinct environment sequences consistent with a single terminal outcome grows at least as*

$$\prod_{k=1}^{n} \left\lceil \frac{|\Delta_{k-1}| \cdot |E_k|}{|P_k|} \right\rceil .$$

Proof. Part (i) is Proposition 3.11(i),(iii). For (ii), per-environment injectivity implies that for fixed e_k the partial map $M_k(\,\cdot\,, e_k)\colon \Delta_{k-1} \to P_k$ is injective, so any single-step collision must satisfy $e_1 \neq e_2$. Iterating across the chain, two distinct input histories consistent with the same terminal outcome p_n must, at some step, produce the same intermediate output from different environment inputs; hence they differ in at least one environment coordinate. Part (iii) follows by combining this constraint with the chain multiplicity bound of Theorem 3.3(ii). □

Proposition 3.14 (Classification of typed boundary-map collision geometry)**.** *Let $M\colon \Delta_{k-1} \times E_k \to P_k$ be the boundary map at step k for any fixed deterministic realization of* Sample*, and assume Axioms ICC-1 and ICC-2(a). Define three collision types:*

- same-E*:* $M(a_1, e_1) = M(a_2, e_2)$ *with* $a_1 \neq a_2$, $e_1 = e_2$*;*
- same-Δ*:* $a_1 = a_2$, $e_1 \neq e_2$*;*
- cross-axis*:* $a_1 \neq a_2$, $e_1 \neq e_2$.

Classify M by its per-axis injectivity properties:

Class	*per-env injective*	*per-input injective*	*Forced collision type*
A	*yes*	*yes*	*cross-axis only*
B	*yes*	*no*	*same-Δ or cross-axis ($e_1 \neq e_2$)*
C	*no*	*yes*	*same-E or cross-axis ($a_1 \neq a_2$)*
D	*no*	*no*	*no constraint beyond existence*

Every boundary map belongs to exactly one class. Each class is non-empty under Axiom ICC-2(a), witnessed by:

- (A) $|\Delta| = |E| = 2$, $|P| = 2$, $M(a, e) = a \oplus e$ *(XOR pattern); unique collisions* $M(0,0) = M(1,1) = 0$ *and* $M(0,1) = M(1,0) = 1$ *are cross-axis.*
- (B) $|\Delta| = |E| = 2$, $|P| = 2$, $M(a, e) = a$*; collisions are same-Δ only.*
- (C) $|\Delta| = |E| = 2$, $|P| = 2$, $M(a, e) = e$*; collisions are same-E only.*

- (D) $|\Delta| = |E| = 2$, $|P| = 2$, M *given by* $\{(0,0)\to 0, (0,1)\to 0, (1,0)\to 0, (1,1)\to 1\}$*; all three collision types appear.*

Proof. The four classes are defined by the truth values of two independent Boolean properties, so they partition the set of boundary maps.

Class A. If $M(a_1, e_1) = M(a_2, e_2)$ with $(a_1, e_1) \neq (a_2, e_2)$, then $e_1 = e_2$ would contradict per-environment injectivity at e_1, and $a_1 = a_2$ would contradict per-input injectivity at a_1. Hence both $a_1 \neq a_2$ and $e_1 \neq e_2$: the collision is cross-axis.

Class B. Per-environment injectivity excludes same-E collisions by the argument above (first half). The witness $M(a, e) = a$ has both rows constant, producing same-Δ collisions $M(a, 0) = M(a, 1) = a$ for each a.

Class C. Symmetric to Class B.

Class D. With neither injectivity condition, no axis is excluded; the witness map produces all three collision types explicitly.

Non-emptiness of each class is verified by direct inspection of the witnesses; each satisfies Axiom ICC-2(a) ($|P| = 2 < 4 = |\Delta| \cdot |E|$). □

Remark 3.15 (Chain consequences of the classification). The class of each step's boundary map determines which coordinate of compatible input histories can agree. Under Class A at step k, two distinct compatible histories must differ in *both* the Δ and E coordinates at that step. Under Class B (the regime of Theorem 3.13), they must differ in at least one E-coordinate across all steps; the Δ-trajectory may be shared, in which case all chain ambiguity is carried purely by the environment sequence. Under Class C, they must differ in at least one Δ-coordinate; environment sequences may be shared. Under Class D, no coordinate-level constraint is forced; the chain bound of Theorem 3.3 still applies but no typed structure on the location of collisions is guaranteed.

Remark 3.16 (Cascade example). A concrete instantiation is a chain of irreversible two-input one-output Boolean gates (e.g., NAND), with $|\Delta_{k-1}| = 4$, $|E_k| = 1$, $|P_k| = 2$ at each step. The per-step multiplicity bound Equation (1) gives $\lceil 4/2 \rceil = 2$, and the chain bound Equation (2) grows as 2^n after n cascaded gates: at least one terminal output bit is consistent with $\geq 2^n$ distinct input histories. This scaling is formally analogous to the bookkeeping behind Landauer erasure [1, 11], where repeated many-to-one logical operations accumulate erased distinctions. The analogy is only structural: ICC counts compatible input histories, while Landauer theory assigns thermodynamic cost to logical erasure. Since ICC contains no entropy functional, temperature, or heat variable, no thermodynamic conclusion follows here.

ICC captures the combinatorial side of erasure; connecting that loss to energetic cost requires additional physical assumptions outside the present framework.

Corollary 3.17 (Irreversible gate cascades)**.** *A cascade of n two-input one-output irreversible gates with $|\Delta_{k-1}| = 4$, $|E_k| = 1$, $|P_k| = 2$ at each step has at least 2^n distinct input histories compatible with one terminal output.*

Proof. Immediate from Theorem 3.3: $\prod_{k=1}^{n} \lceil 4/2 \rceil = 2^n$. □

Corollary 3.18 (Discrete sequential index)**.** *The chain index $k \in \mathbb{N}$ is a discrete, irreversible, sequential parameter:*

(i) ***Discrete*** *by construction (each step is one application of M).*

(ii) ***Irreversible*** *by Theorem 3.1(ii) (the antecedent is non-recoverable).*

(iii) ***Sequential*** *by the typed pipeline and chain construction (Axiom ICC-3(a), Definition 2.12, Section 2.4): A_k cannot be formed without first committing A_{k-1}, because Φ requires*

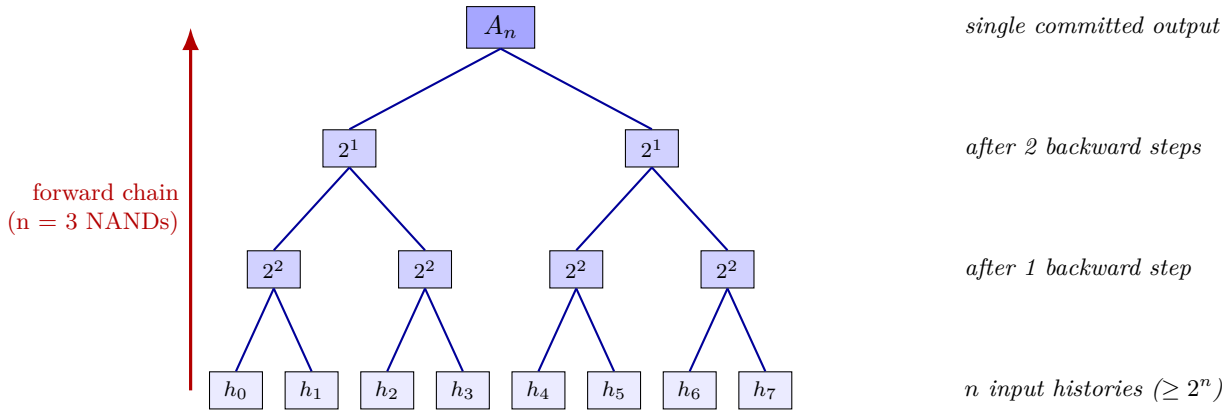

Figure 2: Chain-accumulation bound (Theorem 3.3, Equation (2)) applied to a cascade of n two-input one-output Boolean gates with $|\Delta_{k-1}| = 4$, $|E_k| = 1$, $|P_k| = 2$ at each step. Each backward step doubles the number of consistent input histories: at least 2^n distinct input sequences are consistent with a single terminal output. The figure illustrates $n = 3$, giving $2^3 = 8$ consistent input histories.

> *$\rho_0 = |A_{k-1}\rangle\langle A_{k-1}|$ as its initial condition, and each step of the chain consumes the committed output of the previous step. Well-formedness of ρ_0 uses the fixed basis of Definition 2.23(b).*

Proof. Discreteness: $k \in \mathbb{N}$ by construction (Definition 2.12, Section 2.4). Irreversibility: Theorem 3.1(ii). Sequentiality: by Definition 2.12, $A_k = \sigma_{C(e_k)}(\rho_0)$ where $\rho_0 = |A_{k-1}\rangle\langle A_{k-1}|$. Φ requires ρ_0 (Definition 2.2); Activation requires the output of Φ (Definition 2.6); Sample requires the output of Activation (Definition 2.8). By Axiom ICC-3(a), σ is the sequential composition of these stages. A_{k-1} must be committed before σ can begin. □

Corollary 3.19 (Cumulative non-recoverability)**.** *Let $A_0, A_1, \ldots, A_n$ be a commitment chain and let $j < n$. Even given the present state A_n, the full environment sequence $(e_1, \ldots, e_n)$, and complete pipeline knowledge (the admissible generator family $\{C(e_k)\}$, the interval Δt, and the fixed basis), the earlier state A_j is not in general recoverable. In particular, this maximal information set does not determine A_j.*

Proof. This is a composition of per-step obstruction, not a new mechanism emerging at the chain level. The argument has two parts: per-step non-recoverability is unconditional under the maximal information set, and the chain composes those per-step failures.

Per-step non-recoverability is unconditional. Fix any step $k \in \{1, \ldots, n\}$. By Theorem 3.1(ii), obstacle (c): the committed outcome A_k does not determine the distribution μ_k from which it was drawn (Axiom ICC-3(c)). This obstruction is unconditional: knowing e_k specifies $C(e_k)$ and hence $\Phi^{\Delta t}_{C(e_k)}$, but does not reconstruct μ_k from the single draw A_k. Knowing the full pipeline (the generator family, Δt, and the basis) fixes the map $\rho_0 \mapsto \mu_k$; it does not invert Sample's non-deterministic selection. By Theorem 3.1(ii), obstacle (b): even if μ_k were known, A_{k-1} is not recoverable because Activation is many-to-one (Axiom ICC-3(b)); distinct density matrices may produce the same μ_k, so the information needed to identify $\rho_0 = |A_{k-1}\rangle\langle A_{k-1}|$ uniquely is destroyed at this stage. In the diagonal-extraction realization, the destroyed information is the off-diagonal entries of $\rho(\Delta t)$; knowing any element of the maximal information set does not supply it. Per-step non-recoverability of A_{k-1} from A_k therefore holds regardless of the supplementary data.

Propagation across the chain. Recovery of A_j from A_n requires, at minimum, recovering A_{n-1} from A_n, then A_{n-2} from A_{n-1}, and so on down to A_j. Each such step is the single-step

recovery problem just analysed, applied to a consecutive pair (A_{k-1}, A_k) with the corresponding e_k and pipeline data available. Since single-step non-recoverability holds unconditionally at every step (no element of the maximal information set removes it), the entire chain of recovery attempts fails. No matter which $j < n$ is targeted, A_j cannot be determined from the maximal information set.

Multiple compatible histories. The non-injectivity established in Theorem 3.1(i) implies that, at each step, multiple input configurations produce the same committed outcome. The cumulative effect over n steps is that multiple distinct history sequences $(A_0, e_1, \ldots, e_j, A_j, \ldots)$ are consistent with the same present state A_n and the same environment record $(e_1, \ldots, e_n)$. The antecedent state A_j is therefore not uniquely determined even when the environment sequence is fully known. □

3.1 Dependency localization

The following propositions show that different structural consequences localize to different axiom subparts by examining what survives under three targeted weakenings of the axioms.

Proposition 3.20 (Deterministic truncation)**.** *Consider the truncated pipeline*

$$\tau_{C(e_k)} := \text{Activation} \circ \Phi^{\Delta t}_{C(e_k)} \colon \mathrm{DM}(N) \to (P_k, \mu_k).$$

Under Axiom ICC-1, Axiom ICC-2(a,b), and Axiom ICC-3(b), the following hold:

(a) *Strict boundary compression:* $|P_k| < |\Delta_{k-1}| \cdot |E_k|$.

(b) *Deterministic non-invertibility at the distribution level: distinct inputs may yield the same* (P_k, μ_k)*, so the prior input is not recoverable from the activated output.*

(c) *Sequential dependency up to activation:* (P_k, μ_k) *cannot be formed without first committing* A_{k-1} *and evolving* $\rho_0 = |A_{k-1}\rangle\langle A_{k-1}|$ *through* Φ*.*

What does not survive is any result stated at the level of a committed outcome $A_k \in \Delta_k$ *or of a continuing commitment chain. In particular, without Sample there is no committed history to which Theorem 3.1(iii) or Corollary 3.19 can apply.*

Proof. Parts (a) and (b) follow from Axiom ICC-1, Axiom ICC-2(a,b), and Axiom ICC-3(b) by the same arguments as Theorem 3.1(i) and the first two obstacles of Theorem 3.1(ii), none of which invoke Sample. Part (c) follows from the typing of M (Definition 2.12) and the chain construction (Section 2.4): Φ requires ρ_0 as initial condition, and ρ_0 requires a committed A_{k-1}. The chain requires $A_k \in \Delta_k$ to form the next initial state; the truncated pipeline produces only (P_k, μ_k), which is not a valid input to M. □

Proposition 3.21 (Deterministic commitment)**.** *Suppose Axiom ICC-3(c) is replaced by the requirement that Sample is deterministic: the same* μ_k *always yields the same* A_k*. Then:*

(a) *Theorem 3.1(i) (boundary compression) is unaffected.*

(b) *Corollary 3.18 (sequential dependency of the index) is unaffected: the chain proceeds as before.*

(c) *Theorem 3.1(ii) (non-recoverable antecedent) loses one of its three obstacles: obstacle (c), which relies on the non-determinism of Sample, no longer applies. The remaining two obstacles (pipeline non-invertibility and Activation being many-to-one) still block recovery of the input from* A_k *at distinct pipeline stages.*

(d) *Theorem 3.1(iii) (non-deterministic irreversibility) does not hold: the same input now produces the same committed outcome. (Note: the deterministic/non-deterministic distinction in Theorem 3.1(iii) is substantive only when μ_k is non-degenerate; a point mass renders it moot regardless of whether Sample is deterministic.) The deterministic pipeline remains non-invertible, so the input is still not recoverable from A_k.*

Proof. Parts (a) and (b) do not depend on Axiom ICC-3(c). For (c), obstacles (a) and (b) of Theorem 3.1(ii) cite Axiom ICC-2(b) and Axiom ICC-3(b) respectively; neither involves Sample. Part (d): deterministic Sample with deterministic flow and deterministic Activation makes the full pipeline deterministic, so the same input yields the same outcome. Non-invertibility of Activation $\circ\, \Phi$ (Axiom ICC-2(b)) still prevents recovery. □

Proposition 3.22 (Without pipeline non-invertibility)**.** *Drop Axiom ICC-2(b): the deterministic pipeline* Activation $\circ\, \Phi$ *is no longer required to be non-invertible as a separate axiom-level constraint. Then:*

(a) *Theorem 3.1(i) (boundary compression) is unaffected: it depends only on Axiom ICC-2(a) and Axiom ICC-1.*

(b) *Corollary 3.18 (sequential dependency of the index) is unaffected.*

(c) *Theorem 3.1(ii) (non-recoverable antecedent) loses obstacle (a). Recovery of μ_k from A_k is still blocked by Sample (Axiom ICC-3(c)), and recovery of off-diagonal information is still blocked by Activation (Axiom ICC-3(b)). But the deterministic pipeline is no longer guaranteed to be an additional independent obstacle. In this regime, Axiom ICC-2(a) (boundary compression) remains the sole guaranteed source of non-injectivity at the input–output boundary. Equivalently: Axiom ICC-2(b) is not required for non-recoverability; it strengthens but does not generate it.*

(d) *Theorem 3.1(iii) loses the clause "neither application can be inverted": Sample is still non-deterministic, so repeated application may yield different outcomes, but without Axiom ICC-2(b) the deterministic portion of the pipeline may in principle be invertible.*

Proof. Parts (a) and (b) do not cite Axiom ICC-2(b). For (c), obstacles (b) and (c) of Theorem 3.1(ii) cite Axiom ICC-3(b) and Axiom ICC-3(c) respectively; neither involves Axiom ICC-2(b). For (d), the non-determinism of Sample follows from Axiom ICC-3(c) alone; the invertibility claim in Theorem 3.1(iii) cites Axiom ICC-2(b), which is no longer available. □

Proposition 3.23 (Without Activation many-to-one)**.** *Drop Axiom ICC-3(b): the deterministic extraction of μ_k from $\rho(\Delta t)$ is no longer required to be many-to-one. Then:*

(a) *Theorem 3.1(i) (boundary compression) is unaffected: it depends only on Axiom ICC-2(a) and Axiom ICC-1.*

(b) *Corollary 3.18 (sequential dependency of the index) is unaffected.*

(c) *Theorem 3.1(ii) (non-recoverable antecedent) loses obstacle (b). With injective Activation, $\rho(\Delta t)$ is recoverable from μ_k. But recovery from A_k is still blocked by Sample (Axiom ICC-3(c)), and recovery of the input from $\rho(\Delta t)$ is still blocked by pipeline non-invertibility (Axiom ICC-2(b)). The non-recoverability conclusion of part (ii) survives via the remaining two obstacles.*

(d) *Theorem 3.1(iii) is unaffected: it cites Axiom ICC-2(b) and Axiom ICC-3(c), neither of which involves Axiom ICC-3(b).*

Proof. Parts (a), (b), (d) do not cite Axiom ICC-3(b). For (c), obstacle (a) cites Axiom ICC-2(b) and obstacle (c) cites Axiom ICC-3(c); neither involves Axiom ICC-3(b). A concrete countermodel exhibiting injective Activation on a model with non-unitary Φ (so that Axiom ICC-2(b) still holds independently) is verified in Lean 4 against `mathlib4`; see Appendix A. □

Proposition 3.24 (Typed decomposition requires ICC-3(b))**.** *The typed three-stage decomposition of Axiom ICC-3(a), $\sigma = \text{Sample} \circ \text{Activation} \circ \Phi$, with type signature*

$$\Phi\colon \mathrm{DM}(N) \to \mathrm{DM}(N), \quad \text{Activation}\colon \mathrm{DM}(N) \to \Delta(P_k), \quad \text{Sample}\colon (\Delta(P_k), \mu_k) \to P_k,$$

requires Axiom ICC-3(b) for the codomain of Activation to be $\Delta(P_k)$. Without Axiom ICC-3(b), Activation may have codomain $\mathrm{DM}(N)$ rather than $\Delta(P_k)$; in that case Sample's input type is no longer a classical distribution, and the primitive "commit a single element of P_k" has no canonical realization without an additional basis-selection axiom not present in the framework.

Proof. Definition 2.8 types Sample as $(\Delta(P_k), \mu_k) \to P_k$ and requires its first argument to lie in $\Delta(P_k)$. Definition 2.6 under Axiom ICC-3(b) types Activation as $\mathrm{DM}(N) \to \Delta(P_k)$; in the diagonal-extraction realization, this maps $\rho(\Delta t)$ to a probability vector indexed by P_k. Removing Axiom ICC-3(b) does not affect Φ's codomain ($\mathrm{DM}(N)$); if Activation is then permitted to be informationally complete (e.g. the identity on $\mathrm{DM}(N)$), its codomain becomes $\mathrm{DM}(N) \not\subseteq \Delta(P_k)$. The composition Sample $\circ$ Activation then has a type mismatch unless an auxiliary basis-selection primitive is introduced. In particular, Axiom ICC-3(b) is not merely a source of representational loss but a structural condition required to separate continuous dynamics from discrete commitment; without it, the framework admits no canonical commitment boundary. □

Remark 3.25 (Structural role of ICC-3(b) beyond localization). Proposition 3.24 clarifies that Axiom ICC-3(b) plays two distinct structural roles, not one. First, it identifies obstacle (b) of Theorem 3.1(ii) (Proposition 3.23): Activation discards the information distinguishing distinct evolved density matrices. Second, it licenses the typed decomposition itself: without it, the boundary between continuous quantum dynamics and discrete classical commitment is not located at any specific stage, and the three-stage typing degenerates. A reader who concludes from Proposition 3.23 alone that Axiom ICC-3(b) is auxiliary should note that dropping it dissolves the framework's typing, not just its representation-loss obstacle.

Proposition 3.26 (Strict dependency separation)**.** *The three weakenings of Propositions 3.20 to 3.22 produce three pairwise distinct loss patterns. No two structural consequences are lost together under all three weakenings.*

Proof. By inspection of the following table, each column differs from every other column in at least one row.

	Drop Sample (Proposition 3.20)	Det. Sample (Proposition 3.21)	Drop ICC-2(b) (Proposition 3.22)
(i) compression	unaffected	unaffected	unaffected
(ii) non-recoverable antecedent	obstacles (a),(b) hold; (c) N/A	obstacles (a),(b) hold; (c) lost	obstacles (b),(c) hold; (a) lost
Corollary 3.18 sequential	N/A (no committed output)	unaffected	unaffected
(iii) non-det. irrev.	N/A	lost	nondeterminism remains; invertibility clause lost
chain continues	no	yes	yes

The ingredients of non-recoverability are therefore independently removable: committed-history claims depend on Axiom ICC-3(c); deterministic recovery barriers depend on Axiom ICC-2(b); boundary compression depends on Axiom ICC-2(a) and Axiom ICC-1; and chain sequentiality depends on the typing of M and the chain construction. □

3.2 Scope, dependencies, and interpretation

The remarks in this subsection summarize the load-bearing axiom dependencies of the structural results, situate them against the pigeonhole baseline, and isolate interpretive content from formal claims.

Remark 3.27 (Scope). Theorem 3.1 is a structural consequence of the axioms and typed construction. It does not establish a physical law, solve the measurement problem, or make claims about specific physical systems. Whether any physical process satisfies Axioms ICC-1–ICC-3 is a separate question, addressed in Section 4.

Remark 3.28 (Axiom dependency). The following summarizes which axiom subparts are load-bearing for each result:

Result	Dependencies
(i)	Axiom ICC-2(a), Axiom ICC-1
(ii)	Axiom ICC-2(b) (obstacle a), Axiom ICC-3(b) (obstacle b), Axiom ICC-3(c) (obstacle c)
(iii)	Axiom ICC-2(b), Axiom ICC-3(c)
Corollary 3.18	(ii) (irreversibility); Axiom ICC-3(a), Definition 2.12, Section 2.4 (sequential dependency)
Corollary 3.19	(ii)
Proposition 3.20	Axiom ICC-1, Axiom ICC-2(a,b), Axiom ICC-3(b)
Proposition 3.21	all except Axiom ICC-3(c)
Proposition 3.22	all except Axiom ICC-2(b)
Proposition 3.23	all except Axiom ICC-3(b)
Proposition 3.26	Proposition 3.20, Proposition 3.21, Proposition 3.22

Each cited axiom subpart is load-bearing for the stated result.

Remark 3.29 (Pigeonhole and the typed decomposition). A reader will recognize Theorem 3.1(i) as pigeonhole on finite sets: under Axioms ICC-1 and ICC-2(a), the boundary map cannot be injective, so the input pair (A_{k-1}, e_k) is not uniquely determined by the committed outcome A_k. The weak form of Theorem 3.1(ii) follows directly. We adopt this reading without hedging.

The framework's contribution sits on top of pigeonhole, not in place of it. Axiom ICC-2(b) is a continuous-interior claim about Activation $\circ\, \Phi$ being non-invertible as a function on density matrices; Remark 2.19 exhibits an explicit counterexample where Axiom ICC-2(a) holds and Axiom ICC-2(b) fails, confirming the two conditions are independent. Axiom ICC-3(b) names the off-diagonal entries of the evolved density matrix as the specific side information destroyed. Axiom ICC-3(c) supplies non-determinism at the commitment stage. The combination of these three axioms is what extends the pigeonhole-derived weak non-recoverability into the strong claim of Corollary 3.19, where recovery fails even under maximal pipeline information. The typed decomposition is the bookkeeping that makes the per-axiom incremental content traceable, with pigeonhole as the entry-point row. The additional content of the framework lies in (i) quantifying the resulting ambiguity (Theorem 3.3) and (ii) constraining its structure via the typed factorization (Proposition 3.11 and Theorem 3.13), neither of which is captured by the raw pigeonhole argument.

Remark 3.30 (Interpretation). The following is interpretive and not part of the formal results.

The chain index k may correspond to a discrete time parameter if each application of Φ is identified with a physical integration interval Δt (fixed across steps by convention; see Remark 2.4). Whether k constitutes "time" in the physical sense is outside the formal framework. The non-recoverability established in Theorem 3.1(ii) and Corollary 3.19 is a structural property of the typed pipeline; its relationship to thermodynamic irreversibility depends on whether specific physical processes satisfy the axioms.

Remark 3.31 (Structural arrow and locked past). Theorem 3.1 endows the commitment chain with a distinguished forward orientation: each application of the pipeline produces a later committed state through a map that is non-invertible at the committed level. In this sense the framework defines a structural arrow of commitment. The resulting past is not merely unobserved but compressed: distinct prior input configurations may lead to the same committed outcome, so the exact antecedent branch is not in general recoverable from the present committed state. This also sharpens the distinction between an inaccessible past and a uniquely defined past state. In a classical setting, a prior state may be unreachable in practice while still remaining a well-defined antecedent of the present. In ICC, by contrast, a present committed state need not determine a unique antecedent branch: multiple prior input configurations may be compatible with the same present commitment. The obstruction is therefore stronger than epistemic inaccessibility alone. Relative to the committed description, backward return is not simply blocked; it is in general undefined, because there is no uniquely specified prior branch within the committed description to return to.

This is a formal statement about the typed pipeline, not yet a claim about physical time or thermodynamic irreversibility. Figure 3 summarizes this asymmetry schematically; it is not an additional assumption but a visual summary of the structural consequences established above.

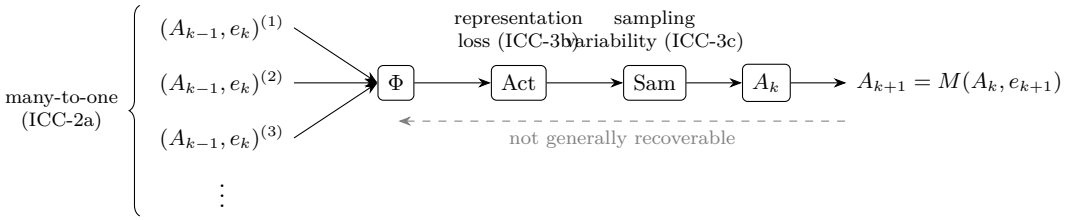

Figure 3: Structural asymmetry of committed history. Multiple distinguishable input configurations (Definition 2.14) feed the pipeline; compression and commitment produce a single committed outcome A_k, which becomes the input to the next step. Exact backward recovery is not generally available: obstacles (a), (b), and (c) of Theorem 3.1(ii) block recovery at distinct pipeline stages. Under per-environment injectivity, all branching in this diagram corresponds to distinct environment sequences (Theorem 3.13).

Remark 3.32 (Stage-local determinism). ICC separates deterministic evolution from committed outcome selection. The flow map and the extraction map are deterministic, while the sampling stage is explicitly non-deterministic at the level of the formal pipeline. The framework therefore does not oppose determinism everywhere; it localizes the breakdown of global determinism, if any, to the commitment step. Non-deterministic sampling is responsible for trajectory variability under the same input (Theorem 3.1(iii)), but is not required for non-recoverability, which is already implied by boundary compression alone.

4 Open Conjecture

This section is interpretive and conjectural. The formal framework does not establish that any physical system instantiates ICC. The main theorem is a structural consequence of the axioms.

The purpose of this section is only to identify what a candidate physical realization would need to satisfy.

The typed pipeline of Section 2.2 admits a candidate reading in open quantum systems. The generator $C(e_k)$ may be read as an environment-dependent dissipative law, for example a Lindblad generator. The induced flow $\Phi^{\Delta t}_{C(e_k)}$ may be read as the corresponding CPTP evolution. Activation may be read as extraction of outcome probabilities in a chosen or decoherence-preferred basis, for example by applying the Born rule to the diagonal of the evolved density matrix in that basis. The sampling stage remains the unresolved step: the mechanism by which a single committed outcome is selected from the distribution μ_k. Decoherence and open-system dynamics can motivate the deterministic stages of the pipeline (flow and extraction) but do not by themselves specify a single-outcome selection mechanism. Consistent histories [4, 6] sits in this gap by assigning classical probabilities to families of decoherent histories rather than committing to single outcomes; ICC differs in committing to one outcome per step and localizing the selection mechanism as the open Sample stage. ICC therefore stands in a complementary relationship to consistent histories: the latter remains agnostic about commitment, the former axiomatizes commitment as a primitive at the cost of inheriting the family-selection question.

The compression claim in Axiom ICC-2 is an abstraction at the level of distinguishable finite possibilities. In a candidate quantum realization, it would correspond to the assertion that the number of distinguishable committed outcomes is strictly smaller than the number of distinguishable input configurations at each step. In standard quantum measurement settings, this inequality is often easy to satisfy because the number of measurement outcomes is typically much smaller than the number of input configurations; but the axiom compares finite boundary cardinalities, not outcome count to Hilbert-space dimension, and its nontriviality depends on the chosen coarse-graining. Whether it holds in any particular regime depends on the system, the environment coupling, the coarse-graining, and the measurement context. The axiom is not a claim about the continuous interior of the density-matrix state space; it is a finite-boundary constraint on the distinguishable sets that bracket the pipeline.

The two conjectures below distinguish two levels of precision at which a physical regime might instantiate ICC. The first requires all five checklist items (Remark 4.3), including an exact single-outcome selection mechanism. The second replaces that requirement with a weaker phenomenological criterion that does not presuppose a resolution of the measurement problem.

Conjecture 4.1 (Exact physical realization)**.** There exists a physically relevant regime (for example, a class of open quantum systems with environment-induced decoherence) for which the generator, induced flow, probability extraction, and single-outcome selection satisfy Axioms ICC-1–ICC-3 exactly, so that the structural conclusions of Theorem 3.1 apply directly to that regime. This conjecture is conditional on the resolution of checklist item (4) in Remark 4.3: it requires a physical mechanism that selects exactly one outcome from the distribution μ_k. Because no agreed-upon such mechanism currently exists, Conjecture 4.1 is currently unsatisfiable in the literal sense: it can be compared against a physical regime only once the measurement problem is resolved.

Conjecture 4.2 (Effective physical realization)**.** There exists a physically relevant regime for which checklist items (1), (2), (3), and (5) of Remark 4.3 hold, and in which the distribution μ_k produced by Activation is operationally sharply peaked: there exists a detection threshold ϵ such that the dominant outcome carries probability at least $1 - \epsilon$, with all other outcomes suppressed below ϵ. Under this phenomenological criterion, the Sample stage is treated as effectively deterministic for that step, and the conclusions of Theorem 3.1 apply at an effective, operational level, with non-determinism suppressed below detection thresholds.

In the class of candidate quantum realizations exhibiting the commutator structure of Appendix D (in particular, $[O_A, H_{\text{int}}^{A\text{-}B}(e)] = 0$), peakedness in this sense is controlled by the gate or measurement-error parameter of the pre-Activation channel rather than by the decoherence

rate; environment-induced decoherence plays the separate role of ensuring that the apparatus state entering Activation is genuinely diagonal in the chosen basis, so that probability extraction is well-defined. Conflating these two roles—taking strong decoherence to imply a peaked diagonal—is incorrect for this class: pure dephasing acting on a coherent superposition preserves the diagonal entries while suppressing only the off-diagonals. This two-control structure is made explicit in the spin-boson model of Appendix D, where the trace distance from the dominant outcome admits an analytical bound that is independent of the decoherence rate in the strong-decoherence regime. In realizations whose system–apparatus and apparatus–bath couplings do not exhibit this commutator structure, the diagonal can in principle acquire bath-dependent contributions; the present conjecture therefore asserts the two-control decomposition for the class of realizations in which it is structurally available, and a worked realization is exhibited in Appendix D, not universally.

This conjecture does not require a solution to the measurement problem; it requires only that a regime exists in which both controls (decoherence sufficient for diagonality, gate or measurement error below the detection threshold) are simultaneously favorable. The non-determinism of Theorem 3.1(iii) applies non-vacuously only when μ_k is non-degenerate; in the effective regime μ_k has nonzero tails of total mass at most ϵ, so the resulting non-determinism is operationally suppressed below the detection threshold. A candidate regime is environment-induced einselection of pointer states in macroscopic apparatuses (cf. [15]). See Appendix D for the worked S+A+B realization.

Scope of the structural results in the effective regime. In the effective regime, the structural results of Theorem 3.1 apply with one important qualification. Theorem 3.1(i) (boundary compression) and Theorem 3.1(ii) (non-recoverability of prior inputs from committed outcomes) hold by their original arguments: both rely on the boundary cardinality condition Axiom ICC-2(a) and the deterministic stages Φ, Activation, neither of which is altered by the effective interpretation of Sample. Theorem 3.1(iii) (non-deterministic irreversibility) becomes *operationally vacuous*: when Sample is treated as effectively deterministic at confidence $1 - \epsilon$, repeated application to the same input produces the dominant outcome on each run, and the trajectory variability that makes (iii) substantive is suppressed below the detection threshold. The chain accumulation bound of Theorem 3.3 survives at the structural level (it depends on cardinalities, not on the Sample mechanism), but its physical interpretation in the effective regime is that, with high probability, each chain run produces the unique dominant trajectory. This trade-off is the price of replacing the exact commitment criterion with a phenomenological sharpening condition; it is not a hidden cost but a deliberate scope choice. In particular, Theorem 3.1(iii) becomes vacuous in this regime, but Theorem 3.3 and Proposition 3.11 remain fully operative: the structural ambiguity and its typed collision geometry persist even when non-determinism is operationally suppressed.

The two conjectures are not competing: a resolution of the measurement problem consistent with Conjecture 4.1 would also satisfy Conjecture 4.2, since any exact selection mechanism implies the effective criterion. The effective conjecture is the weaker and more immediately testable claim.

Neither conjecture is a claim of proof here. The point is that ICC may serve as an abstract target against which candidate physical regimes can be compared. Any claimed realization would need to satisfy the following diagnostic checklist.

Remark 4.3 (Realization checklist). A candidate physical regime instantiates ICC only if it justifies the following items. Items (1), (2), (3), and (5) are required by both Conjecture 4.1 and Conjecture 4.2. Item (4) is required only by Conjecture 4.1; Conjecture 4.2 replaces it with the phenomenological sharpening criterion stated in that conjecture.

(1) **Finite boundaries.** The sets Δ_k and E_k are well-defined and finite for the system in question.

(2) **Motivated coarse-graining.** The inequality $|P_k| < |\Delta_{k-1}| \cdot |E_k|$ holds under a physically motivated coarse-graining, not an ad hoc one.

(3) **Basis selection.** A preferred or decoherence-selected basis exists for the Activation stage.

(4) **Commitment criterion** (required by Conjecture 4.1 only). A criterion identifies the transition from distribution to single outcome. This is the measurement problem; no agreed-upon physical mechanism currently exists. Conjecture 4.2 replaces this item with a phenomenological sharpening condition (see that conjecture).

(5) **Pipeline non-invertibility.** The constraint Axiom ICC-2(b) holds for the relevant boundary inputs (exactly under Conjecture 4.1; at an effective level under Conjecture 4.2).

Even if the exact axioms fail in nature, Conjecture 4.2 identifies a weakened but independently coherent realization target.

This section does not claim that the candidate mapping above is unique. Other typed decompositions of physical processes may yield overlapping or competing axiomatic packages. The framework is neutral on which physical interpretation, if any, is correct.

5 Scope and Limitations

5.1 Strength of axioms and structural scope

The axioms should be read as a sufficient package, not as weak premises from which non-recoverability unexpectedly emerges. Axiom ICC-2 encodes the boundary compression claim, and Axiom ICC-3 fixes the typed decomposition of the commitment process. These are strong, load-bearing assumptions. Several parts of the main theorem are therefore structural consequences once the pipeline is fixed. The value of the framework lies in isolating that sufficient structure, tracking which assumptions each result uses, and providing a clear target for weakening, comparison, or attempted realization.

The theorem is structural. It applies to any process satisfying the stated axioms, but it does not show that any physical system instantiates ICC. Whether quantum decoherence and measurement, or any other physical regime, satisfies Axioms ICC-1–ICC-3 is a separate empirical and interpretive question. The candidate mapping in Section 4 is illustrative, not verified.

The compression claim in Axiom ICC-2 is made at the finite boundary level: it concerns the cardinalities of distinguishable input configurations and distinguishable outcomes. It does not make claims about the continuous density-matrix interior of the pipeline. The flow map Φ acts on $\mathrm{DM}(N)$, which is a continuous space; the compression inequality constrains only the finite sets that bracket the pipeline at input and output. This distinction is important: the pigeonhole argument in the proof of Theorem 3.1(i) applies because both sides of the inequality are finite.

5.2 The measurement problem

The framework does not solve the measurement problem; it localizes the unresolved step to the sampling stage. The formal results hold for any mechanism satisfying Axiom ICC-3(c), whether objective-collapse, branching, hidden-variable, or otherwise. Whether the chain index k corresponds to physical time depends on identifications outside the formal system.

5.3 Basis and pointer-observable assumptions

The quantum realization of Section 2.6 involves a choice of basis for the Activation stage. In the decoherence setting, this basis may be selected by environment-induced superselection, but the axioms do not specify how it is chosen. Dependence on basis choice is a limitation of the quantum interpretation, not of the abstract framework: the latter treats Activation as an

abstract many-to-one map. More broadly, the framework's definitions use density matrices and Dirac notation; the abstraction is that the specific dynamics and sampling mechanism are left open, not that the formalism is independent of quantum state-space structure.

Two additional framework-level hypotheses constrain the class of candidate realizations. First, the generator C must be admissible (Definition 2.1): the initial value problem must have a unique solution for the flow map Φ to be well-defined. Second, a single fixed orthonormal basis must be maintained across all steps of the chain (Definition 2.23(b)), ensuring that the identification $P_k \cong \Delta_k$ preserves well-typedness.

The single-basis requirement of Definition 2.23(b) defines the class of systems to which the framework applies: those with a fixed pointer observable across all environments and steps. Systems whose committed outcomes index incompatible observables, such as a qubit chain alternating between σ_z and σ_x measurements, fall outside this class. The per-step relaxation to a basis $\{|i\rangle\}^{(k)}$ at each step extends the framework's reach to such systems and preserves the per-step structural results of Theorem 3.1, since each step's Activation, Sample, and pigeonhole argument depend only on the basis chosen at that step. A full development of the per-step variant, including how the identification $P_k \cong \Delta_k$ propagates labels across changing bases, is left to future work.

5.4 Markov and multi-channel extensions

A second relaxation worth flagging is the *non-Markov chain variant*. The pipeline as stated reinitializes $\rho_0 = |A_{k-1}\rangle\langle A_{k-1}|$ at the start of each step, which closes the Markov property: the next step depends on the previous committed state and nothing else. A non-Markov variant would carry forward additional state between steps (e.g., a memory register or correlated environment). The boundary cardinality argument of Theorem 3.1(i) applies in either case, since it concerns the sets Δ_{k-1}, E_k, P_k at the single-step boundary; the chain accumulation bound of Theorem 3.3 requires the chain identification $P_k \cong \Delta_k$ and may take a modified form when memory is carried between steps. Together with the per-step basis variant, these two relaxations identify the most physically pressing extensions of the present framework.

The pipeline as stated treats the commitment chain as singular: at each step, one Activation, one Sample, one committed outcome. In systems with multiple concurrent decoherence channels, this corresponds to an implicit choice of consistent family in the sense of [7]. The axioms do not specify how this family is selected, nor whether different families yield mutually compatible chains. Extension to multi-channel settings would require either an explicit family-selection criterion or a treatment of compatibility conditions across families; this is left open.

5.5 Categorical placement

The deterministic stages of the pipeline (Φ and Activation) embed cleanly as morphisms in standard probabilistic categorical settings such as FinStoch [3]. The single-outcome commitment of Sample is the structural obstruction: Markov kernels return distributions rather than sampled elements, and the transition $(P_k, \mu_k) \to A_k$ therefore exits the Markov-categorical setting. Effectus-theoretic treatments [10], which admit commitment as a primitive, are a natural direction for formalization.

5.6 Comparison to Bennett–Landauer

The framework's placement of structural irreversibility also differs from the Bennett–Landauer treatment [1, 11] of information-thermodynamic irreversibility. ICC localizes structural irreversibility at the Activation step (where the off-diagonal information distinguishing distinct evolved states is discarded by the many-to-one extraction μ_k). Bennett–Landauer localizes the thermodynamic cost at the memory-erasure step (resetting a recorded outcome to a blank

state requires at least $kT \ln 2$ of work). These are not competing claims about the same step; they are claims about different stages of the same cycle, in different currencies (cardinality vs energy). Bridging them would require an entropy-cardinality correspondence, which the present framework does not provide. This gap is intentional: ICC does not propose to derive the Landauer bound, and the structural results above hold without invoking thermodynamic primitives.

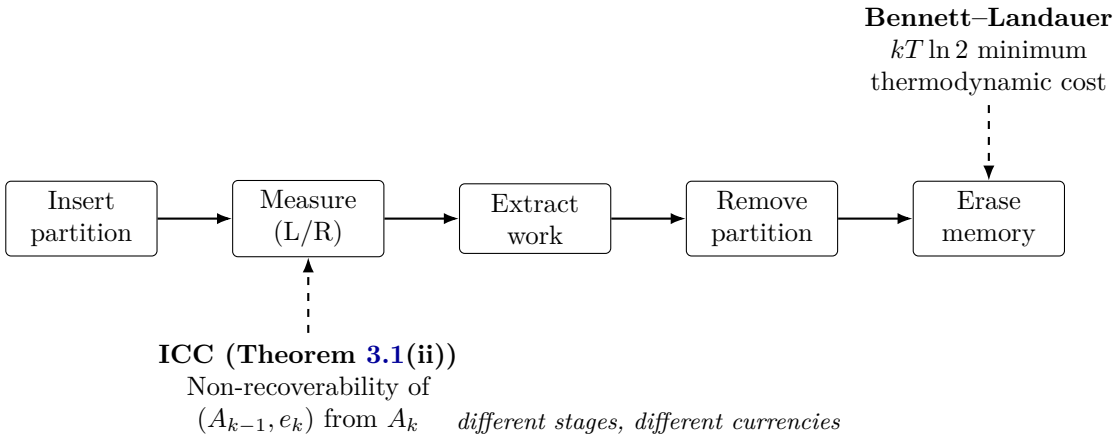

Figure 4: ICC and Bennett–Landauer localize the irreversibility of the Szilard cycle at different stages. ICC's structural non-recoverability (Theorem 3.1(ii)) is generated at the measurement step, where Activation discards the information distinguishing distinct evolved states. Bennett–Landauer's thermodynamic bound is incurred at the memory-erasure step. Neither claim implies the other; the present framework does not include the entropy–cardinality bridge that would relate them.

5.7 Regime classification and verification scope

The framework also excludes familiar reversible cases. For example, reversible classical permutation dynamics with a bijective boundary map $(\Delta_{k-1} \times E_k) \to P_k$ do not satisfy Axiom ICC-2(a). Likewise, closed-system unitary evolution without an Activation/Sample stage is outside ICC: the deterministic evolution is invertible and no committed boundary compression is imposed. These cases are not counterexamples to ICC; they delimit the class of processes the framework is meant to describe.

Read this way, Axiom ICC-2(b) functions as a *regime classifier* rather than a universal physical law: it separates the class of pipelines for which the deterministic stages Activation $\circ\, \Phi$ are non-invertible (ICC applies and contributes obstacle (a) of Theorem 3.1(ii)) from the class where they are invertible (ICC's representational results do not apply in the same form, even when the boundary cardinality condition Axiom ICC-2(a) is satisfied). The counterexample of Remark 2.19 inhabits the second class explicitly. Treating Axiom ICC-2(b) as a regime condition makes its role transparent: it is part of the framework's scope description, not a posit about all physical processes.

The dependency check implemented in SymPy verifies the propositional implication structure only. It confirms that the axiom-to-theorem implication graph is internally consistent, that each axiom subpart is used, and that removing a single axiom permits at least one theorem part to fail. It does not verify the mathematical proofs, semantic independence, or any physical claim.

The present paper does not claim uniqueness or finality for this axiom set. Different decompositions of the commitment pipeline, or weaker variants of the compression axiom, may yield overlapping or distinct structural consequences. The framework may remain useful even if future work revises, weakens, or replaces individual axioms, provided an analogous typed-pipeline

structure is retained. Identifying the weakest sufficient conditions for each theorem part is a natural direction for further work.

6 Conclusion

This paper introduced ICC, a typed axiomatic framework for irreversible commitment processes. Each step is decomposed into three stages with explicit domain and codomain types: deterministic flow on density matrices, deterministic many-to-one extraction of outcome probabilities, and non-deterministic sampling to a single committed outcome. The compression claim is stated at the finite boundary level of distinguishable inputs and outputs, not on the continuous interior state space.

Under Axioms ICC-1–ICC-3, the main theorem is structural. It establishes strict boundary compression, non-recoverability of prior inputs from committed outcomes, and non-deterministic irreversibility; a corollary captures the discreteness, irreversibility, and sequentiality of the chain index. These consequences do not arise from weak premises: they follow once a strong, typed compression-commitment package is fixed. The paper's contribution is therefore not a new dynamical law, but the formal isolation of that package and the explicit tracking of which assumptions support which conclusions. Dependency localization is the methodological core; the typed decomposition is what makes it precise.

The framework does not solve the measurement problem, and it does not establish that any physical system instantiates ICC. The candidate bridge to open quantum systems is articulated as two distinct conjectures (Conjectures 4.1 and 4.2): an exact realization (conditional on a resolution of the measurement problem) and an effective, for-all-practical-purposes realization in which the dominant outcome of μ_k is operationally definite through gate or measurement-error suppression, while environment-induced decoherence ensures the apparatus state is genuinely diagonal in the chosen basis. The framework's value is methodological: it provides a clear abstract target against which candidate realizations, weakened variants, and neighboring formal frameworks can be compared.

ICC is best understood as a compact structural lens. It decomposes irreversible commitment into separable contributors: boundary compression is the structural backbone, and the remaining axioms localize where in the pipeline information loss occurs. The framework also provides a diagnostic for which ingredients are sufficient for which irreversible behaviors. What the present work establishes is *independence* of the axiom subparts (each subpart is load-bearing for some result, witnessed by countermodels formalized in Lean 4 against `mathlib4`), not *minimality* of the axiom set as a whole. Whether the package admits a strictly smaller sufficient subset (after accounting for the typing dependencies made explicit in Proposition 3.24) is open. Two directions therefore remain: identifying the weakest sufficient conditions for each structural consequence (a proof-theoretic question), and determining whether any physically relevant regime satisfies them (an empirical one).

References

[1] Charles H. Bennett. "The Thermodynamics of Computation—a Review". In: *International Journal of Theoretical Physics* 21.12 (1982), pp. 905–940. DOI: 10.1007/BF02084158.

[2] E. Brian Davies and John T. Lewis. "An operational approach to quantum probability". In: *Communications in Mathematical Physics* 17.3 (1970), pp. 239–260. DOI: 10.1007/BF01647093.

[3] Tobias Fritz. "A synthetic approach to Markov kernels, conditional independence and theorems on sufficient statistics". In: *Advances in Mathematics* 370 (2020), p. 107239. arXiv: 1908.07021.

[4] Murray Gell-Mann and James B. Hartle. “Quantum mechanics in the light of quantum cosmology”. In: *Complexity, Entropy, and the Physics of Information*. Ed. by Wojciech H. Zurek. Vol. VIII. Santa Fe Institute Studies in the Sciences of Complexity. Addison-Wesley, 1990, pp. 425–458.

[5] Vittorio Gorini, Andrzej Kossakowski, and E. C. G. Sudarshan. “Completely Positive Dynamical Semigroups of N-Level Systems”. In: *Journal of Mathematical Physics* 17.5 (1976), pp. 821–825. DOI: 10.1063/1.522979.

[6] Robert B. Griffiths. “Consistent histories and the interpretation of quantum mechanics”. In: *Journal of Statistical Physics* 36.1–2 (1984), pp. 219–272. DOI: 10.1007/BF01015734.

[7] Robert B. Griffiths. “Choice of consistent family, and quantum incompatibility”. In: *Physical Review A* 57.3 (1998), pp. 1604–1618. DOI: 10.1103/PhysRevA.57.1604. arXiv: quant-ph/9708028.

[8] Austin Hulse, Benjamin Schumacher, and Michael D. Westmoreland. “Axiomatic Information Thermodynamics”. In: *Entropy* 20.4 (2018), p. 237. DOI: 10.3390/e20040237. arXiv: 1801.05015.

[9] Mathieu Huot and Sam Staton. “Quantum Channels as a Categorical Completion”. In: *Electronic Notes in Theoretical Computer Science* 341 (2019), pp. 31–54. DOI: 10.1016/j.entcs.2019.04.003. arXiv: 1904.09600.

[10] Bart Jacobs. “New directions in categorical logic, for classical, probabilistic and quantum logic”. In: *Logical Methods in Computer Science* 11.3 (2015), pp. 1–76. DOI: 10.2168/LMCS-11(3:24)2015.

[11] Rolf Landauer. “Irreversibility and heat generation in the computing process”. In: *IBM Journal of Research and Development* 5.3 (1961), pp. 183–191.

[12] Göran Lindblad. “On the Generators of Quantum Dynamical Semigroups”. In: *Communications in Mathematical Physics* 48.2 (1976), pp. 119–130. DOI: 10.1007/BF01608499.

[13] Roland Omnès. “Logical reformulation of quantum mechanics. I. Foundations”. In: *Journal of Statistical Physics* 53.3–4 (1988), pp. 893–932. DOI: 10.1007/BF01014230.

[14] Arthur J. Parzygnat and Francesco Buscemi. “Axioms for Retrodiction: Achieving Time-Reversal Symmetry with a Prior”. In: *Quantum* 7 (2023), p. 1013. DOI: 10.22331/q-2023-05-23-1013. arXiv: 2210.13531.

[15] Maximilian Schlosshauer. “Decoherence, the Measurement Problem, and Interpretations of Quantum Mechanics”. In: *Reviews of Modern Physics* 76.4 (2004), pp. 1267–1305. DOI: 10.1103/RevModPhys.76.1267.

[16] Wojciech H. Zurek. “Decoherence, Einselection, and the Quantum Origins of the Classical”. In: *Reviews of Modern Physics* 75.3 (2003), pp. 715–775. DOI: 10.1103/RevModPhys.75.715.

A Propositional Dependency Check

For completeness, the following note records a propositional sanity check of the proof skeleton.

The axiom subparts and structural premises of the framework were encoded as propositional variables: Axiom ICC-1 as ICC-1, Axiom ICC-2(a) and (b) as ICC-2a and ICC-2b, Axiom ICC-3(a), (b), and (c) as ICC-3a, ICC-3b, and ICC-3c, plus two structural premises (the typing of M from Definition 2.12 and the chain construction of Section 2.4). Theorem parts (i)–(iii) and Corollary 3.18 were encoded as propositional variables, with an implication from their cited premises. Corollary 3.19 and the axiom-drop propositions were not separately encoded; the latter are addressed implicitly by the drop-table of CHECK 3 below.

The SymPy satisfiability checker verifies three properties at the propositional level:

1. **CHECK 1.** The full system (all premises asserted, all dependency implications active) is satisfiable, and all result variables evaluate to true.

2. **CHECK 2.** No result variable can be false when the premises and dependency implications hold. This confirms that no result is vacuously encoded.

3. **CHECK 3.** Dropping any single premise allows at least one result variable to become false. This confirms that each premise is load-bearing at the dependency-graph level. The drop-table records, for each premise, which results become unentailed; for example, dropping Axiom ICC-3(c) breaks Theorem 3.1(iii), consistent with Proposition 3.21.

This check verifies propositional structure only. It does not verify the semantic correctness of the mathematical arguments, the naturalness of the axioms, or any physical claim. It is a sanity check on the proof skeleton, not a substitute for mathematical review. The encoding of Theorem 3.1(ii) is a single proposition entailed by the conjunction of all three obstacle-axioms; the obstacle-level layering of the proof is a feature of the prose argument, not of the propositional artifact. Independence of the three obstacles within the encoding is structural (the three axioms appear as atomic variables) rather than separately verified; the semantic independence of Axiom ICC-2(b) from Axiom ICC-2(a) together with Axiom ICC-3(b), established by counterexample in Remark 2.19, is outside the scope of the propositional check. Non-degeneracy (whether μ_k is a point mass) is a semantic property outside propositional reach; it is addressed in prose in Axiom ICC-3(c) and Proposition 3.21. Three classes of formal artifact have been independently developed in Lean 4 against `mathlib4`:

(i) *Algebraic check.* The pairwise distinctness of the four diagonal values used in Remark 2.19 is verified.

(ii) *Countermodel suite.* Four concrete countermodels (one for each axiom subpart Axiom ICC-1, Axiom ICC-2(a), Axiom ICC-3(b), and Axiom ICC-3(c)) are formalized: each exhibits a model in which the dropped subpart fails and the relevant theorem part (or pipeline obstacle) is demonstrably absent while the remaining axioms hold. The countermodel for Axiom ICC-3(b) corresponds to Proposition 3.23.

(iii) *Theorem-level formalization.* Theorem 3.1(i) is formalized as `T1i_noninjective` and `T1i_collision` using `Fintype.card_le_of_injective` from `mathlib4`; the per-step bound of Theorem 3.3(i) is formalized as `per_step_multiplicity_bound` via an averaging argument over fibers, with concrete instances verified for $\mathrm{Fin}\,4 \to \mathrm{Fin}\,2$ and $\mathrm{Fin}\,6 \to \mathrm{Fin}\,2$; and the dependency-localization content of Proposition 3.23 is formalized as `summary_prop_drop_icc3b`, exhibiting an injective Activation on a model with non-unitary flow Φ such that obstacle (b) fails while obstacle (a) and Theorem 3.1(i) survive. The chain-product form of Theorem 3.3(ii) is left to future work (the per-step composition step `fiber_comp_max_ge` that drives the induction is formalized).

Together, items (i)–(iii) provide model-theoretic and theorem-level witnesses that complement the propositional dependency check above: the propositional artifact tracks which axiom subparts are cited by which results; the Lean artifacts verify that the cited content holds under the standard mathematical interpretation.

B Load-Bearing Dependency Map

The table below summarizes the support structure of each formal result. It reproduces the dependency information from Remark 3.28 in expanded form for reference.

Result	Load-bearing premises
Theorem 3.1(i)	Axiom ICC-1 (finiteness of both sides), Axiom ICC-2(a) (strict boundary inequality). Pigeonhole applies because both sets are finite.
Theorem 3.1(ii)	Axiom ICC-2(b) (deterministic pipeline non-invertible), Axiom ICC-3(b) (Activation many-to-one), Axiom ICC-3(c) (Sample non-deterministic). Layered obstruction at distinct reverse-pipeline stages.
Theorem 3.1(iii)	Axiom ICC-2(b) (deterministic pipeline non-invertible), Axiom ICC-3(c) (Sample non-deterministic). Same distribution, possibly different outcome; neither application invertible.
Corollary 3.18	Theorem 3.1(ii) (irreversibility); Axiom ICC-3(a) (decomposition of σ), Definition 2.12 (typing of M: ρ_0 from A_{k-1}), Section 2.4 (recursive chain construction) (sequential dependency). Sequential composition of typed stages.
Corollary 3.19	Theorem 3.1(ii) (non-recoverability at each step, composed over the chain).
Proposition 3.20	Axiom ICC-1, Axiom ICC-2(a,b), Axiom ICC-3(b). Deterministic core without Sample; chain cannot continue.
Proposition 3.21	All except Axiom ICC-3(c). Deterministic Sample: loses (iii) entirely, loses one obstacle of (ii); remaining obstacles still prevent recovery.
Proposition 3.22	All except Axiom ICC-2(b). No pipeline non-invertibility axiom: loses obstacle (a) of (ii) and invertibility clause of (iii); Sample non-determinism and Activation many-to-one still block recovery.
Proposition 3.23	All except Axiom ICC-3(b). No Activation many-to-one axiom: loses obstacle (b) of (ii); pipeline non-invertibility (a) and Sample non-determinism (c) still block recovery.
Proposition 3.26	Proposition 3.20, Proposition 3.21, Proposition 3.22 (pairwise column comparison).

Each cited axiom subpart is load-bearing for at least one stated result.

C Worked Example: Two-State System

The following example is illustrative only. It is consistent with the framework definitions but does not constitute evidence for physical realization.

Consider a system with $|\Delta_0| = 2$ (two distinguishable prior states) and $|E_1| = 2$ (two distinguishable environment inputs). The input space has $|\Delta_0| \cdot |E_1| = 4$ distinguishable configurations. Suppose the outcome set has $|P_1| = 2$. Then $|P_1| < |\Delta_0| \cdot |E_1|$, satisfying Axiom ICC-2(a). By pigeonhole (Theorem 3.1(i)), at least two input configurations correspond to the same outcome class.

The pipeline proceeds as follows. Given input (A_0, e_1), the initial density matrix $\rho_0 = |A_0\rangle\langle A_0|$ is formed. The flow map $\Phi^{\Delta t}_{C(e_1)}$ evolves ρ_0 to $\rho(\Delta t) \in \mathrm{DM}(N)$. Activation extracts the diagonal: $\mu_1(i) = \langle i|\rho(\Delta t)|i\rangle$ for $i \in P_1$, producing a distribution on two outcomes. Sample draws a single committed outcome $A_1 \in P_1$ according to μ_1.

Recovery of (A_0, e_1) from A_1 faces the three obstacles of Theorem 3.1(ii). Concretely, suppose the evolved diagonal yields $\mu_1 = (0.7, 0.3)$ on $P_1 = \{0, 1\}$. Obstacle (a) (pipeline non-invertibility): suppose a different input configuration (A'_0, e'_1) also produces $\mu_1 = (0.7, 0.3)$ after evolution and activation. Then μ_1 does not determine which input produced it. Obstacle (b) (Activation many-to-one): both inputs may have evolved to density matrices with different off-diagonal entries but the same diagonal, so μ_1 does not determine $\rho(\Delta t)$. Obstacle (c) (Sample non-determinism): repeated application of M to the same input may produce different committed outcomes, since Sample may draw different outcomes from the same μ_1. One application of Sample returns $A_1 = 0$ (drawn with probability 0.7); a second application to the same input returns $A_1 = 1$ (drawn with probability 0.3). Both are valid committed outcomes drawn from the same distribution, illustrating Theorem 3.1(iii).

- *What is coarse-grained.* Four distinguishable input configurations (two prior states times two environment inputs) are compressed to two committed outcomes.
- *Why distinct histories collapse.* The two input configurations that map to the same μ_1 are indistinguishable at the committed level; the coarse-graining discards the difference.
- *What survives without Sample.* The truncated pipeline (Proposition 3.20) still produces μ_1, so boundary compression and deterministic non-invertibility hold, but no committed outcome is selected and the chain cannot continue.
- *Why this is not evidence for physics.* The example is consistent with the definitions but does not exhibit a physical system satisfying the axioms; the generator, the coarse-graining, and the commitment mechanism are stipulated, not derived from dynamics.

D Constructive Effective Realization in a Spin-Boson Measurement Model

This appendix exhibits a concrete physical realization satisfying the effective-realization conjecture Conjecture 4.2. The model identifies the structural ingredients required for a candidate realization (apparatus, dissipation, commuting pointer observable) and verifies that the checklist items Remark 4.3(1)–(3) and (5) are met for one step; a chain-extension proposition (Proposition D.2) gives the conditions under which the per-step result extends to an n-step chain. This is a minimal-conditions analysis, not a full characterization; other realizations may exist with different physical structure.

D.1 Model setup

We consider a system–apparatus–bath (S+A+B) measurement model in which the system S is measured by an apparatus A that subsequently dephases against an environment bath B. The setup is a textbook von-Neumann–style measurement: a system entangles with a pointer apparatus during a brief pre-measurement coupling, after which the apparatus dephases against a bath.

Components.

- System S: a single qubit with computational basis $\{|0\rangle_S, |1\rangle_S\}$.
- Apparatus A: a single pointer qubit with basis $\{|0\rangle_A, |1\rangle_A\}$, initial state $|0\rangle_A$.
- Bath B: a collection of N_{env} harmonic oscillator modes with creation/annihilation operators $a_\ell^\dagger, a_\ell$ and frequencies ω_ℓ.

Hamiltonian. The total Hamiltonian over the step interval Δt is

$$H_{\text{total}} = H_S + H_A + H_B + H_{\text{int}}^{S\text{-}A} + H_{\text{int}}^{A\text{-}B},$$

with system–apparatus pre-measurement coupling

$$H_{\text{int}}^{S\text{-}A} = g_{SA}\, \sigma_z^S \otimes \sigma_x^A$$

(switched on for a brief interval at the start of the step, with $g_{SA}\, \Delta t_{\text{pre}} = \pi/2$ giving a controlled-NOT-style entangling unitary), and apparatus–bath dephasing coupling

$$H_{\text{int}}^{A\text{-}B}(e_k) = \sigma_z^A \otimes \sum_{\ell=1}^{N_{\text{env}}} \lambda_\ell(e_k)\, (a_\ell + a_\ell^\dagger),$$

where the environment input $e_k \in E_k$ selects a set of bath couplings $\{\lambda_\ell(e_k)\}_\ell$.

Step dynamics. A single step proceeds in two phases:

(P1) *Pre-measurement entanglement.* The system–apparatus interaction acts for time Δt_{pre}, producing an entangling unitary U_{SA} that maps $|0\rangle_S|0\rangle_A \mapsto |0\rangle_S|0\rangle_A$ and $|1\rangle_S|0\rangle_A \mapsto |1\rangle_S|1\rangle_A$. An imperfect implementation introduces a per-step gate-error rate p, so that the post-pre-measurement reduced apparatus state for system input $|j\rangle_S$ is $\rho_A^{(j)} = (1-p)\,|j\rangle\langle j|_A + p\,|1{-}j\rangle\langle 1{-}j|_A + \mathcal{O}(\sqrt{p})$ off-diagonal.

(P2) *Bath dephasing.* The apparatus–bath coupling acts for the remainder of the step, generating Lindblad-form pure dephasing on the apparatus in the σ_z^A eigenbasis. Off-diagonal coherences of ρ_A decay as $\exp(-\Gamma(e_k)\Delta t)$, where $\Gamma(e_k)$ is the bath spectral density integrated against the step interval and depends on e_k through the bath couplings.

ICC primitives. The model maps onto the typed pipeline of Section 2 as follows:

- $\Delta_0 = \{|0\rangle_S, |1\rangle_S\}$ indexed by $j \in \{0, 1\}$, giving $|\Delta_0| = 2$.
- $E_1 = \{e_{\text{weak}}, e_{\text{strong}}\}$ a two-element set of bath couplings, giving $|E_1| = 2$.
- $P_1 = \{|0\rangle_A, |1\rangle_A\}$ the pointer basis of the apparatus, giving $|P_1| = 2$.
- $\Phi_{C(e_1)}^{\Delta t}$: the composite (P1) + (P2) CPTP map on DM(N) with $N = 2$.
- Activation: extraction of the diagonal of the reduced apparatus state $\rho_A(\Delta t)$ in the σ_z^A eigenbasis.
- Sample: single-outcome commitment to $A_1 \in \{|0\rangle_A, |1\rangle_A\}$ drawn according to μ_1.

D.2 Single-step checklist verification

We verify items (1), (2), (3), and (5) of Remark 4.3 for the model above at a single step. Item (4) is replaced by the phenomenological peakedness criterion of Conjecture 4.2, which is addressed in Appendix D.4.

Item (1): Finite boundaries. $|\Delta_0| = 2$, $|E_1| = 2$, both finite by construction.

Item (2): Motivated coarse-graining. $|P_1| = 2 < 4 = |\Delta_0| \cdot |E_1|$, satisfying Axiom ICC-2(a). The coarse-graining is the pointer-basis Born rule on the reduced apparatus state. The pointer basis is selected by the apparatus–bath coupling structure (see Appendix D.3) and is the natural measurement-outcome basis. The coarse-graining is not ad hoc.

Item (3): Basis selection. The pointer basis $\{|0\rangle_A, |1\rangle_A\}$ is the σ_z^A eigenbasis, which is decoherence-free under the apparatus–bath interaction (see Appendix D.3). This satisfies Definition 2.23(b) for the abstract framework.

Item (5): Pipeline non-invertibility. The four boundary inputs $(j, e) \in \Delta_0 \times E_1$ map under Activation $\circ\, \Phi^{\Delta t}_{C(e_1)}$ to two distinct distributions on P_1:

$$\begin{aligned}
(0, e_{\text{weak}}) &\mapsto \mu_1 = (1-p, p), \\
(0, e_{\text{strong}}) &\mapsto \mu_1 = (1-p, p), \\
(1, e_{\text{weak}}) &\mapsto \mu_1 = (p, 1-p), \\
(1, e_{\text{strong}}) &\mapsto \mu_1 = (p, 1-p).
\end{aligned}$$

The two configurations $(0, e_{\text{weak}})$ and $(0, e_{\text{strong}})$ produce identical μ_1; likewise for $j = 1$. The deterministic pipeline is therefore non-invertible on the boundary inputs, satisfying Axiom ICC-2(b). The load-bearing feature is that the bath coupling strength g does not enter the apparatus diagonal in this model: bath dephasing erases off-diagonal coherences uniformly across e_k, leaving the diagonal independent of the coupling strength.

D.3 Basis-rigidity resolution

The compatibility of Definition 2.23(b) (single fixed basis throughout the chain) with environment-induced basis selection turns on a commutator condition.

The pointer observable $O_A = \sigma_z^A$ commutes with the apparatus–bath interaction Hamiltonian:

$$[\sigma_z^A, H_{\text{int}}^{A\text{-}B}(e_k)] = \left[\sigma_z^A,\ \sigma_z^A \otimes \sum_\ell \lambda_\ell(e_k)(a_\ell + a_\ell^\dagger)\right] = 0,$$

since $[\sigma_z^A, \sigma_z^A] = 0$ and the bath operator $\sum_\ell \lambda_\ell(e_k)(a_\ell + a_\ell^\dagger)$ acts on a tensor factor disjoint from σ_z^A. This identity holds for every $e_k \in E_k$, regardless of the coupling values $\{\lambda_\ell(e_k)\}_\ell$. The pointer basis $\{|0\rangle_A, |1\rangle_A\}$ is therefore decoherence-free under the interaction at every environment input; the bath dephases off-diagonal entries of ρ_A in this basis but does not rotate the basis itself.

Remark D.1. This is a sidestep of the basis-rigidity tension noted in Section 5, not a resolution. The model is chosen so that the apparatus–bath coupling operator coincides with the pointer observable, ensuring $[O_A, H_{\text{int}}^{A\text{-}B}] = 0$. A more general apparatus–bath coupling (e.g., one with a σ_x^A component) would not have this property, and the pointer basis would depend on the coupling structure in a way that requires the per-step basis relaxation discussed in Section 5.

D.4 Peakedness

Conjecture 4.2 requires μ_k to be sharply peaked, with the dominant outcome selected operationally below a detection threshold ϵ. In the present model, peakedness has two distinct controls that should not be conflated.

Diagonal asymmetry: controlled by p. For system input $|j\rangle_S$, the post-step apparatus diagonal is $\mu_1 = (1-p, p)$ if $j = 0$ or $(p, 1-p)$ if $j = 1$, where p is the gate-error rate of the pre-measurement entangling unitary. The dominant outcome is selected operationally provided $p \ll \epsilon$. For high-fidelity solid-state qubit gates with $p \in [10^{-4}, 10^{-3}]$ and detection thresholds $\epsilon \sim 10^{-2}$, this condition is met by several orders of magnitude.

Coherence suppression: controlled by N_{env}. The apparatus–bath dephasing rate $\Gamma(e_k)$ scales with the number of effective bath modes N_{env} and the bath spectral density. For sufficient N_{env} and Δt, the off-diagonal coherences of ρ_A decay as $\exp(-\Gamma\,\Delta t)$ to numerical zero, ensuring that the apparatus state is genuinely diagonal at the end of the step.

Distinction. Diagonal asymmetry (peakedness of μ_1) is controlled by p and is independent of N_{env}. Bath size N_{env} controls the suppression of off-diagonal coherences, ensuring ρ_A is diagonal, but does not affect the diagonal entries. Conflating these two controls would suggest that more bath modes produce more peaked distributions; this is incorrect. The role of the bath in this model is to ensure that ρ_A is genuinely classical (diagonal) at the end of the step, not to peak the distribution over outcomes.

Scope of the two-control structure. The two-control decomposition relies on the commutator condition $[O_A, H_{\text{int}}^{A\text{-}B}(e)] = 0$ (Proposition D.2, hypothesis (i)) and on the strong-decoherence regime in which off-diagonal coherences are suppressed below numerical threshold. Numerical evidence supporting the scope of these assumptions was generated for this paper across four representative model classes: a single qubit with pure Lindblad dephasing, two qubits with combined local and correlated dephasing, a damped harmonic oscillator with thermal bath, and a finite-mode spin–boson model with exact unitary evolution. The two models within the commutator class (single-qubit and two-qubit dephasing) leave the diagonal entries unchanged under the dephasing parameter, confirming the two-control structure. The two models outside the commutator class (thermal-bath energy relaxation and the σ_x-coupled spin–boson model with finite bath) exhibit diagonal peaking driven by the system–bath coupling, in the directions predicted by the model structure and consistent with the conjecture's hedged scope. The hedge in Conjecture 4.2 is therefore not stylistic but empirically necessary.

D.5 Chain extension

The single-step verification of Appendix D.2 extends to an n-step chain under two hypotheses.

Proposition D.2 (Chain extension)**.** *Let the S+A+B model of Appendix D.1 be applied iteratively for n steps, with environment inputs $e_1, \ldots, e_n$ at each step and a constant per-step error rate p. If*

(i) *the pointer observable satisfies $[O_A, H_{\text{int}}^{A\text{-}B}(e)] = 0$ for every $e \in \bigcup_{k=1}^{n} E_k$, and*

(ii) *the per-step error rate p and chain length n satisfy $np \ll \epsilon$ for detection threshold ϵ,*

then the n-step chain satisfies Definition 2.23(b) and the conclusions of Corollary 3.19 hold under Conjecture 4.2.

Proof. Hypothesis (i) ensures that the pointer basis is e_k-invariant at every step, satisfying Definition 2.23(b). The single-step checklist verification of Appendix D.2 therefore applies unchanged at each step.

Hypothesis (ii) controls operational definiteness across the chain. Let X_k denote the event that step k exits the operationally definite regime (the dominant outcome is not selected at confidence above $1 - \epsilon$). Each step has $\Pr(X_k) \leq p$ by the single-step analysis of Appendix D.4. By the union bound (which holds regardless of independence between steps), the probability that any step in an n-step chain exits the operationally definite regime is at most np. Under hypothesis (ii), $np \ll \epsilon$, so this probability is below the detection threshold for the chain as a whole.

The non-recoverability of Corollary 3.19 under Conjecture 4.2 then follows from the per-step verification together with the cumulative bound. □

Remark D.3. Hypothesis (i) is a structural condition on the model: the apparatus–bath coupling must commute with the pointer observable for every environment input encountered in the chain. In the model of Appendix D.1, this is satisfied by construction since $H_{\text{int}}^{A\text{-}B}(e_k) \propto \sigma_z^A$ for every e_k. A model in which different environments select different pointer bases would require the per-step basis relaxation discussed in Section 5. Hypothesis (ii) is a quantitative condition on chain length and per-step error rate, giving a finite-n bound on operational definiteness; it does not require $p = 0$.

D.6 What is not claimed

This appendix establishes the following: the S+A+B model of Appendix D.1 satisfies the structural axioms ICC-1 through ICC-3 at each step and meets the realization checklist of Remark 4.3(1)–(3) and (5); the chain-extension proposition (Proposition D.2) gives explicit conditions under which the per-step result extends to an n-step chain. The following are explicitly not claimed.

- *Sample mechanism.* The model does not specify a mechanism for the single-outcome selection performed by Sample. The measurement problem remains open; this realization is for Conjecture 4.2, not Conjecture 4.1.
- *Generic open systems.* The model relies on the apparatus–bath coupling commuting with the pointer observable ($[O_A, H_{\text{int}}^{A\text{-}B}] = 0$). Generic open systems with non-commuting couplings would not satisfy Definition 2.23(b) directly; the per-step basis relaxation of Section 5 would be required.
- *Quantum Darwinism.* We do not claim that bath fragments redundantly encode the pointer-basis outcome. The role of the bath in this analysis is to suppress off-diagonal coherences, not to peak the diagonal.
- *Specific experimental fitting.* Numerical values for g_{SA}, $\lambda_\ell(e_k)$, ω_ℓ, and p are not fitted to any specific experimental platform. The model is a minimal-conditions analysis; realization in a particular physical setup requires separate parameter justification.
- *Characterization.* This is not a characterization of the realizations satisfying Conjecture 4.2. Other models with different physical structure may also satisfy the conjecture; this appendix verifies sufficiency for one model class, not necessity of the chosen ingredients.

The appendix's contribution is to demonstrate that Conjecture 4.2 has at least one explicit, concrete realization satisfying the relevant checklist items, with the structural ingredients (apparatus, dissipation, commuting pointer observable) identified clearly and the chain-extension conditions stated formally.