Causal Propagation Requires Fact-Commitment

Keith Taylor
VERSF Theoretical Physics Program

For the General Reader

Physics uses the word "cause" constantly. A photon causes a detector to click. A force causes an acceleration. A signal causes a downstream event. But what does "cause" actually require, at the level of fundamental physics?

The standard story is that physics propagates amplitudes — wave functions, field values, probability currents — and causation is something we read off from that propagation after the fact. This paper argues that story is incomplete in a precise sense. Amplitude propagation alone is not enough to supply physical causation. Something more is required: a *committed fact* must travel from cause to effect.

The argument is not against quantum mechanics or reversible microphysics. It accepts that amplitudes evolve unitarily. What it claims is that unitary evolution, by itself, cannot make "A caused B" a physically meaningful, observer-comparable statement — because in a purely reversible world, no outcomes are ever settled, and unsettled outcomes cannot serve as causal antecedents. Causation requires a ledger. Ledgers require irreversible entries. Irreversible entries are exactly what the commitment framework calls facts.

Central claim: Physical causation is not amplitude propagation. It is the propagation of *committed distinguishability* — stable, irreversible records that persist long enough to constrain downstream outcomes. The causal graph of physics is built from commitment events and the constraints they carry forward, not from the unitary evolution that underlies them.

This has a sharp practical consequence. When a referee objects that "light propagates as amplitudes, so your commitment-layer budget argument doesn't touch the speed of light," the correct response is: that is precisely the point. If causation is amplitude propagation, you do not yet have a causal theory — you have a wave equation. A causal theory must specify how committed facts propagate and close. That layer is what this framework is about.

Table of Contents

Abstract

We analyze the minimal structure required for observer-comparable causal ancestry in a finite-resource physics. *Scope: results address pre-geometric frameworks where causal order is not presupposed as background structure but must be operationally grounded in records; fixed-background QFT is outside this scope by construction.* We distinguish (i) weak dynamical influence (D1a), representable in amplitude-based theories via commutators and microcausality, from (ii) strong causal ancestry on recorded outcomes (D1b), which is the form of causation used in empirical science — counterfactual dependence between stable records comparable across observers. Under finite distinguishability (FD) and irreversible commitment (IC) as operational admissibility constraints, we prove:

Theorem A (Record Necessity). If $A \prec B$ holds in the strong sense D1b, then there exists a stable record R whose value depends on admissible interventions at A and remains operationally distinguishable ($TV > \varepsilon$) into the causal domain of B. D1b causal propagation is necessarily record-carrying propagation at the admissible layer.

Theorem B (Amplitude Trilemma — No-Go for D1b). No purely reversible amplitude ontology can define observer-comparable D1b causal order without either (i) dictionary/subspace dependence, (ii) a non-injective operational closure map whose inverse is not admissible (thereby introducing operational commitment), or (iii) abandoning settled outcomes. (Pre-geometric scope: causal order is not presupposed as background structure.)

Causal Ledger Theorem. Measurement-localized commitment and microcausality alone are insufficient to ground apparatus-independent causal ancestry in pre-geometric settings. Any observer-comparable causal order must be instantiated on committed records — the causal ledger. Three lemmas establish the components: (1) fact-causal relations require stable record carriers sustaining $TV > \varepsilon$ in transit; (2) measurement-localized commitment provides no structural guarantee of apparatus-invariant fact partitions across all admissible observers; (3) any commitment map R is either injective (dictionary-dependent) or non-injective with no admissible recovery (operational IC) — microcausality does not resolve the choice.

Corollary. The causal graph of an admissible physics is built from commitment events and the constraints they propagate forward. Causal propagation = propagation of committed distinguishability.

These results do not deny unitary dynamics. They specify what must supplement reversible dynamics before D1b causal claims become physically meaningful. A dimensional consequence (Section 7) follows conditionally on Assumptions A6′ (no free creation of independent committed constraints) and A7 (spatial locality of commitment events): primitive causal ledger propagation fails admissibility for $d \geq 4$ spatial dimensions, constraining the effective dimension of the commitment substrate.

Referee Clarification Note on Ontological Layers

The present framework distinguishes explicitly between two structural layers:

Layer 1 — Reversible amplitude dynamics. Described by unitary evolution on a Hilbert space (or equivalent reversible process structure), from which standard continuum field equations (Maxwell, Schrödinger, etc.) are recovered in the infrared. This layer is compatible with Lorentz invariance and conventional quantum field theory. No claim is made against it.

Layer 2 — Fact-commitment (causal ledger) dynamics. The operational layer at which alternatives become non-recombinable within the admissible process class $\mathcal{O}(R, \varepsilon)$ and stable records are formed. This layer is constrained by finite distinguishability (FD) and irreversible commitment (IC) and is the proper domain of the dimensional admissibility arguments.

The non-admissibility theorem for primitive causal ledger propagation (Section 7) applies strictly to Layer 2. No claim is made that Maxwell electrodynamics is incorrect as an effective theory. The claim is that the causal substrate responsible for objective record formation must satisfy admissibility constraints that restrict its effective dimensionality.

Objections that invoke microcausality or unitary QFT as sufficient accounts of causal structure conflate operator-level signaling constraints (Layer 1, D1a) with observer-comparable fact-causal order (Layer 2, D1b). The Causal Ledger Theorem shows that microcausality alone cannot ground apparatus-independent causal ancestry without an additional commitment map. Once such a map exists, it is subject to the FD/IC admissibility constraints analyzed here.

Lorentz invariance at Layer 1 and a two-dimensional commitment substrate at Layer 2 are not logically incompatible (Appendix L). Any residual symmetry violation from the discrete microgeometry of Layer 2 is symmetry-controlled and renormalization-group suppressed. A companion phenomenology paper proposes a 6Ω anisotropy signature as a falsifiable probe of that substrate structure.

---

1. Introduction

Causation is the backbone of physical reasoning. Every experimental claim takes the form "preparing system A in state s caused outcome distribution p at detector B." Every conservation law is a constraint on what causes what. Every notion of signal propagation is a causal claim.

Yet the fundamental dynamical laws of physics are reversible. Quantum mechanics evolves wave functions unitarily. Classical mechanics evolves phase-space points by diffeomorphisms. In neither case does the fundamental law contain an asymmetric causal arrow. The causal asymmetry we observe — causes precede effects, signals travel forward in time, interventions have downstream but not upstream consequences — is not a feature of the evolution law itself.

The standard response is that causation is "emergent": it arises from thermodynamic asymmetry, from the low-entropy initial condition, from decoherence, from coarse-graining. These responses are not wrong, but they leave a structural gap. They explain why the world *looks* causal to observers embedded in it. They do not explain what physical structure grounds the asymmetric relation $A \prec B$ as a fact about the world rather than a perspective adopted by observers.

This paper closes that gap by identifying the minimal physical structure causation requires: a *stable record* carrying information from the causal domain of A to the causal domain of B. We prove this is necessary (Theorem A) and that purely reversible amplitude evolution cannot supply it (Theorem B). We further prove that D1b reusability is not an added assumption but is forced by the minimal conditions for empirical lawhood under admissible record exchange (Theorem C). The commitment framework — the operational picture in which irreversible facts arise at a codimension-1 interface Σ in the state space of a reversible substrate — is precisely the structure that fills this gap.

The argument does not require a specific physical theory. It operates at the level of what any admissible account of causation must contain, given finite distinguishability and the requirement that causal claims be observer-comparable. This makes the results robust: they apply to classical mechanics, quantum mechanics, and any future framework that aspires to make causal claims.

Foundational assumption (stated explicitly). The results of this paper are conditional on one methodological commitment: that physics is an operational enterprise — that physical claims must be underwriteable by intersubjective record comparison under finite resources. This is not derived here; it is the stance from which D2 (admissibility) is defined. Readers who regard causal structure as an ontological primitive independent of any operational grounding are outside the scope of this paper. The paper's target audience is frameworks that seek to derive causal and spacetime structure from physically checkable relations — pre-geometric programmes in the sense made precise in the Scope note at Theorem B.

Why pre-geometric scope addresses a genuine problem. Several research programmes in quantum gravity — including canonical approaches with diffeomorphism constraints, causal-set-like ontologies, and some formulations of holographic reconstruction — face a version of the derivation problem: they require causal order to emerge from or be reconstructed from more primitive structure, rather than being stipulated as background. The claim is not that all such approaches are committed to the pre-geometric stance, nor that the stance is compulsory — the holographic case in particular retains background structure on the boundary, and canonical gravity admits multiple interpretations. The claim is more modest: the derivation problem is real for a non-trivial class of serious frameworks, and the results here directly constrain any such framework by identifying what operational structure observer-comparable causal order minimally requires.

---

2. Definitions

D1 (Physical Causation — Two Senses). Two distinct notions of causal influence must be distinguished; the results of this paper apply specifically to the stronger:

D1a (Weak causal influence): A dynamical signaling relation — a region A has weak causal influence on region B if unitary operations at A can change expectation values of observables at B. This is the sense operative in quantum field theory microcausality, Lieb–Robinson bounds, and commutator-based signaling: the amplitude dynamics itself can represent this relation without invoking any commitment.

D1b (Reserved label — derived notion): We reserve the label "D1b causal ancestry" for the reusability property that will be derived as Theorem C from the requirements of empirical lawhood under admissible record exchange. Until Theorem C is established, we use only D1a. The independence clause and formal readout conditions below anticipate the structure that Theorem C will derive.

Independence clause (formal). Independence is defined relative to the admissible process class $\mathcal{O}(R, \varepsilon)$: a readout Π_2 at t_2 is *independent of* Π_1 if it is implementable by an operation in $\mathcal{O}(R, \varepsilon)$ that is localized to B's domain at t_2, using only admissible ancillae prepared after t_1, without any nonlocal joint unitary across the full A–B system. The resource bound R on $\mathcal{O}(R, \varepsilon)$ fixes what counts as "local": operations exceeding R are not admissible, so "global control" is operationally defined as any operation outside $\mathcal{O}(R, \varepsilon)$. This excludes quantum-eraser scenarios requiring non-admissible joint control across A–B.

Relation to Woodward's interventionism. Woodward's account treats counterfactual dependence under interventions as sufficient for causation — $A \prec B$ holds if interventions on A change outcome distributions at B. OLC-2 and OLC-3 below are the specifically physical closure requirements: intersubjective testability under finite admissible resources and representation invariance under record exchange. Theorem C shows that these closure requirements, together with Woodward's base counterfactual notion, force reusability. The result is not a philosophical strengthening of Woodward; it is what his criterion becomes when subjected to the minimal constraints of empirical physics.

Definition C.0 (Operational Law Closure). A theory is *law-admissible* if it satisfies:

(OLC-1) Finite operationality. All physically meaningful claims are statements about outcome distributions of processes in $\mathcal{O}(R, \varepsilon)$.

(OLC-2) Replicable testability. For any claimed intervention–outcome relation do(A=s) ↦ P_B(· | do(A=s)), there exist two admissible experimental instances (not necessarily simultaneous) producing records that can be exchanged and compared without requiring any non-admissible global control of the joint system.

(OLC-3) Representation invariance under record exchange. If two admissible observers perform admissible interventions and readouts and then exchange records, the truth of the law statement

does not depend on arbitrary representation choices (basis, subsystem dictionary) made by either observer.

Theorem C (Reusability is Forced by Empirical Law Closure). Assume a theory satisfies OLC-1–OLC-3 and admits at least one nontrivial intervention effect:

$\exists\, s \neq s'$ such that $\mathrm{Dist}(P\_B(\cdot|do(A=s)), P\_B(\cdot|do(A=s'))) > \varepsilon$.

Then any causal relation "A ≺ B" that is intended to be lawlike — assertable as an observer-comparable physical claim — must satisfy: there exist two admissible readouts at B, executed at times $t_1 < t_2$, each localized to B and each implementable within $\mathcal{O}(R, \varepsilon)$, such that both readouts witness the intervention-dependence above ε. This derived property is what we call *D1b causal ancestry*.

Proof.

By OLC-2, a law statement must be confirmable by independent admissible experimental instances whose records can be compared without non-admissible global control. Suppose the influence of do(A=s) on B exists only as reversible coherence requiring global recombination to extract. Then any later attempt to confirm the claim either (i) requires a non-admissible global operation — violating OLC-2 — or (ii) depends on an apparatus or dictionary choice not fixed by the propagation substrate — violating OLC-3. Either way the claim is not law-admissible.

Therefore, for the intervention-dependence to be law-admissible, it must persist in a form accessible to a localized admissible readout at B at t_1, and remain accessible to a second independent localized admissible readout at B at $t_2 > t_1$, without requiring joint control of the full A–B system. Formally: there exist $\Pi_1, \Pi_2 \in \mathcal{O}(R, \varepsilon)$, localized to B, such that $\mathrm{TV}(P(\Pi_i \mid do(A=s)), P(\Pi_i \mid do(A=s'))) > \varepsilon$ for i = 1, 2. This is D1b. ■

Interpretation. D1b is not a philosophical strengthening of Woodward; it is the minimal condition for an interventionist causal claim to be a law statement rather than a one-shot coherence extraction that fails closure under admissible record exchange.

Scope note. Theorems A and B are claims about D1b. Amplitude-only theories can represent D1a (dynamical influence) via commutators and microcausality. They cannot represent D1b (law-admissible causal ancestry) without a commitment layer. When a referee argues that "causation in QFT is defined by microcausality without records," they are invoking D1a. The paper does not contest this. It establishes that D1b — as derived in Theorem C — requires a carrier that persists independently accessible to multiple readouts, and Theorem A shows this forces a stable record.

D2 (Admissibility). An account of physical causation is *admissible* if it satisfies:

- *Finite Distinguishability (FD):* causal claims can only invoke distinctions that are operationally resolvable at finite resources R and tolerance ε.
- *Irreversible Commitment (IC):* some outcome events — those that serve as causal antecedents — must be irrecoverable within the admissible process class: there is no

admissible operation in $\mathcal{O}$(R, ε) that restores the pre-event microstate. This upgrades a robust record (D3) to a committed record (D3C).

- *Observer-Comparability (OC):* different observers, performing experiments on A and B independently, must be able to compare their results and agree on whether A ≺ B held in a given instance.

Process class $\mathcal{O}$(R, ε). The admissible process class $\mathcal{O}$(R, ε) consists of all physically realizable operations implementable within resource bound R and discrimination tolerance ε. In the quantum setting, every operation in $\mathcal{O}$(R, ε) is a completely positive trace-preserving (CPTP) map — the standard mathematical description of any physically realizable quantum process, including unitary evolution, measurement, and open-system dynamics. In the classical limit, CPTP maps reduce to stochastic maps. This is not an additional assumption: CPTP is the general structure of any linear, trace-preserving, positivity-preserving operation on a density matrix, which any physically realizable process must satisfy. The contractivity of trace distance (and TV distance in the classical case) under CPTP maps — the key tool in Lemma A.1 — therefore applies to all operations in $\mathcal{O}$(R, ε) by construction.

Remark on observer-comparability. This condition is not a metaphysical axiom. It is a structural requirement for representation-invariant law: any account of causation whose causal relations depend on arbitrary observer representation choices cannot support invariant law statements. The formal statement is A0′ (Operational Closure Under Record Exchange) in the Assumptions Ledger, where three structural justifications are given and a Proposition connects A0′ violation directly to Case 1 (Dictionary/Subspace Dependence) of Theorem B.

D3 (Robust Record / R-Record). A *robust record* is a physical variable that has taken a definite value and retains it over an operationally relevant horizon T with error ≤ ε. Formally, a record is a readout map $\Pi : \mathcal{X} \rightarrow \{0, 1\}$ satisfying the metastability condition:

$$P(\ \Pi(\Phi_t(x)) = \Pi(x)\ \forall\ t \in [0, T]\ ) \geq 1 - \varepsilon$$

A robust record is stable and independently readable but does not yet require irrecoverability. Theorem A proves that D1b reusable influence forces a robust record. The IC condition in D2 then upgrades it.

D3C (Committed Record / C-Record). A *committed record* is a robust record (D3) whose pre-record microstate is not recoverable within $\mathcal{O}$(R, ε): no admissible operation inverts the commitment map R_M within the operational domain. Committed records are the outputs of commitment events in the full D2 sense.

D3A (Commitment Map). We denote by R_M : S → $\mathscr{R}$ any map from the amplitude-level state description S to a discrete record/event space $\mathscr{R}$ induced by a measurement coupling M (pointer observable, coupling Hamiltonian, environment structure). A commitment map is *operationally committing* if and only if: (i) R_M is non-injective (many amplitude microstates map to the same record), and (ii) the inverse R_M^{-1} is not in the admissible process class $\mathcal{O}$(R, ε) — no admissible operation can restore the pre-map microstate from the record. Non-injectivity alone (coarse-graining) is not commitment; irrecoverability of the inverse within $\mathcal{O}$(R, ε) is what makes it so.

D4 (Amplitude-Based Propagation). An account of propagation is *amplitude-based* if the fundamental ontology consists solely of reversibly evolving amplitudes (wave functions, field configurations, probability currents) and no irreversible commitments occur at the fundamental level. In this picture, anything that resembles "cause/effect" is an emergent, observer-relative description that arises only after a measurement or post-selection layer is introduced from outside the fundamental dynamics.

---

3. Main Results

Theorem A (Reusable Causal Influence Requires a Stable Record Carrier)

Let $A \prec B$ be a law-admissible causal relation satisfying D1b (as derived in Theorem C — reusable influence under OLC-1–OLC-3). Then there exists at least one robust record R (in the sense of D3) that:

(i) Carries information about the intervention at A — the value of R depends on which admissible intervention was applied at A.

(ii) Persists from A's causal domain into B's — R remains stable and independently readable throughout the transit.

(iii) Is accessible at finite resources — R is distinguishable by any admissible observer bounded by (R, ε).

When D2/IC also holds, R is a committed record (D3C). The theorem is non-circular: "record" does not appear in D1b; it is derived from the reusability condition.

Proof.

Step 1: Reusable influence requires a physical carrier. D1b requires two independent readouts Π_1, Π_2 — applied at separate times without joint global control — both to detect the A-intervention above ε. For the influence of A's intervention to be independently readable twice, there must be some physical degrees of freedom between A and B whose state encodes the intervention and can be accessed by each readout independently. If no such carrier exists, the second readout cannot access what the first did, and reusability fails.

Step 2: The carrier must persist above the distinguishability threshold. Between the two readout times, the carrier must maintain the A-intervention-dependent influence. If it is scrambled below ε before the second readout — by decoherence, thermalization, or recombination requiring global control — then the second readout cannot independently confirm the first, and D1b fails. Formally, the total variation distance must remain above ε throughout:

$$\mathrm{TV}\big(P(\Pi_i \mid do(A = s)), P(\Pi_i \mid do(A = s')) \big) > \varepsilon \text{ for } i = 1, 2 \text{ at their respective times}$$

Step 3: By Lemma A.1, persistent TV separation above ε over horizon T constitutes a robust record (D3). A carrier sustaining TV > ε across two independent readouts separated in time, without global control, is precisely a variable whose value constrains outcome distributions independently at resolution ε — which is D3. This is not assumed; it follows from what D1b's reusability condition requires of any physical carrier.

Lemma A.1 (ε-Ancestry Requires Channel-Stable Distinguishability). *Let interventions s, s′ at A induce two carrier state ensembles at time t_1, and let $\mathscr{E}$_{$t_1 \to t_2$} be the admissible propagation map for the carrier C from t_1 to t_2. If $A \prec B$ holds in the sense of D1b — i.e., two independently applied readouts $\Pi_1, \Pi_2 \in \mathcal{O}(R, \varepsilon)$ both detect the intervention above ε at their respective times — then the underlying carrier distributions must remain separated above ε under $\mathscr{E}$_{$t_1 \to t_2$}. This persistent separation constitutes a robust record (D3) in C.*

Proof. The key tool is contractivity: for any classical stochastic map (or CPTP quantum channel) $\mathscr{E}$,

$$\mathrm{TV}(\ \mathscr{E}(\rho\_s),\ \mathscr{E}(\rho\_\{s'\})\ ) \le \mathrm{TV}(\ \rho\_s,\ \rho\_\{s'\}\ )$$

(trace distance contractivity in the quantum case). This is the data processing inequality: no admissible local operation on the carrier can increase the distinguishability between the two intervention-conditioned states.

Suppose TV(ρ_s(t), ρ_{s′}(t)) falls below ε at some intermediate time $t \in (t_1, t_2)$. Then by contractivity, TV remains below ε for all subsequent times under any admissible map — no choice of predicate at t_2 can restore separation above ε once it has been lost. The second readout Π_2, applied at t_2, cannot detect above ε regardless of which separating observable it uses. D1b then fails.

Therefore, if D1b holds — if both readouts independently detect above ε — the carrier distributions must maintain TV > ε throughout $[t_1, t_2]$ under $\mathscr{E}$_{$t_1 \to t_2$}. This is not a statement about which predicate is used; it is a monotonicity condition on the carrier's distinguishability under the propagation channel. Persistent TV separation above ε further implies the existence of at least one binary readout observable Π on the carrier whose outcome distributions differ by > ε at both t_1 and t_2: this follows from the variational characterization of total variation distance as a supremum over measurable partitions, TV(P, Q) = sup_Π |P(Π=1) − Q(Π=1)|. That observable Π, stably readable from the carrier's local degrees of freedom throughout $[t_1, t_2]$, satisfies D3: the metastability condition is the requirement that such a Π exists and remains operationally accessible at resolution ε over the horizon $T = t_2 - t_1$. ■

Note. This proof does not require the optimal separating predicate to be the same at t_1 and t_2 — the predicate may evolve, and the second readout may apply a different observable than the first. What the data processing inequality establishes is that the *underlying distinguishability of the carrier states* cannot be recovered by predicate choice once lost. D1b therefore forces the carrier's distinguishability to be maintained by the propagation dynamics itself, which is the content of D3.

Step 4: Observer-comparability (D2/OC) requires the robust record to be independently verifiable. By D2/OC, the causal claim A ≺ B must be comparable across observers. This requires that the record R can be read by different observers and yield consistent results — which is what D3 provides (metastability, finite-resource readout). When D2/IC also holds — when the pre-record microstate is irrecoverable in $\mathcal{O}$(R, ε) — the robust record is upgraded to a committed record (D3C).

Therefore, D1b requires a robust record R satisfying (i)–(iii), and D2/IC upgrades it to a committed record. ■

Theorem B (Amplitude Trilemma — No-Go for D1b in Amplitude-Only Ontologies)

Assume fundamental dynamics is purely reversible (D4) and that causal order is not presupposed as background structure (pre-geometric scope — see Scope Condition below). Any attempt to define observer-comparable causal ancestry on recorded outcomes (D1b) within a purely amplitude-based ontology falls into exactly one of:

(Case 1 — Dictionary/Subspace Dependence): The event partition (hence causal order) depends on a representation choice — violating Observer Comparability (D2). Different observers with different but equally valid representations disagree on which events are causally ordered.

(Case 2 — Non-Injective Operational Closure): The account introduces a commitment map R_M (D3A) that is non-injective and whose inverse is not in $\mathcal{O}$(R, ε) — thereby introducing operational irreversible commitment. The ontology is no longer amplitude-only.

(Case 3 — No Settled Outcomes): No operationally non-recombinable records exist. D1b causal ancestry is undefined; the practice of physics as record comparison is groundless.

Therefore, a commitment layer is required for D1b causal order. This theorem does not contest D1a dynamical influence; amplitude-only theories handle D1a via commutators and microcausality. It targets D1b causal ancestry on records, which is the form used in empirical science.

Scope (pre-geometric setting). This paper addresses frameworks in which causal order is not presupposed as background structure but is required to be grounded in physically checkable relations between events (records) under admissibility D2. In background-spacetime formulations of QFT, causal order is stipulated by the manifold structure; our results do not contest the internal consistency or utility of that stipulation. The results here apply to programmes that seek to derive causal order — and spacetime structure — rather than assume it.

Proof.

Step 1 (Reversible dynamics preserve distinguishability). Under purely reversible evolution, the state transformation U(t) is injective and distinguishability-preserving: no information is lost at the fundamental layer. No outcome is *intrinsically* settled — non-recombinability arises only relative to the admissible process class $\mathcal{O}$(R, ε), not from the dynamics themselves. The practical impossibility of reversal (Loschmidt-type objections, resource constraints) is handled precisely by the IC condition in Case 2: what matters is whether the inverse is in $\mathcal{O}$(R, ε), not whether it is cosmically difficult. Step 1 establishes only that at the fundamental layer, prior to any resource-bounded commitment condition, no record is permanently settled — non-recombinability is an operational property of $\mathcal{O}$, not an intrinsic property of the dynamics.

Step 2 (Causal claims require stable records). By Theorem A, any physically meaningful A ≺ B requires intermediate degrees of freedom that (i) depend on A, (ii) persist long enough to constrain B, and (iii) are robust to perturbations at finite resources. These are stable records in the sense of D3. In pure amplitude evolution (Step 1), no such stable records exist at the fundamental layer — any apparent record can be unwound by the invertible dynamics. Therefore Operational Causation (D1) cannot be satisfied within the amplitude-only ontology without additional structure.

Step 3 (Case 1 — Injective definitions become representation-dependent). Suppose one refuses non-injective reduction and attempts to define causal events purely within reversible amplitude evolution. Then event structure must be defined via chosen observables, subsystem decompositions, or coarse-graining conventions. Without irreversible commitment fixing equivalence classes of histories, these choices are not canonical: the same physical evolution will be partitioned into different "causal events" by different choices of basis, subspace, or subsystem decomposition. Causal relations become dictionary-dependent. Two observers using different but equally valid representations may disagree on whether A ≺ B holds — violating Observer Comparability (D2).

A natural objection is that einselection or decoherence physics canonically selects a preferred pointer basis, making the event partition unique. The worked example in Lemma 2 directly answers this: M_z and M_x are both admissible measurement couplings for the same incoming state — both satisfy the finite-resource, stable-pointer conditions of $\mathcal{O}$(R, ε) — yet they induce incompatible fact partitions. Einselection stabilizes a particular pointer basis *given a specific apparatus-environment coupling*, but it does not structurally exclude other admissible couplings. Different experimenters using different admissible apparatuses (with different environment structures) may still disagree on which event occurred. The event partition remains apparatus-dependent until the propagation substrate itself fixes equivalence classes before any coupling acts. □ Case 1.

Step 4 (Case 2 — Non-injective definitions introduce operational commitment). Suppose instead one defines causal influence using a reduction map R : S → S′ that produces definite outcome records. For R to produce distinct, stable records from different input states, it must be non-injective: many microstates must map to the same record state.

Non-injectivity alone is not sufficient for commitment — a referee may correctly note that any coarse-graining is non-injective without being irreversible. The additional required condition is *no admissible recovery*: there is no admissible operation (bounded by resources R and tolerance ε) that inverts R and restores the pre-map microstate. Formally: R constitutes operational commitment relative to the admissible process class $\mathcal{O}$(R, ε) if and only if R is non-injective *and* the inverse R^{-1} is not in $\mathcal{O}$(R, ε).

This combined condition is precisely IC (Irreversible Commitment) as defined in D2: alternatives become non-recombinable within the operational domain. An account that uses such an R has therefore introduced a commitment layer — not merely a coarse-graining. The ontology is no longer amplitude-only. □ Case 2.

Step 5 (Case 3 — Refusal of both options destroys operational meaning). If one refuses both dictionary-dependent event structure (Case 1) and non-injective operational closure (Case 2), then no stable facts exist in the ontology. No record is ever settled in the required sense.

A referee may object: "Reversible dynamics can preserve logical bits indefinitely — topological codes and integrable systems protect information without any commitment." This is true but does not rescue causal ancestry. A reversible protected memory is not a committed fact: it is a *recombinability-preserving* encoding in which the stored information can be undone, coherently recombined, and interfered. Commitment requires not merely persistence of information but *irrecoverability of alternatives within the admissible domain* — non-recombinability under $\mathcal{O}$(R, ε). Reversible protection preserves bits; commitment closes the history. They are structurally distinct. An account that offers only reversible protected memory cannot ground D1b causal ancestry, because the "record" can always in principle be unwound and the counterfactual undone.

Therefore, in an ontology without operational commitment, no causal ancestry relation can be stably compared between observers. □ Case 3.

Step 6 (Trilemma is exhaustive). Cases 1–3 cover all logically possible responses to the requirement of admissible causal claims in a purely reversible ontology. Any proposed account must either: define events via representation choice (Case 1), use a non-injective record-forming map (Case 2), or forgo settled outcomes entirely (Case 3). No fourth option exists. In each case, the amplitude-only ontology fails to supply admissible causal propagation. The commitment layer is therefore not optional — it is forced by the joint requirements of Operational Causation and Observer Comparability. ■

The Causal Ledger Theorem

The Amplitude Trilemma (Theorem B) establishes that some commitment layer is required for admissible causal order. The Causal Ledger Theorem establishes the stronger and more precise claim: any observer-comparable causal order must be grounded in a commitment layer at the *structural propagation level. Measurement-localized commitment and microcausality alone are each insufficient. The three lemmas below establish the components; the theorem collects them.*

Definitions for the Causal Ledger Theorem

Definition CLT.1 (Microcausality). A theory satisfies *microcausality* if for spacelike-separated regions A and B, $[\mathcal{O}_A, \mathcal{O}_B] = 0$, so operations in A cannot alter outcome statistics in B unless mediated by timelike propagation. Microcausality constrains operator influence (D1a). It does not by itself define settled outcomes.

Definition CLT.2 (Fact-Causal Order). A *fact-causal order* is a partial order on committed records E_i such that $E_i \prec E_j$ means E_i is a necessary causal antecedent of E_j in the counterfactual sense: admissible interventions on E_i alter the distribution of outcomes at E_j. Fact-causal order must satisfy Operational Causation (D1b), Observer Comparability (D2), and Admissibility (FD/IC).

Definition CLT.3 (Commitment Map). A *commitment map* is a reduction $R : S \to \mathscr{R}$ from the amplitude-level state space S to a record space $\mathscr{R}$ of discrete events. R may be injective or non-injective. If R is non-injective *and* $R^{-1} \notin \mathcal{O}(R, \varepsilon)$, then R constitutes operational irreversible commitment (IC) at the admissible layer.

Definition CLT.4 (Measurement-Localized Commitment). A theory satisfies *measurement-localized commitment* if: (i) fundamental propagation is reversible; (ii) irreversible commitment occurs only during interactions with macroscopic measurement devices; (iii) the causal order between non-measurement events is defined by amplitude-level microcausality.

Lemma 1 (Causation Requires Stable Record Carriers)

If $E_i \prec E_j$ is a physically meaningful fact-causal relation (D1b), then there exists a stable record carrier R whose state depends on E_i and persists above the distinguishability threshold ε long enough to constrain E_j.

Proof. Counterfactual dependence (D1b) requires that some physical degrees of freedom encode information about E_i and influence E_j. The A-dependent mutual information $I(E_i; R)$ must remain above ε throughout transit from E_i's causal domain to E_j's. If all influence from E_i is reversibly recombinable before reaching E_j — if the carrier is scrambled below ε — then interventions at E_i cannot systematically alter outcome distributions at E_j, and no counterfactual dependence survives. Therefore stable records above the ε-threshold are required. This is Theorem A restated for the fact-causal order. ■

Lemma 2 (No Guaranteed Apparatus-Invariant Fact Partition Without Structural Commitment)

In a measurement-localized commitment model (CLT.4), there is no guaranteed apparatus-invariant partition of amplitude histories into fact-events within the admissible domain $\mathcal{O}(R, \varepsilon)$. Any invariance achieved is contingent on dynamical details of the measurement coupling and therefore cannot serve as a foundation for observer-comparable fact-causal order.

Proof (sketch with worked example). Let $R\_M : S \to \mathscr{R}$ denote the record map induced by measurement coupling M (pointer observable, coupling Hamiltonian, environment structure). In a measurement-localized model, reversible propagation between measurements supplies no intrinsic event partition; the partition is induced only when R_M is applied. For two admissible couplings M_1, M_2 (both within $\mathcal{O}(R, \varepsilon)$), there is no structural guarantee that the induced partitions $R\_\{M_1\}$ and $R\_\{M_2\}$ coincide on the same amplitude histories.

Worked example. Let the system state arriving at a detector be $|\psi\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. Define two admissible measurement couplings to pointer states:

- M_z: $|0\rangle|P^*\rangle \to |0\rangle|P_0\rangle$ and $|1\rangle|P^*\rangle \to |1\rangle|P_1\rangle$ (Z-basis pointer coupling)
- M_x: $|+\rangle|P^*\rangle \to |+\rangle|P_+\rangle$ and $|-\rangle|P^*\rangle \to |-\rangle|P_-\rangle$ (X-basis pointer coupling)

Both couplings are admissible: each produces stable, independently readable pointer states within $\mathcal{O}(R, \varepsilon)$. Under M_z, the committed event is "spin-z outcome 0 or 1." Under M_x, the committed event is "spin-x outcome + or −." These are incompatible partitions of the same amplitude history $|\psi\rangle$: two admissible observers using M_z and M_x disagree on which event occurred, and neither partition is derivable from the amplitude history alone — it depends on which coupling was applied. The fact partition is therefore apparatus-dependent.

For apparatus-independence, the event partition must be fixed *before* either coupling acts — by a feature of the propagation substrate that assigns non-recombinable equivalence classes to amplitude histories independent of the measurement apparatus. That is structural commitment at the propagation level (D3A), which is what the Causal Ledger Theorem requires.

Pointer-basis selection mechanisms (decoherence, einselection) may produce strong practical stability in specific physical situations, but this stability is a property of the particular coupling M and its environment — not a general structural guarantee across all admissible couplings. Therefore, without prior structural commitment, the fact partition is not guaranteed to be apparatus-invariant in the sense required by D2. ■

Remark (what is not claimed). This lemma does not claim that pointer states never stabilize or that decoherence fails. It claims only that, absent structural commitment, apparatus-invariant fact partitions are not guaranteed across all admissible couplings and therefore cannot be taken as the primitive grounding of observer-comparable causal order. Where decoherence does produce stable pointer states, it does so by effectively implementing a non-injective R_M with no admissible recovery — which is operational commitment in the sense of D3A, appearing as Case 2 in Theorem B.

Lemma 3 (Microcausality Cannot Define Fact Order Alone)

Assume fundamental dynamics is reversible and satisfies microcausality (CLT.1). To define a fact-causal order (CLT.2), one must introduce a commitment map R (CLT.3). Every such map either (Case 1) is injective — making event structure representational and dictionary-dependent — or (Case 2) is non-injective with no admissible recovery — constituting operational irreversible commitment. Therefore microcausality alone cannot ground fact-causal order; commitment is required.

Proof.

Step 1 (Presupposition). In algebraic quantum field theory, microcausality is defined relative to a background spacetime: local operator algebras $\mathcal{A}(A)$ are assigned to regions of a manifold with given causal structure. In the pre-geometric programme, this causal structure is what we seek to derive from committed-record propagation. One cannot use algebraic causal order to ground the causal facts if algebraic causal order already presupposes the answer. This limits microcausality to frameworks with a background spacetime — outside the pre-geometric setting, it is not available as a grounding structure.

Step 2 (Operators vs outcomes). Even given a background spacetime, the condition $[\mathcal{A}(A), \mathcal{A}(B)] = 0$ establishes that operations at A cannot signal to B via expectation values. This is D1a — dynamical influence. It does not establish which outcomes *occurred*. The same commutation structure is compatible with any distribution of committed facts: a measurement at A yields a committed record r_A regardless of the commutation relations with $\mathcal{A}(B)$. The algebra constrains what could happen; only committed records establish what did happen.

Step 3 (Case 1 — injective R). If the commitment map R is injective, it is a relabelling — many amplitude microstates map to distinct records, no alternatives are collapsed. Event structure is representational: which "event" occurred depends entirely on the choice of R. Different admissible observers choosing different injective R maps will partition the amplitude history into different event sequences. Fact-causal order becomes dictionary-dependent, violating Observer Comparability (D2).

Step 4 (Case 2 — non-injective R with no admissible recovery). If R is non-injective and $R^{-1} \notin \mathcal{O}(R, \varepsilon)$, then multiple amplitude microstates map to a single committed record, and the record cannot be undone within the admissible domain. This is operational IC — the alternatives are non-recombinable within $\mathcal{O}(R, \varepsilon)$. The commitment layer has been introduced into R. The ontology is no longer amplitude-only.

Step 5 (No third option). Cases 1 and 2 are exhaustive for any commitment map R: either it is injective (representational, dictionary-dependent) or it is non-injective (commitment present, implicitly or explicitly). Microcausality does not select between them; it governs only operator commutation, not the choice or character of R. Therefore microcausality alone cannot ground fact-causal order. ■

Scope remark. Lemma 3 applies within the pre-geometric setting (Theorem B scope note): frameworks where causal order is not presupposed as background structure. In algebraic QFT with a fixed background spacetime, local operator algebras $\mathcal{A}(A)$ are assigned to spacetime regions whose causal structure is stipulated by the manifold, and within that setting microcausality plus the given spacetime structure does define a consistent causal order. Lemma 3 does not contest this. It establishes that microcausality alone cannot *derive* causal order — which is only a requirement if one is in the business of deriving it. Readers working in fixed-background QFT are outside the scope of this lemma by construction, not by refutation.

Causal Ledger Theorem (Pre-Geometric Scope)

Theorem (Causal Ledger). *In any framework that requires causal order to be physically instantiated and observer-comparable rather than stipulated as background structure, fact-causal order cannot be grounded in reversible amplitude dynamics and microcausality alone. A commitment layer producing operationally non-recombinable records is necessary, and the event order is therefore defined on committed records — the causal ledger.*

Proof. From Lemma 1, any fact-causal relation $E_i \prec E_j$ requires a stable record carrier sustaining $TV > \varepsilon$ in transit. From Lemma 2, measurement-localized commitment provides no structural guarantee of apparatus-independent fact partitions; contingent pointer-state selection is insufficient for general D2 compliance. From Lemma 3, any commitment map R is either injective (dictionary-dependent, Case 1) or non-injective with no admissible recovery (operational IC, Case 2); microcausality does not resolve this choice. Therefore admissible causal order requires structural commitment at the propagation level. ■

Consequence for propagation geometry. The Causal Ledger Theorem establishes that the dilution argument in Section 7 targets the causal ledger layer, not the reversible amplitude layer. Reversible amplitude evolution remains valid as an effective description between commitment events. It does not define the causal substrate. The dimensional constraint is conditional on Assumptions A6′ and A7 (Assumptions Ledger).

Summary of the result chain. Theorem A: D1b requires stable records. Theorem B: amplitude-only D1b fails in three classifiable ways. Causal Ledger Theorem: structural commitment is required; measurement-endpoint localization and microcausality are jointly insufficient.

---

4. Corollary: The Causal Graph Is Built from Commitment Events

Corollary. *In any admissible physics, the causal graph — the directed graph whose nodes are events and whose edges are causal relations $A \prec B$ — is built from commitment events and the* *stable records they propagate. Causal propagation is the propagation of committed distinguishability.*

Proof. By Theorem A, every edge A ≺ B in the causal graph requires a stable record R carrying A-information into B's causal domain. Each such record is the output of a commitment event — an irreversible transition from the uncommitted to the committed region (a crossing of Σ in the state-space picture). The record persists as a committed fact until it either decoheres below operational resolution or is consumed/overwritten by a subsequent commitment event. The causal graph is therefore the graph of commitment events and committed-record propagation. Unitary amplitude evolution occurs between commitment events and determines the probability amplitudes that govern which commitment events occur and with what probability — but the causal edges themselves are grounded in commitments, not amplitudes. □

Remark on the two-layer structure. This gives a clean two-layer picture of physics:

- *Reversible substrate layer (void/tick layer):* unitary or diffeomorphic evolution of amplitudes. This layer is causally neutral — it determines what is possible, but not what is actual.
- *Commitment layer (Σ and beyond):* irreversible transitions that produce stable records. This layer is where causal relations are grounded. The causal graph lives here.

The reversible layer is not causally inert — it determines transition probabilities between commitment events. But it is not the *carrier* of causal relations. Those relations are carried by committed records propagating forward from Σ.

---

5. Implications and the Referee Escape Hatch

A persistent objection in discussions of commitment-based physics takes the following form:

"Light propagates as amplitudes. The speed of light is a feature of the wave equation, not of any commitment layer. Therefore budget arguments about commitment events (dilution, bit-cost, etc.) do not constrain c or light propagation."

Theorems A and B expose the equivocation in this objection. It conflates two different things:

1. **The amplitude propagation of the electromagnetic field** — the solution to Maxwell's equations, which is indeed reversible, unitary, and commitment-free.
2. **The causal claim that "a light signal from A caused a detection event at B"** — which requires, by Theorem A, a stable record at B that carries A-dependent information.

The first is amplitude propagation. The second is causal propagation. They are not the same thing. The electromagnetic wave propagates amplitudes. The *causal relation* A ≺ B is grounded in the committed detection event at B — the irreversible click of the detector, the committed photoionisation, the stable record that persists beyond the moment of interaction.

If a referee wants to say that causation is amplitude propagation alone — that A $\prec$ B holds in virtue of the electromagnetic wave passing through the region between A and B, with no commitment layer at B — then their account of causation is not observer-comparable (different observers can never agree on whether B "detected" the signal, because no outcome is ever settled) and is therefore inadmissible by D2.

The correct summary is: **the wave equation governs what arrives at B; the commitment layer decides what B records.** Both matter. Budget arguments about commitment events govern the second layer — and the second layer is where causal claims live.

---

6. Anticipated Objections

Objection 1: Everettian Quantum Mechanics Avoids Commitment

"In Everettian (many-worlds) quantum mechanics, global unitarity is maintained permanently. No branch is ever preferred and no collapse occurs. The no-go lemma therefore does not apply."

Response. This objection is addressed without taking a position on the metaphysics of many-worlds. The lemma does not require fundamental non-unitarity. It requires *operational IC* — that within the admissible process class $\mathcal{O}(R, \varepsilon)$, certain alternatives become non-recombinable. Whether this reflects fundamental collapse or merely resource-bounded inaccessibility is irrelevant to the operational argument.

The operative question is: are branch distinctions operationally recoverable within $\mathcal{O}(R, \varepsilon)$?

If yes — if an observer within an Everettian branch can coherently recombine with other branches and detect interference using resources bounded by R — then no commitment has occurred at the admissible layer, and Case 3 applies: no settled outcomes, no admissible causal claims of type D1b.

If no — if branch recombination requires resources exceeding R, or if environmental entanglement makes branches non-interferable at the admissible scale — then the branching event constitutes effective IC at the admissible layer. Alternatives are operationally non-recombinable. The lemma applies regardless of whether the global wavefunction is unitary.

The Everettian account of records (branch-relative stable outcomes) is therefore compatible with the lemma: branch-relative records are the commitment events in our sense, localised to the admissible domain. The lemma is interpretation-agnostic; it operates at the layer of admissible operations, not fundamental ontology. Everett commits at the admissible layer even if it does not commit globally.

Objection 2: Decoherence Explains Effective Irreversibility Without Genuine Commitment

"Decoherence explains why macroscopic systems behave irreversibly without introducing genuine commitment. The environmental correlations are real but inaccessible; no fundamental irreversibility is needed."

Response. Decoherence provides stability of reduced descriptions — robust records (D3) — generically and efficiently. This paper does not dispute that. The point is more precise.

Decoherence displaces coherence into environmental degrees of freedom. The reduced density matrix becomes effectively diagonal; the system behaves classically. This gives a D3 robust record: the pointer state is stable and independently readable. What decoherence does not automatically provide is operational closure — the condition that no admissible operation in $\mathcal{O}(R, \varepsilon)$ can recover the pre-decoherence microstate and restore coherence. That additional condition is IC.

If the environmental degrees of freedom are formally inaccessible within $\mathcal{O}(R, \varepsilon)$ — if reconstructing the full system-plus-environment state is beyond admissible resource bounds — then the decoherence map is operationally committing: it is a non-injective R_M with no admissible inverse. In that case, decoherence confirms the commitment framework rather than evading it.

The novelty of this paper is not "decoherence fails" but:

1. The D1a/D1b separation: decoherence grounds D1a dynamical influence; the additional IC condition is what upgrades it to D1b causal ancestry.
2. The trilemma classification: any amplitude-only account either imports IC covertly (decoherence as Case 2), becomes representation-dependent (Case 1), or forfeits settled outcomes (Case 3).
3. The apparatus-invariance requirement: not just that pointer states stabilize in one model, but that the event partition is structurally guaranteed across all admissible couplings — which requires IC at the propagation level, not at the measurement endpoint alone.

Decoherence is the mechanism; operational closure is the condition. The paper specifies when the mechanism meets the condition.

Objection 3: Causation Can Be Defined Purely Relationally Without Settled Outcomes

"Causal structure can be defined purely relationally — as a partial order on events — without requiring any event to be 'settled' or 'committed.' Causal set theory, for instance, defines spacetime as a partial order without invoking records."

Response. Purely relational definitions of causal order — including causal set approaches — still require persistent relational structure. The partial order A ≺ B must be a fact about the world that different observers can agree on and use to make predictions.

If the partial order is itself a transient amplitude structure — if the relation A ≺ B can be undone by admissible inverse evolution — then it does not constitute a stable causal fact and cannot support counterfactual claims of the form "if A had been different, B would have been different." The partial order must be committed to be operationally usable.

More precisely: defining "event A" in a purely relational causal set requires that A is a definite, settled node in the partial order — not a superposition of nodes that might interfere. If events are not settled, the partial order is not defined (or is representation-dependent, falling into Case 1). If events are settled, we have stable records. The relational framework presupposes committed events; it does not eliminate them.

The no-go lemma concerns *admissible* physical causation of type D1b — causation that supports counterfactual dependence on recorded outcomes, finite distinguishability, and observer comparability. Abstract relational ordering that lacks these properties is not a theory of physical causal ancestry in the sense of D1b; it is a mathematical structure that may or may not correspond to physical reality. For the correspondence to hold, the relational nodes must be physically settled — which is commitment.

A further objection deserves direct treatment: in algebraic quantum field theory, a canonical local algebra of observables can define events as projectors in a local algebra, without an arbitrary dictionary choice. Does this escape Case 1?

It does not, for the following reason. Algebraic QFT constructs its causal structure by presupposing a background spacetime with a given causal order. The local algebras are assigned to regions of a *pre-given* causal structure. If one is asking how that causal structure itself becomes physically authoritative — how it comes to be the structure that "counts" as the physical causal order rather than one of many equally valid algebraic representations — then one cannot help oneself to a background causal order at the outset. Our claim is not that AQFT is wrong; it is that AQFT does not solve the derivation problem for causal order because it assumes it. The programme of this paper is pre-geometric: causal order is what we seek to derive from admissibility, not what we assume in order to apply it.

Remark on indefinite causal order (quantum switch). The process-matrix framework of Oreshkov, Costa & Brukner (2012) exhibits correlations inconsistent with any definite causal order between events — a phenomenon realized experimentally in quantum switch scenarios. A referee may ask where such scenarios fall within the present framework. The present paper operates under the assumption that commitment events have a definite causal ordering — specifically, that the committed records constituting the causal ledger are partially ordered by the record-causal ancestry relation. This is implicit in the structure of D1b (which requires a definite directed relation A ≺ B) and the Causal Ledger Theorem (which builds a causal graph on committed events). Quantum switch scenarios, where the causal order between operations is itself in quantum superposition, therefore fall outside the present framework's direct scope: they raise the question of whether the commitment layer itself can exhibit indefinite causal structure. The most natural landing point in the trilemma is a variant of Case 3: not that local outcomes at A and B are unsettled — they may each be individually committed — but that the causal ancestry relation $A \prec B$ itself lacks a committed value. D1b requires a definite directed relation; in the quantum switch, that relation is in superposition, so D1b causal ancestry is undefined until a further commitment event settles which ordering obtained. The indefiniteness is in the causal structure, not in the event structure — a subtler form of Case 3 where causal ancestry rather than event occurrence is what remains uncommitted. Whether the process-matrix framework can be extended to include a commitment layer that collapses causal indefiniteness into definite ledger entries is a substantive open question that the present results constrain but do not resolve.

Objection 4: Reversible Error-Correcting Codes Provide Stable Records Without Commitment

"Topological quantum codes, integrable systems, and other reversible substrates can protect logical bits indefinitely against errors without any irreversible commitment. These provide stable records that could ground causal ancestry."

Response. A reversible protected memory and a committed fact are structurally distinct in the way that matters for causal ancestry.

A reversible protected memory preserves a logical bit against local perturbations, but the bit remains *recombinable*: a sufficiently powerful adversary with admissible resources could in principle apply the inverse encoding, recover the full microstate, and undo the "record." The bit is stable against noise but not against reversal. In the framework of this paper, a record grounds causal ancestry only if its alternatives are non-recombinable within $\mathcal{O}(R, \varepsilon)$ — if no admissible operation can restore the pre-record microstate. This is the IC condition. Reversible protection satisfies stability; it does not satisfy non-recombinability.

Put concretely: suppose Alice uses a topological code to "record" the outcome of an event A. The code protects the logical bit from local noise, but Bob, with sufficient quantum hardware within resource bound R, can decode and re-encode the state in a different basis, access the full amplitude correlations, and in principle detect interference that would not exist if the record had been genuinely committed. The code has not committed the fact — it has preserved an amplitude configuration. The distinction between "stable amplitude" and "committed fact" is exactly what the IC condition formalises.

Objection 5: The Commitment Layer Need Not Be the Propagation Substrate

"Commitments occur only at detectors/interaction points; amplitudes propagate in 3D, so commitment arguments do not constrain propagation geometry."

Response. This objection is correct for D1a weak dynamical influence: amplitude propagation is compatible with many geometries and is governed by the effective field equations. However, the present paper concerns D1b fact-causal order: observer-comparable causal ancestry on recorded outcomes. By the Causal Ledger Theorem, D1b causal order is instantiated on committed records, and therefore any claim of causal ancestry — including between detector events — must be grounded in record propagation and closure at the admissible layer. The question of propagation geometry is therefore a question about the commitment ledger, not about reversible amplitudes. Dimensional implications for record propagation are treated separately and conditionally in Section 7.

Assumptions Ledger

The paper's results depend on the following assumptions. The first two are premises of the foundational framework; the third applies only to Section 7.

A0′ (Operational Closure Under Record Exchange). A physical theory is admissible only if it supports the following closure property: for any two admissible observers O_1 and O_2 performing interventions and readouts within $\mathcal{O}(R, \varepsilon)$, the joint distribution of their recorded outcomes admits a consistent refinement under record exchange that does not depend on arbitrary representation choices. That is: if O_1 measures A and O_2 measures B and they later compare results, there must exist a representation in which both agree on the causal relation $A \prec B$ without invoking non-admissible operations.

Justification of A0′. A0′ is not a metaphysical claim about ontology; it is a structural requirement for the existence of empirical law. If two admissible observers cannot, upon exchanging records, arrive at a representation-independent account of which events occurred and how they are causally related, then the theory lacks closure under record exchange and cannot support invariant law statements. In such a theory, causal relations would be observer-relative artifacts of representation rather than physically grounded structures. A0′ therefore expresses a minimal consistency condition for empirical physics: causal order must be stabilizable under admissible record comparison. A0′ is not a preference about how physics should be done; without it, D1b causal ancestry is undefined — there is no representation-independent causal graph to state.

Three structural justifications. (1) *Predictive consistency:* without A0′, predictions are observer-relative, law statements cannot be globally formulated, and no single causal graph exists — the theory lacks a global predictive structure. (2) *Law invariance:* a physical law must be invariant under admissible observer transformations; if causal order depends on the measurement basis or decomposition chosen by an observer, the law is representation-dependent, violating physical invariance. (3) *D1b requires it:* counterfactual dependence requires that outcomes can be stably compared after the fact; if they cannot, causal ancestry collapses.

Proposition (Operational Closure Requirement). Any theory lacking A0′ cannot define a representation-invariant causal graph over recorded events.

Proof sketch. Without A0′, two observers performing admissible interventions may disagree on the event partition induced by their measurements. Disagreement on event partitions implies disagreement on which events are causally ordered. The causal graph therefore depends on the observer's representation choice — it is not a representation-invariant structure of the theory. This is exactly the structure of Case 1 (Dictionary/Subspace Dependence) in Theorem B: the causal order becomes representation-dependent. ■

Scope note. The claim is conditional: if one wants D1b (empirical causal ancestry on recorded outcomes), then A0′ is structurally required. This paper takes no position on theories that do not aspire to D1b. Whether Everettian QM or relational QM satisfies A0′ is what Objection 1 contests and is not resolved here by stipulation — those interpretations must establish on their own terms that their branch-relative or relational facts satisfy the consistent-refinement condition. The paper specifies what success requires; it does not adjudicate whether it is achieved.

FD + IC (from D2). Causal claims invoke only operationally resolvable distinctions (FD), and outcome events serving as causal antecedents are irrecoverable within $\mathcal{O}$(R, ε) (IC). Together with A0′, FD + IC ground the core results (Theorems A, B; Causal Ledger Theorem). Everything through Section 6 follows from A0′, FD, IC, and the pre-geometric scope of D1b.

A6′ (No Free Broadcasting of Committed Distinguishability). The *committed constraint rank* J(ρ) is the minimal number of irrecoverable binary constraints required to encode the committed distinction in state ρ under the equivalence classes of $\mathcal{O}$(R, ε). J is discrete and structural — it counts irrecoverable distinctions, not thermodynamic entropy or Shannon bits. J is invariant under admissible reversible transformations and increases only under commitment events.

A6′ (formal statement): J is a monotone under admissible propagation. If $\mathscr{E}$ is an admissible propagation map that does not include new commitment events, then:

$J(\mathscr{E}(\rho)) \leq J(\rho)$

Broadcasting would require creating additional irrecoverable constraints — either violating IC (irrecoverability) or exceeding admissibility bounds. Neither is permitted without new commitment events.

Corollary (MI bound). Since J(ρ) bounds the number of independently extractable committed distinctions, and mutual information $I(S; Y) \leq J(\rho) \cdot \log 2$ for any admissible readout Y (the J constraints are binary), it follows that $I(S; Y_{1:N(r)}) \leq J_{src} \cdot \log 2 = I_{src}$. The data processing inequality then gives $I_r \leq I_{src}$ as used in Proposition 7.1.5. Mutual information is a diagnostic lower bound derived from J, not the primary monotone.

Thermodynamic support (Landauer's principle, no-broadcasting) follows as a consequence: each unit increase in J requires dissipation and a new committed sink, because irrecoverable constraints cannot be created for free. The thermodynamic reading is not the definition of J; it is what J's structural monotonicity implies physically.

One-sentence intuition: A6′ is the ledger-level analogue of no-cloning/no-broadcasting — a committed fact cannot be freely replicated into many independently addressable committed facts without additional irrecoverable constraints and corresponding physical cost.

A7 (Spatial Locality and Finite Causal Reach of Commitment Events). Commitment events are spatially localized: each is associated with a definite spatial region of size ≥ ℓ_min, and the causal influence of a commitment event propagates at finite speed v with a bounded interaction radius per unit time. Independently usable committed records occupy disjoint spatial support regions — a single committed bit cannot be nonlocally encoded across arbitrarily many spatially disjoint addresses without additional commitment events. A7 is grounded in Lieb–Robinson-type locality bounds for finite-speed causal propagation: in any system with finite propagation speed and bounded local interaction range, the number of causally independent spatial regions reachable from a source at radius r in time r/v scales as the surface area ~ r^{d-1} of the light cone cross-section. A7 converts this locality bound into the addressability count used in the dilution argument. It is an additional structural assumption about the commitment substrate; it is not implied by the core results. The dilution argument in Section 7 requires both A6′ and A7.

Strategic note. A reader who accepts Theorems A, B and the Causal Ledger Theorem but is skeptical of A6′, A7, or the geometric scaling argument in Section 7 loses nothing. The foundational claim stands without Section 7.

Falsification criteria: see Falsification E for A6′ (Section 8).

---

7. Dimensional Admissibility — Conditional Implication

The results of Sections 3–6 establish that observer-comparable causal order is defined on committed records (the causal ledger). Section 7 draws a further conditional implication, dependent on A6′ (Assumptions Ledger) and a geometric scaling argument. It can be read independently of, and is not required by, the core results. Readers whose interest is in the foundations result alone may proceed directly to Section 8.

The connection. Causal propagation is the propagation of committed distinguishability (Corollary, Section 4). Committed distinguishability is carried by stable records. Stable records are produced at commitment events — crossings of the interface Σ. The causal graph is therefore a graph on commitment events, and the spatial structure of causal propagation is the spatial structure of Σ and the committed records emanating from it.

Assumption A6′ (No Free Broadcasting of Committed Distinguishability). Restated for this section: the total committed distinguishability about a source S cannot increase under admissible propagation without additional commitment events. Each independently readable cell at radius r is a separate addressable channel for extracting S-dependence; A6′ bounds the total across all such channels.

Dimensional clarification. The argument below uses d to refer to the effective ambient spatial dimension — not the intrinsic dimension of Σ itself. The companion paper establishes $\dim(\Sigma) = 2$ (codimension-1 in the 3D ambient). The dilution argument asks: in how many ambient dimensions can a single committed bit remain independently addressable? The two are consistent: the dilution result is what the companion paper's codimension-1 structure is designed to satisfy.

Section 7.1 — Local Dispersion Bound and Dilution of Committed Distinguishability

This subsection provides the formal derivation of the dilution scaling. Three ingredients suffice: (i) independently readable committed records require disjoint spatial support at operational resolution; (ii) propagation is local with bounded interaction radius; (iii) committed distinguishability cannot be freely broadcast without additional commitment events (A6').

Definition 7.1.1 (Operational Support and Disjoint Readout Cells). Fix operational resolution $\delta = \delta(R, \varepsilon)$ and let $\ell_{min} \geq \delta$ denote the minimum spatial scale at which an admissible readout can localize and stably discriminate a record (Assumption A7). A committed record instance R at time t has *operational support* in region $D \subset \mathbb{R}^d$ if there exists an admissible readout $U_D \in \mathcal{O}(R, \varepsilon)$ acting only on degrees of freedom in D such that the outcome distribution distinguishes the record value above ε. Two record instances are *independently addressable* at time t if they admit readout regions D_1, D_2 with $\mathrm{dist}(D_1, D_2) \geq 2\ell_{min}$, so that admissible readouts on D_1 and D_2 are disjoint and do not require joint global control. Independently addressable readouts are defined operationally such that no admissible observer can exploit correlations across disjoint support regions without exceeding resource bounds — any operation that would require joint access to both D_1 and D_2 simultaneously is by definition outside $\mathcal{O}(R, \varepsilon)$.

Definition 7.1.2 (Bounded Interaction Radius). There exists an interaction radius $\rho > 0$ (of order ℓ_{min}) such that over one admissible time step Δt, information about the committed record carrier can only influence degrees of freedom within distance ρ. The forward evolution map $\mathcal{E}_{\Delta t}$ is ρ-local: the reduced state on any region D at time $t + \Delta t$ depends only on the state on the ρ-neighborhood $N_\rho(D)$ at time t. This is the operational content of finite-speed causal propagation.

Lemma 7.1.3 (Packing Bound). Let $S_r \subset \mathbb{R}^d$ be a sphere of radius r and let $\{D_i\}_{i=1}^{N(r)}$ be a collection of disjoint readout cells of diameter $\sim \ell_{min}$ with $D_i \subset N_\rho(S_r)$. Then:

$$N(r) \leq C(d, \rho/\ell_{min}) \cdot (r/\ell_{min})^{(d-1)}$$

Proof. The (d−1)-dimensional surface area of S_r scales as $c_d r^{(d-1)}$. Each disjoint readout cell of linear scale ℓ_{min} intersects the shell $N_\rho(S_r)$ in a patch of (d−1)-dimensional measure $\Omega(\ell_{min}^{(d-1)})$ up to constants depending on ρ/ℓ_{min}. Since cells are disjoint, the total (d−1)-measure of intersected patches is at most the (d−1)-measure of the shell, which is $O(r^{(d-1)})$. Dividing gives $N(r) \leq C(r/\ell_{min})^{(d-1)}$. ■

In a d-dimensional ambient geometry, the number of independently addressable disjoint readout regions at radius r grows at most like $r^{(d-1)}$. This is the precise mathematical content of "new addresses grow like surface area."

Definition 7.1.4 (Total Committed Information at Radius r). Let $S \in \{0, 1\}$ denote the committed source bit at the origin. Let $Y_{1:N(r)} = (Y_1, \ldots, Y_{N(r)})$ be the joint outcomes of all independently addressable disjoint readout cells on the shell of radius r. Define:

$$I_r := I(S; Y_{1:N(r)})$$

This is the total committed distinguishability available at radius r — the only information-theoretically meaningful conserved quantity under admissible propagation. No additivity across cells is assumed; correlations between cell outcomes are permitted.

Proposition 7.1.5 (Ledger-Level Data Processing Bound). Under A6′ (no free broadcasting of committed distinguishability) and bounded interaction radius (Definition 7.1.2):

$$I_r \leq I_{src}$$

where I_{src} is the committed information content of the source event (of order one bit for a primitive commitment).

Proof. The propagation from the source commitment to the shell at radius r is an admissible CPTP map in $\mathcal{O}(R, \varepsilon)$. By the data processing inequality:

$$I(S; Y_{1:N(r)}) \leq I(S; C_{src})$$

A6′ asserts that no additional committed sinks are created during propagation and no new independent committed constraints are injected. Therefore $I_r \leq I_{src}$. ■

Dilution via independent addressability. To serve as a reusable causal ancestor (D1b), a source event must support multiple independently readable downstream observers. Let each independently addressable cell require at least $\varepsilon_I > 0$ mutual information with the source: $I(S; Y_i) \geq \varepsilon_I$. Since the joint information I_r upper-bounds the information extractable by any single observer network, and each disjoint observer extracts information from disjoint spatial support, the necessary condition is:

$$N(r) \cdot \varepsilon_I \leq I_r \leq I_{src}$$

Combining with the packing lemma (Lemma 7.1.3), $N(r) \sim r^{(d-1)}$:

$$\varepsilon_I \lesssim I_{src} \cdot r^{(-(d-1))}$$

Absent new commitment events, the per-address committed information budget decays as $r^{(-(d-1))}$.

Lemma 7.1.7 (Transverse Constraint Independence). In a k-dimensional causal boundary, the number of independent locally distinguishable transverse gradient directions equals k. Each independent transverse gradient direction corresponds to an independent committed constraint channel at the boundary.

Proof sketch. Local distinguishability of committed records at a boundary point requires variation in at least one independent coordinate direction. In a k-dimensional manifold, there exist exactly k independent local coordinate directions (the tangent space has dimension k). Each coordinate direction supports independent perturbation modes — variations of the committed record value along that direction that are operationally distinguishable from variations along any other direction. Orthogonal perturbation modes require independent J constraints because a single irrecoverable binary constraint can encode a distinction only along one direction in state space: resolving a second orthogonal distinction requires a second constraint whose inversion is also outside $\mathcal{O}(R, \varepsilon)$, and no admissible operation can collapse two orthogonal committed distinctions into one without losing information about at least one of them. Under admissibility, each such independent perturbation mode therefore corresponds to an independent J-contribution: to encode two independently distinguishable constraint states that differ in orthogonal directions, at least two irrecoverable binary constraints are required. Therefore a k-dimensional causal boundary supports exactly k independent committed constraint channels. ■

Admissibility Criterion (Minimal Non-Degenerate Causal Boundary). The dimensional threshold is structural. A 0-dimensional boundary cannot propagate. A 1-dimensional boundary (a curve) has exactly one independent transverse gradient direction (Lemma 7.1.7), supporting only one independent committed constraint channel — it cannot sustain multi-directional distinguishable causal structure for independently addressable downstream events. The minimal dimension of a non-degenerate causal boundary — one that supports at least two independent committed constraint channels and therefore nontrivial spatial separation in distinct causal directions — is 2.

Consequence 7.1.6 (Dimensional Inadmissibility — structural). The argument runs as a linear chain of six steps:

1. *Non-degenerate causal propagation requires ≥ 2 independent transverse constraint channels.* A boundary with fewer than 2 independent J-channels cannot support independently addressable committed records in distinct causal directions. This is the minimal condition for D1b reusable causal ancestry across spatially separated events.
2. *By Lemma 7.1.7, a k-dimensional causal boundary supports exactly k independent constraint channels.* The channel count equals the boundary dimension.
3. *Therefore the causal boundary must have intrinsic dimension ≥ 2.* $d - 1 \geq 2$, hence $d \geq 3$. This is the lower bound.
4. *A6′ prohibits silent increase of J-capacity.* Each additional transverse boundary dimension beyond 2 introduces an additional independent committed constraint channel (Lemma 7.1.7). That additional channel constitutes additional J-capacity at the boundary. By A6′, J-capacity cannot be created without new commitment events. A primitive commitment event has J = 1 — exactly one irrecoverable binary constraint — which can supply at most one independent J-channel; sustaining two independent channels therefore

requires J ≥ 2, which for a primitive J = 1 source is achievable only if both channels share the single constraint, possible only when the boundary dimension is exactly 2 and no surplus independent direction exists.

5. *Therefore the causal boundary cannot have intrinsic dimension > 2 for a primitive carrier.* A boundary of dimension 3 (d = 4) would introduce one surplus independent J-channel that A6′ prohibits creating silently. A primitive one-bit commitment event cannot supply J-capacity for 3 independent channels without injecting additional commitment events — which violates the primitive carrier definition.
6. *Therefore d − 1 = 2, hence d = 3.*

The threshold is not asymptotic or scaling-dependent. It is a structural consequence of Lemma 7.1.7 (boundary dimension equals independent J-channel count) and A6′ (J is monotone non-increasing under admissible propagation without new commitment events). The argument holds at any radius; no $r \to \infty$ limit is required.

The conclusion is not "spacetime is 3-dimensional" — that requires the companion paper's codimension-1 result dim(Σ) = d_ambient − 1. Combined, one obtains the effective 3D substrate conclusion. Amplitude propagation (D1a) remains dimension-agnostic throughout. Maintaining global reusability in d ≥ 4 would require commitment events injected at a rate scaling with the extra transverse J-channels, converting primitive propagation into super-extensive ledger creation and violating the primitive carrier definition.

Remark on violations. Violations correspond exactly to Falsification E: a process producing N independently readable committed copies at large radius without additional commitment events or dissipation scaling with N.

The wave equation governs amplitude propagation in any dimension. The commitment ledger governs causal propagation — and the ledger has a preferred dimension.

---

8. Failure Modes and Falsification

Falsification A (Causal Claims Without Records). If a physically meaningful causal relation A ≺ B can be demonstrated in which no stable record carries A-information toward B — where the influence is purely through transient, reversible amplitude correlations, with no committed intermediate states — then Theorem A fails. This would require showing that counterfactual dependence can be established and observer-verified without any committed intermediate facts.

Falsification B (Admissible Pure-Amplitude Causation). If an account of causation can be given that satisfies D1–D2 in full (including observer-comparability and finite distinguishability) without invoking any commitment layer — and without falling into any of the three cases of the Amplitude Trilemma — Theorem B fails. Any such account must identify: (i) how outcome events become stable enough to serve as causal antecedents, (ii) how different observers agree on which outcomes occurred without a shared committed record, and (iii) how causal directionality arises from reversible amplitudes without representation dependence.

Falsification C (Commitment Without Causation). If commitment events occur but the committed records they produce do not propagate into causal domains of future events — if facts are produced but immediately become causally isolated — then the corollary would fail. The causal graph would not be the graph of commitment propagation. This would require a physical mechanism that produces stable records while simultaneously shielding them from influencing downstream events.

Falsification D (The Transcendental Premise). If the claim "facts exist" is denied — if the practice of physics (performing experiments, recording outcomes, comparing results) is groundless — then Step 5 of Theorem B is blocked. The many-worlds interpretation does not deny causation; it relocates it to the branch-relative structure of the wave function. Addressing this case requires specifying what "observer-comparable" means in a no-collapse ontology — a non-trivial extension of D2 that is not attempted here, and which must in any case grapple with the Everettian response in Section 6.

Falsification E (A6′ — Record Self-Broadcasting). Exhibit a physical process where a single committed bit at a source S produces N spatially separated committed copies at distance r without (i) additional record sinks (new apparatus or environment degrees of freedom that themselves become committed), (ii) additional free-energy dissipation scaling with N (violating the Landauer bound on reliable record creation), or (iii) additional commitment events (further crossings of Σ). Specifically: show that the committed distinguishability I_comm(S; ·) grows superlinearly with the number of independently addressable spatial regions spanned, without any of (i)–(iii). If such a process exists, the data processing inequality for committed distinguishability is violated, the dilution monotone fails, and the dimensional inadmissibility result does not follow.

---

9. Relation to the Codimension-1 Programme

This paper is logically downstream of the codimension-1 commitment surface paper but does not require its results. The relationship is:

- The **codimension-1 paper** establishes that commitment events occur at a codimension-1 interface Σ in state space, that this interface has effective ambient dimension 3 (with dim(Σ) = 2 from the trilogy), and that no silent extra transverse dimensions exist.
- **This paper** establishes that causal propagation requires committed records — that the causal graph of physics lives in the commitment layer, not the amplitude layer.

Together they give a complete structural picture: commitment events occur at Σ (codim-1 paper); causal relations are grounded in the committed records these events produce (this paper); the propagation of these records through the 3-dimensional effective substrate constitutes physical causation.

The connection to the hex interface framework is direct: the dilution argument for the commitment budget applies to the *causal* propagation of committed bits, not to the amplitude propagation of the electromagnetic field. Theorem A explains why: only the committed layer carries the causal relations. Dilution of committed distinguishability is therefore a constraint on causal propagation, not on amplitude propagation.

---

10. Discussion

The core logical chain is:

1. Two senses of causation must be distinguished: D1a (weak dynamical influence, representable by amplitude commutators) and D1b (strong causal ancestry on recorded outcomes, the subject of this paper) (Section 2).
2. D1b requires a physical carrier from A to B that persists above operational resolution ε (Theorem A, Steps 1–2).
3. Persisting, resolvable, A-dependent carrier = stable record in the sense of D3, not merely a stable amplitude (Theorem A, Steps 3–4).
4. Observer-comparable D1b causal claims require committed records — stable under non-recombinability, not merely reversible protection (Theorem A, Step 4; Objection 4).
5. Any attempt to define observer-invariant causal order in a pure amplitude ontology falls into one of three named failures: representation-dependence, non-injective operational closure, or no settled outcomes — where Case 2 requires both non-injectivity *and* non-admissible recovery (Theorem B, hardened).
6. The three standard escape hatches map onto trilemma cases: Everett = operational IC at admissible layer; decoherence = Case 2 with environmental trace; purely relational / algebraic QFT = Case 1 (Objections 1–3).
7. Reversible protected memory does not constitute commitment because it lacks non-recombinability (Objection 4).
8. Measurement-localized commitment — commitment only at apparatus endpoints — cannot yield apparatus-independent causal order, because event partitions are not structurally guaranteed to be invariant across admissible couplings. Structural commitment at the propagation level is required (Lemma 2).
9. Algebraic locality (microcausality, commutators) is insufficient to ground fact order: it presupposes background spacetime and constrains operators rather than grounding which outcomes occurred. It is a necessary but insufficient component; structural commitment is also required (Lemma 3).
10. The commitment layer therefore cannot be decoupled from the causal propagation substrate in a pre-geometric framework (Objection 5).
11. The causal graph is therefore built from commitment events and committed-record propagation (Corollary).
12. The geometry of causal propagation is constrained by commitment-capacity scaling, yielding dimensional inadmissibility for $d \geq 4$ under A6′ and A7 (Section 7).
13. The existence of representation-invariant physical law presupposes D1b, not merely D1a.

Steps 8 and 9 are the structural additions of this version. The earlier argument (Theorems A and B) established that some commitment is necessary and that amplitude-only accounts fail in three classifiable ways. Lemma 2 establishes *where* commitment must occur: at the propagation substrate level, not localized to measurement endpoints, because event partitions must be structurally fixed before apparatus acts. Lemma 3 closes the microcausality escape: algebraic locality is a constraint on dynamical influence (D1a), not on which outcomes occurred (D1b); the map from commutation structure to committed facts requires the commitment layer, not the other way around.

No Law Without Reusability. A physical law relates intervention classes to stable outcome distributions — it states that if A is prepared in state s, outcome distribution p at B reliably follows. But for such a relation to be a *law* rather than a one-time observation, the outcome effect must be reusable: independently accessible at multiple times and by multiple admissible observers. If outcome effects are not reusable — if they exist only as transient amplitude correlations that cannot be independently accessed, compared, and aggregated into a stable statistical pattern — then there is no stable object for the law to relate to. Outcome distributions measured by one observer cannot be confirmed by another; statistical regularities cannot be established across trials; the relation between intervention and outcome becomes observer-relative and representation-dependent. Law statements thereby lose invariance. D1b (reusable causal ancestry on recorded outcomes) is therefore not a stronger condition that physics might optionally require — it is the condition presupposed by the existence of any representation-invariant law over recorded events. D1a (dynamical influence) is necessary but not sufficient: an influence that leaves no reusable record cannot ground a law, only a wave equation. The commitment requirement is not an optional metaphysical addition; it is structurally implied by the existence of invariant law statements, which is what physics is.

Together, Theorems A, B and the Causal Ledger Theorem (Lemmas 1–3) form an interlocking closure argument. Every known amplitude-primitive escape route maps onto a named failure mode. The argument is interpretation-agnostic throughout — it operates at the level of operational admissibility conditions (D2), not at the level of metaphysical commitments about collapse or many-worlds. The commitment layer is not posited for interpretive convenience; it is the unique structural feature that simultaneously satisfies finite distinguishability, observer comparability, and apparatus-independence.

The two-layer picture that emerges — reversible substrate plus structural commitment layer — is not new as a physical proposal. What is new is that it is operationally forced, and that the forcing is classifiable: the trilemma (Theorem B) names the failure modes of amplitude-only accounts; Lemma 2 (apparatus dependence) names the failure of measurement-localization; Lemma 3 names the failure of algebraic-locality substitution. A referee who wishes to reject the argument must identify which failure mode their proposed escape route is exempt from, and explain why. No such exemption is currently known.

Appendix L — Lorentz Compatibility of a 2D Commitment Ledger

L.1 Scope and Claim

A standard objection to any programme that grounds causal structure in a 2D commitment interface is:

"If the fundamental substrate is a 2D interface, Lorentz invariance must be violated. But Lorentz invariance is experimentally confirmed to extreme precision."

This appendix shows that the objection conflates two distinct layers and that Lorentz invariance at the effective (amplitude) layer is fully compatible with a 2D commitment ledger at the substrate layer.

Compatibility statement. The programme distinguishes:

- *Reversible amplitude dynamics* — an effective continuum description (unitary evolution, wave propagation) in which Lorentz invariance can hold exactly or to arbitrarily high precision.
- *Causal ledger dynamics* — the substrate layer of fact-commitment and record propagation, constrained by admissibility and finite distinguishability, whose primitive propagation arena is a 2D interface.

Lorentz invariance is a symmetry of the effective event manifold and amplitude equations (Layer 1), not a statement about the microscopic architecture of the commitment ledger (Layer 2). The two are compatible provided the ledger admits an isotropic universality class in the IR and the mapping from ledger to effective events preserves cone structure.

L.2 Two Notions of Causal Structure

A key source of confusion is the conflation of:

1. **Causal influence cone in the effective continuum description** — the Lorentz cone, a feature of the amplitude layer.
2. **Record-causal ancestry in the commitment ledger** — which events are committed and which commitments are ancestors of which, a feature of the causal ledger layer.

In the record-theoretic derivation of Lorentz structure (the Alexandrov–Zeeman route), Lorentz symmetry emerges from cone preservation applied to admissible observers on the event manifold — from requirements on record-causal order and observer comparability. Lorentz symmetry is therefore properly a symmetry of the event manifold of committed records once such a manifold exists, not a statement that the microscopic ledger must itself be a 3D Euclidean bulk.

L.3 How Lorentz Symmetry Can Emerge From a Discrete Substrate

L.3.1 Continuum Lorentz symmetry from isotropic finite-speed dynamics.

A discrete, local, finite-speed update rule with no preferred directions can yield the continuum wave equation and Lorentz symmetry in the long-wavelength limit. This is the standard emergence logic: discrete microstructure → continuum PDE with isotropic cone. Lorentz invariance is the continuum symmetry of isotropic finite-speed interface dynamics, with discrete residue suppressed at long wavelengths. Separately, Lorentz structure can be derived from cone preservation and observer admissibility conditions, emphasizing that Lorentz transformations are forced once causal order is preserved on an appropriate-dimensional event manifold.

L.3.2 Why a 2D ledger is not in conflict with a 3D Lorentzian effective world.

The commitment ledger is not "space" in the naive sense; it is the arena where committed distinguishability propagates and becomes causal ancestry. The effective 3D world is reconstructed from interface correlations and coarse-graining depth (a scale index). The compatibility claim is:

- Lorentz symmetry constrains the event-level causal cone and observer transformations.
- The ledger supplies the discrete causal update substrate that generates the event cone in the IR.

This is the same conceptual separation already present in the ticks-vs-bits architecture and the admissibility programme: reversible between commitments, irreversible at the commitment boundary.

L.4 Residual Lorentz Violation and Why It Can Be Tiny

A discrete interface substrate generically breaks continuous rotation symmetry down to a point group. In the hexagonal case, the symmetry group is C_{6v} (the dihedral group of order 12, combining six-fold rotation with reflection). The key group-theoretic fact is that any rotation-invariant observable on a C_{6v} lattice must transform as a scalar representation of C_{6v}. The lowest-order angular harmonics $\cos(n\theta)$ compatible with C_{6v} symmetry are those with n divisible by 6 — i.e., $\cos(6\theta)$, $\cos(12\theta)$, ... This is because C_{6v} includes a $\pi/3$ rotation, which maps $\theta \to \theta + \pi/3$ and thus maps $\cos(n\theta) \to \cos(n(\theta + \pi/3))$: invariance requires $n\cdot(\pi/3) = 2\pi k$ for integer k, i.e. n divisible by 6. Terms with $n = 2$ or $n = 4$ (which would arise from square C_{4v} or rectangular symmetry) are forbidden by the hexagonal point group. The leading anisotropy is therefore $\cos(6\theta)$, not $\cos(2\theta)$ or $\cos(4\theta)$.

The suppression order $p \geq 4$ follows from a renormalization group argument: the dispersion relation near the isotropy-restoration fixed point takes the form $\omega^2 = c^2k^2(1 + \alpha(ka)^2 + \beta(ka)^4$

$\cos(6\theta) + \cdots$). Terms odd in ka vanish by time-reversal symmetry. The first isotropic correction is $O((ka)^2)$; the first anisotropic correction compatible with C_{6v} enters at $O((ka)^4)$, giving $p = 4$ as the minimal suppression order. Higher-order terms and RG running can increase p further, but $p < 4$ is forbidden by the combined constraints of C_{6v} symmetry and time-reversal invariance.

$$\delta c/c(\theta) = \varepsilon(ka)^p \cos(6\theta) + \cdots, \quad p \geq 4$$

This structure has two important implications:

Lorentz invariance can hold to arbitrarily high precision at accessible scales. For long wavelengths $ka \ll 1$, the anisotropic correction is parametrically small.

The residual is symmetry-protected in harmonic order. A companion phenomenology paper proposes a 6Ω discriminator (vs 4Ω for a square lattice) as a robust observable even when ε is very small. Under that framework, null results in Lorentz-violation experiments constrain the parameter combination εa^p — not the logical possibility of Lorentz emergence.

L.5 Compatibility With High-Precision Constraints

A frequent objection invokes multimessenger bounds (GW170817-style) constraining $|c\_GW - c\_\gamma|/c$ to parts in 10^{15}. In the framework:

The effective cone speed is the emergent invariant speed measured by rods and clocks built from the same substrate, and is shared across admissible observers. Residual lattice anisotropy affects direction-dependent corrections to propagation, but in the hexagonal model those corrections are symmetry-suppressed and lie below current sensitivity for sufficiently small lattice scale a and/or sufficiently large suppression order p.

Crucially, the framework does not require a large Lorentz violation; a companion phenomenology paper predicts a specific harmonic structure of any residual anisotropy. Isotropy-violation searches should look for 6Ω angular dependence rather than a uniform shift in propagation speed.

L.6 Compatibility With Unitary Quantum Dynamics

A referee might argue: "If commitments are fundamental, is unitarity lost?"

No. The admissibility architecture explicitly separates:

- *Reversible evolution between commitments* — governed by a unitary group with a self-adjoint generator (Stone–Wigner structure). Any admissible reversible evolution must be unitary: this is a theorem of the admissibility framework, not an additional postulate.

- *Irreversible commitment processes* — operational constraint at the causal ledger layer, modelled as a CPTP semigroup structure at the boundary.

Irreversible commitment is an operational constraint layer, not a denial of unitary microdynamics. Lorentz-compatible QFT remains the correct effective description of reversible dynamics. The causal ledger describes how facts become operationally non-recombinable at Σ.

L.7 Falsifiers

This appendix commits to concrete falsification criteria:

F1 (Wrong harmonic order). Detection of a clean 4Ω angular anisotropy without a 6Ω component would falsify hexagonal substrate selection (indicating a square lattice instead).

F2 (No IR Lorentz universality). Observation of unsuppressed low-order anisotropy that does not RG-decay with wavelength or scale would falsify the "emergent Lorentz in IR" claim. This would require an alternative substrate architecture with a different universality class.

F3 (Observer comparability failure). Any empirical demonstration that causal order depends on apparatus choice in a way not removable by admissible equivalence transformations would violate the assumptions of the Causal Ledger Theorem — and would simultaneously undermine standard physics practice as record comparison.

L.8 Summary

Lorentz invariance is a symmetry of the effective event manifold derivable from cone preservation and admissible observer transformations. A 2D commitment ledger is compatible with Lorentz invariance because the ledger is not a 3D bulk: it is the substrate for fact-commitment whose IR universality class generates an isotropic cone and therefore Lorentz structure at the event level. Residual Lorentz violation from discrete microgeometry is symmetry-controlled and RG-suppressed, with a leading 6Ω signature in the hexagonal case detailed in a companion phenomenology paper. Unitary dynamics and Hamiltonian structure remain intact as the reversible layer between commitments.

References

- Pearl, J. *Causality: Models, Reasoning, and Inference.* Cambridge University Press, 2000. (Counterfactual causation, do-calculus, intervention semantics.)

- Zurek, W. H. "Decoherence, einselection, and the quantum origins of the classical." *Reviews of Modern Physics* 75, 715 (2003). (Pointer states, stable records, open-system classicality.)
- Haag, R. *Local Quantum Physics: Fields, Particles, Algebras.* Springer, 1996. (Algebraic QFT, local operator algebras, microcausality — the framework addressed by Causal Ledger Theorem Lemma 3.)
- Spekkens, R. W. "Evidence for the epistemic view of quantum states: A toy theory." *Physical Review A* 75, 032110 (2007). (Operational constraints on physical ontology.)
- Janzing, D. et al. "Thermodynamic cost of reliability and low temperatures: Tightening Landauer's principle and the second law." *International Journal of Theoretical Physics* 39, 2217 (2000). (Thermodynamic grounding of stable record formation.)
- Woodward, J. *Making Things Happen: A Theory of Causal Explanation.* Oxford University Press, 2003. (Interventionist causation; grounding of counterfactuals in physical manipulability.)
- Butterfield, J. & Isham, C. J. "Spacetime and the philosophical challenge of quantum gravity." In *Physics Meets Philosophy at the Planck Scale*, eds. Callender & Huggett, Cambridge University Press, 2001. (The problem of time in canonical quantum gravity; why background causal structure is unavailable in the fundamental description — directly relevant to the pre-geometric scope of this paper.)
- Chiribella, G., D'Ariano, G. M. & Perinotti, P. "Informational derivation of quantum theory." *Physical Review A* 84, 012311 (2011). (Derivation of quantum mechanics from operational/informational axioms; a paradigm case of the pre-geometric programme to which this paper's results apply.)
- Oreshkov, O., Costa, F. & Brukner, Č. "Quantum correlations with no causal order." *Nature Communications* 3, 1092 (2012). (Process matrices and the quantum switch; demonstrates observable causal indefiniteness, providing a natural test case for whether commitment structure can ground causal order when quantum superpositions of causal order are present.)