Proto-Time and Emergent Lorentz Invariance from Irreversible Commitment: A Structural Framework and Comparative Analysis

Keith Taylor VERSF Theoretical Physics Programme

For the General Reader

Physics has two different relationships with time, and they contradict each other.

In quantum mechanics, time is a fixed backdrop — a stage on which particles move, but not itself part of the action. In general relativity, time is a dimension of spacetime, bent by mass and energy but still taken for granted as something that simply exists. When physicists try to combine the two theories, the contradiction surfaces starkly: the central equation of quantum gravity, the Wheeler-DeWitt equation, contains no time at all. The universe it describes is frozen. Yet here we are, experiencing one moment after another.

This paper proposes a different starting point. Rather than asking how time fits into physics, it asks: *what generates time?* The VERSF framework answers with a three-layer picture. At the deepest level, microscopic evolution proceeds through reversible steps called ticks — a bookkeeping of how distinguishability accumulates in the universe's fabric, governed by the Ticks-Per-Bit framework. These ticks are entirely reversible: they describe the formation of potential distinctions, not actual facts. Above this sits proto-time: a global ordering parameter that tracks the reversible quantum evolution of the entire universe. Neither ticks nor proto-time is physical time. Physical time — the kind clocks measure and we experience — emerges only from a third layer: *irreversible commitment events*, the moments at which an accumulated distinction crosses a threshold, stabilises as a permanent record, and cannot be undone. Every such event generates a quantum of physical time at that location. Time is not a background. It is a cumulative product of irreversible facts.

This creates an immediate puzzle. If time is generated by events in a sequence, what determines that sequence? Proto-time provides the ordering — but it is not itself measurable. No clock can read it. No experiment can detect it.

The central result of this paper is that proto-time's ordering of events is, in a precise mathematical sense, *outside the observable algebra of physics*. Two histories that differ only in which spacelike-separated event happened first in proto-time are completely identical in everything measurable — the same records, the same causal structure, the same distances between events. They are physically the same history. This is the Commitment Reordering

Equivalence: proto-time ordering is to physics what gauge redundancy is to electromagnetism — a freedom in the description that leaves no imprint on anything measurable. And the symmetry group that emerges from this — the group of transformations that preserve what is measurable — is the Lorentz group. Relativistic covariance is not assumed; it is derived from the structure of irreversible records.

The paper also shows that the committed-record sector is locally Minkowskian — not because Minkowski geometry is inserted as an assumption, but because it is the minimal structure forced by four operational requirements: finite-speed signal propagation, radar-definable intervals, flat-limit isotropy, and independence from proto-temporal ordering. Minkowski geometry is the weakest local completion compatible with these requirements.

The paper then compares this framework with the three leading approaches to quantum gravity and time — loop quantum gravity, causal set theory, and relational quantum mechanics — identifying where VERSF agrees, where it differs, and why the differences matter. It closes with two experimental predictions: a suppression of decoherence in superconducting qubits at temperatures below ten millikelvin, and a precise numerical relationship ($\sigma_\tau/\sigma_{peak} = \sqrt{2 \ln 2} \approx 1.18$) between the temporal and optical observable signatures of a localised region of suppressed record-formation capacity — a ratio that holds at leading order for any smooth, single-scale profile of such a region, not just for one specific functional form. Both predictions are testable with current technology.

Table of Contents

Abstract

The problem of time in fundamental physics has two faces: quantum mechanics treats time as an external parameter, while the Wheeler-DeWitt equation of canonical quantum gravity contains no time at all. The VERSF (Void Energy-Regulated Space Framework) programme approaches this not by quantising gravity but by deriving both gravity and time from a more primitive ontology (as detailed elsewhere in the programme [3, 8]): a three-layer temporal structure in which microscopic reversible tick-dynamics (TPB framework [1]) generate distinguishability, proto-time orders the global reversible quantum evolution of possibilities, and physical time emerges locally only from the subset of tick-transitions that cross the commitment threshold and stabilise as irreversible records. The commitment density $\lambda(s) = \Sigma(s)/(k_B \ln 2)$ governs the rate of this third layer, and physical time accumulates as $t(s) = \int_0^s \lambda(s') \, ds'$. Reversibility lives at the tick level; irreversibility is introduced precisely and only at the commitment threshold.

This structure raises an immediate tension: proto-time imposes a global ordering on commitment events, whereas relativistic spacetime admits no preferred foliation. We resolve this tension through a quotient construction. Two proto-temporal histories are defined as physically indistinguishable if they produce identical committed records, identical causal relations, and identical invariant signal-cone structure. Observable physics is defined on the quotient under this equivalence, and we show that all observables on the quotient are Lorentz invariant — proto-time ordering is an unobservable degree of freedom, structurally analogous to gauge redundancy. The committed-record sector is shown to be locally Minkowskian not by postulate but as the minimal flat operational completion compatible with finite-speed signal propagation, radar-definable intervals, and flat-limit isotropy.

The equivalence relation is not introduced ad hoc. It is motivated and structurally necessitated by the Pre-Factual Algebraic Reversibility condition (PAR): before any physical fact is formed, no compositional transition in the space of possibilities may be observationally irrecoverable, since an irrecoverable transition would introduce irreversibility without a fact to ground it. Three independent derivation routes — from admissibility logic, information-theoretic reasoning, and Landauer-type thermodynamics — converge on PAR; the formal derivation that CRE is PAR's observable-sector expression is established in companion work [5]. Lorentz invariance is its geometrical corollary, given the Minkowskian causal structure of committed records.

We compare VERSF with loop quantum gravity, causal set theory, relational quantum mechanics, and the Connes-Rovelli thermal time hypothesis, identifying structural distinctions in each case. We present two quantitative predictions: (1) a decoherence rate anomaly in superconducting qubits at sub-10 mK temperatures — the existence and direction of the anomaly (suppression below the Lindblad value) are parameter-free predictions, while the scale of onset depends on the empirically constrained commitment threshold λ_c; and (2) a cross-channel scale ratio $\sigma_\tau / \sigma_{peak} = \sqrt{2 \ln 2} \approx 1.18$, where σ_{peak} is the peak-deflection impact parameter of the optical channel, holding at leading order for any smooth, isotropic, single-scale commitment-density profile with controlled corrections from higher-order profile moments. These results establish VERSF as a structurally distinctive and experimentally discriminable approach to the emergence of spacetime structure and physical time from irreversible commitment.

---

1. Introduction

The problem of time in fundamental physics has two faces. In quantum mechanics, time enters as a fixed external parameter in the Schrödinger equation — a background against which systems evolve, not a property of those systems. In general relativity, time is a dimension of the spacetime manifold, bent and dilated by mass-energy but still presupposed as a structure within which physics occurs. When the two theories are combined, the problem becomes acute: the Wheeler-DeWitt equation, the central equation of canonical quantum gravity, contains no time at all. The universe it describes is static. Yet we experience relentless temporal progression, memory of the past, and causally structured physical events.

Existing approaches address this in structurally different ways. Loop quantum gravity discretises geometry and seeks to recover time from relational observables among spin-network excitations. Causal set theory takes the discrete partial order of spacetime events as primitive and derives continuous spacetime from it by random sprinkling. Relational quantum mechanics treats temporal facts as observer-dependent and seeks to recover consistency among observers without a global time. Each approach makes progress; none provides a complete account of how an irreversible, directed, locally metered time arises from a time-symmetric fundamental description.

This paper presents and defends the VERSF approach to this problem, which differs from all three in a specific structural way: it identifies **irreversible commitment** — the transition from reversible quantum possibility to permanent physical record — as the primitive from which both time and physical reality are constructed. The novelty is not the idea that time is emergent (this is shared with the programmes above) but the mechanism: time is not emergent from geometry, from causal order, or from relational observation. It is emergent from the local rate of irreversible fact formation.

The central technical contribution of this paper is the Commitment Reordering Equivalence (CRE) principle and its consequence: a demonstration that Lorentz invariance emerges at the level of observable records, given the Minkowskian causal structure of committed records, even though the underlying proto-temporal ordering is not Lorentz covariant. Proto-time is shown to be an unobservable degree of freedom — the temporal analogue of gauge redundancy. Crucially, CRE is not introduced to solve the Lorentz problem; it is the observable-sector expression of the Pre-Factual Algebraic Reversibility condition that is independently motivated within the VERSF commitment architecture. The formal proof of this correspondence lives in companion work [5]; what this paper establishes is the structural argument and its physical motivation.

We then provide a systematic comparison with loop quantum gravity, causal set theory, and relational quantum mechanics, followed by quantitative experimental predictions that distinguish VERSF from each.

---

2. The Two-Layer Temporal Ontology of VERSF

2.1 Proto-Time

VERSF introduces a pre-physical ordering parameter $s \in \mathbb{R}$ with a total order $\prec$. A quantum system evolving under Hamiltonian $\hat{H}$ without irreversible environmental coupling is parametrised by s:

$$i\hbar\, \partial|\psi(s)\rangle/\partial s = \hat{H}|\psi(s)\rangle$$

We call s **proto-time**. It provides a relational ordering of quantum states — $|\psi(s_2)\rangle$ follows $|\psi(s_1)\rangle$ if $s_1 \prec s_2$ — but is formally distinguished from physical time by three criteria:

(C1) Non-measurability. Proto-time s does not appear in the algebra of observables of any physical subsystem. No measurement carried out by any internal clock yields s as a value. It is a parameter of the global wavefunction, not a locally accessible quantity. The reason is structural: s parametrises the evolution of the entire closed system $|\Psi(s)\rangle$, whereas any measurement is an interaction between subsystems — an operation in the local observable algebra of some proper part of the universe. A quantity defined only at the level of the global wavefunction is simply not a member of any subsystem's observable algebra, by the definition of what a subsystem's observable algebra contains. No Hermitian operator on any local Hilbert space corresponds to the global evolution parameter s.

(C2) Global definition. Proto-time is defined with respect to the evolution of the entire closed system |Ψ(s)⟩. Physical time, by contrast, is always locally defined — different regions accumulate it at different rates.

(C3) No conjugate observable. There is no physical-time conjugate to the system Hamiltonian in the usual sense. Physical time t is an emergent local quantity, not a global canonical parameter.

These are not merely terminological distinctions. They have formal consequences for what can and cannot be measured, and they are the structural basis for the resolution of the Lorentz tension developed in Section 4.

2.2 Physical Time

Physical time t is not a background. It is generated locally and incrementally by **commitment events**: the moments at which a quantum possibility becomes a definite physical fact by forming an irreversible record in the environment. The local rate of physical time accumulation is the commitment density λ(s):

$$t(s) = \int_0^s \lambda(s')\, ds'$$

The commitment density is grounded in thermodynamics:

$$\lambda(s) = \Sigma(s) / (k\_B \ln 2)$$

where Σ(s) is the local entropy production rate with respect to proto-time. No free parameters appear. Every quantity is either a universal constant (k_B, ln 2) or a measurable physical observable (Σ).

An important clarification for what follows: Σ is coarse-grained entropy production, defined relative to physically accessible macrostates — not microscopic entropy, which is conserved under unitary evolution. The definition is therefore effective rather than fundamental, but it is operationally measurable and aligns physical time with thermodynamic irreversibility rather than microscopic reversibility. This is the appropriate level of description for commitment events, which are macroscopic record-formation processes. A related concern — that entropy production depends on the choice of coarse-graining, making time observer-dependent — is addressed directly: while the numerical value of Σ depends on coarse-graining, the existence of irreversible record formation is invariant under changes of coarse-graining that preserve macroscopic distinguishability. Physical time is therefore tied to the existence of irreversible records, not to the precise numerical value of entropy assigned to them. Two observers using different coarse-grainings will agree on whether a commitment event has occurred, even if they assign different entropy values to the process. In particular, any coarse-graining that preserves stable macroscopic records yields the same commitment structure, ensuring that physical time is invariant under admissible changes of description. One tension is worth acknowledging explicitly: the coarse-graining invariance argument holds cleanly in the regime λ ≫ λ_c, where commitment events are well-defined macroscopic record-formation processes. Near the threshold

$\lambda \approx \lambda_c$ — precisely the regime targeted by the decoherence prediction of Section 7.1 — the boundary between coarse-graining-invariant commitment and coarse-graining-dependent apparent record-formation becomes less sharp, and the invariance argument requires more careful treatment. This is therefore an open issue directly connected to the paper's primary experimental prediction: the prediction in Section 7.1 relies on the commitment threshold being well-defined even near λ_c, and the coarse-graining analysis in that regime is an acknowledged open theoretical task.

When $\lambda(s) = 0$ — in a perfectly isolated quantum system undergoing unitary evolution — no physical time accumulates. The system exists in proto-time but contributes nothing to any local physical clock. When $\lambda(s) > 0$ — when irreversible interactions occur — physical time is generated at that location, at a rate proportional to the local entropy production.

2.3 Coexistence and the Wheeler-DeWitt Resolution

The Wheeler-DeWitt equation $\hat{H}|\Psi\rangle = 0$ describes a universe with no external time. In the VERSF framework, this is not a paradox but an accurate description of the global proto-time substrate: the global wavefunction $|\Psi(s)\rangle$ has no physical-time parameter because there is nothing external to the universe to record its evolution. Proto-time s advances everywhere, including regions where $\lambda = 0$ and no physical time accumulates. Local physical times t(x) emerge within the universe wherever subsystems interact irreversibly, at rates set by local entropy production. The frozen universe of Wheeler-DeWitt is the correct description of the proto-time substrate; the experienced time of physics is the physical-time layer growing on top of it.

The identification — that $|\Psi(s)\rangle$ satisfies the WdW constraint — is treated fully in [8], but its logical structure is as follows. The VERSF proto-time Hamiltonian $\hat{H}_{VERSF}$ is constructed as the generator of reversible global evolution in the absence of commitment events. Since commitment events are defined precisely as the transitions at which $\lambda > 0$ begins — when the system transitions from global reversible evolution to local irreversible record formation — the global wavefunction in the absence of all commitment events evolves as a timeless superposition. The WdW constraint $\hat{H}|\Psi\rangle = 0$ then holds because $|\Psi(s)\rangle$ is the global state from which physical times are generated by commitment, not a state that evolves in any physical time. The physical content of WdW in VERSF is therefore: global proto-time evolution is timeless from the inside because physical time is entirely a product of the commitment events that happen within it. At this stage the WdW correspondence is presented as structural motivation for the two-layer ontology; the derivation that $\hat{H}_{VERSF}$ satisfies the constraint is given in [8].

This provides a natural structural account of the measurement problem: quantum measurement is not a collapse imposed from outside the formalism, but the generation of physical time — a commitment event in which λ transitions from zero to positive at a specific proto-time value s^*. The first-passage comparison across competing branches that determines which outcome becomes the Bit is defined with respect to proto-time, the global ordering parameter of reversible evolution (Section 2.4). Whether this account fully resolves the measurement problem, or whether additional dynamical specification is required, is a question that depends on the completeness of the commitment formation theory and is not claimed to be settled here.

2.4 Proto-Time and the Ordering of Tick Events

The Ticks-Per-Bit (TPB) mechanism [1] requires a well-defined ordering of candidate microscopic events across decohered branches in order to formulate a first-passage selection rule. In particular, the statement that an outcome is determined by the first threshold-crossing tick presupposes a common ordering relation under which events generated in different branches can be compared. Proto-time s provides precisely this: a global ordering parameter for reversible evolution that carries no metrical structure and corresponds to no clock reading, but supplies a total ordering relation over candidate micro-events in the pre-factual domain. Its role is purely structural — to make "earlier" and "later" well-defined in the pre-factual domain where no irreversible records yet exist. The existence of such an ordering is not an additional dynamical assumption but a structural requirement for defining first-passage processes in the absence of physical time: without a shared ordering across branches, the argmin of Section 2.4 below is undefined.

Ticks as embedded events. A tick is a local, reversible increment of distinguishability within a branch's detector–environment microstate. Ticks are generated by branch-specific dynamics and are therefore local processes, but they do not define their own ordering. Instead, each candidate tick event is assigned a proto-time label s, inherited from the underlying reversible evolution. Let $\mathscr{E} = \{e_i^{(A)}\}$ denote the set of candidate tick events generated across branches A. Each event $e_i^{(A)}$ is associated with a proto-time value $s_i^{(A)}$. Proto-time induces a total ordering over candidate micro-events in the pre-factual domain, while the observable record sector retains only the Lorentzian partial order after quotienting under CRE. This is not a contradiction but the paper's central structural result: the total pre-factual ordering is real but unobservable, and the partial causal order of committed records is what survives into physics.

First-passage selection. Outcome selection is defined by a first-passage rule over proto-time:

$$A_{out} = \mathrm{argmin}_A \{ s_{th}^{(A)} \}$$

where $s_{th}^{(A)}$ is the proto-time at which branch A first produces a threshold-crossing tick — the event that becomes a Bit. This definition is well-posed because proto-time provides a shared ordering across all branches. In continuous systems, the probability that two distinct branches produce threshold-crossing ticks at exactly the same proto-time value is measure zero, so the argmin is generically unique; exact degeneracy is an acknowledged edge case whose tie-breaking behaviour is deferred to the dynamical completion of the framework [8]. Ticks themselves remain reversible micro-events; irreversibility enters only at the level of threshold crossing. The Bit — the first threshold-crossing tick — constitutes a macroscopic, thermodynamically irreversible record and is precisely the commitment event of Section 2.2.

Relation to TPB. In the TPB framework, microscopic dynamics are parametrised by a tick-ordering variable λ, which is explicitly not physical time. The present construction identifies λ as a local manifestation of proto-time within a subsystem: TPB dynamics sample the global proto-time ordering through local microdynamic processes. The commitment density λ(s) of Section 2.2 therefore represents not the rate of ticks themselves, but the rate at which tick-driven threshold crossings occur — the rate at which the first-passage selection rule fires and generates irreversible records.

This yields a consistent four-level hierarchy:

- **Proto-time s:** global ordering of reversible evolution (unobservable, no metrical structure)
- **Ticks (λ-ordered events):** local, reversible micro-events embedded in proto-time
- **Bits (threshold crossings):** irreversible commitment events generating records
- **Physical time t:** accumulation of Bits, $t(s) = \int_0^s \lambda(s')\, ds'$

Reversibility lives at the tick level; irreversibility is introduced precisely and only at the commitment threshold. This resolves the apparent tension between TPB dynamics (reversible) and physical time (irreversible): they operate at different layers of the same structure, connected by the threshold condition λ_c.

Conceptual note. Proto-time should not be interpreted as a hidden physical clock or background time parameter. It is the minimal ordering relation required to define first-passage selection among competing tick processes. Ticks populate this ordering with reversible micro-events; Bits select a subset as irreversible records; and physical time emerges only from the accumulation of those records.

---

3. The Proto-Time / Lorentz Tension

In VERSF, irreversible commitment events are ordered in proto-time, forming a sequence:

$C_1 \prec C_2 \prec C_3 \prec \cdots$

At first sight, this global ordering appears to introduce a preferred temporal sequence incompatible with relativistic covariance. This tension is the central problem the paper resolves; the resolution — that proto-time ordering is physically unobservable by construction — is established in Section 4. For the present section, we characterise the tension precisely. Relativistic covariance requires:

- There is no universal simultaneity
- Spacetime is invariant under Lorentz transformations
- No preferred foliation exists

If proto-time were directly observable — if any experiment could determine the proto-temporal order of spacelike-separated commitment events — it would define a preferred frame and contradict the empirical success of special relativity. The resolution requires proving that proto-time ordering is physically unobservable, so that no experiment can access it and no preferred frame is defined at the level of measurable physics.

4. Commitment Reordering Equivalence and the Emergence of Lorentz Invariance

4.1 The Equivalence Principle

Definition (Commitment Reordering Equivalence). Let $\mathscr{H}$_proto denote the space of proto-temporal histories — complete specifications of the sequence of commitment events and their proto-temporal ordering. Define an equivalence relation ~ on $\mathscr{H}$_proto by:

$H_1 \sim H_2 \Leftrightarrow \mathscr{R}(H_1) = \mathscr{R}(H_2)$

where $\mathscr{R}(H)$ is the **committed record structure** of history H, defined by three jointly necessary conditions:

1. **Same committed records:** the set of irreversible commitment events and their physical content are identical.
2. **Same causal relations:** the partial order among commitment events — which events are in the causal past or future of which others — is identical.
3. **Same signal-cone structure:** the invariant spacetime intervals between commitment events are identical.

Physical history is then defined as the quotient:

$\mathscr{H}$_phys = $\mathscr{H}$_proto / ~

Formal definitions of $\mathscr{H}$_proto, $\mathscr{R}$, the equivalence relation ~, and the observable algebra $\mathcal{O}$ are collected in Appendix A for readers wishing to verify the mathematical structure of this construction.

4.1A Why the Committed-Record Sector is Locally Minkowskian

The Lorentz-emergence argument of Section 4.2 requires that the committed-record sector admit a causal and metrical structure on which automorphisms can act. This structure is not introduced as an arbitrary metric ansatz. Rather, it is the minimal operational completion of the record sector under four assumptions:

(i) **Finite-speed signal propagation.** Commitment events are connected by physical signals that travel at or below a finite limiting speed. This is the only ingredient needed to define a causal partial order on commitment events — which can signal which.

(ii) **Radar-definable intervals.** Coincidence, duration, and distance are definable by local clock-and-signal procedures: two observers exchange light signals, record emission and reception times on their own clocks, and derive interval assignments from those readings. This is the operational content of the invariant-interval notion — it requires no prior global coordinate system.

(iii) **Homogeneity and isotropy in the flat limit.** In the limit of vanishing fold-density gradients (the source-free, flat-field regime), the committed-record sector has no preferred location, direction, or frame among inertial observers. This is the flatness condition — it excludes curvature in the source-free limit without assuming curvature is absent in general.

(iv) **Independence from proto-temporal ordering.** No observable depends on the proto-temporal ordering of commitment events; the record sector is defined entirely by what is measurable, not by the global evolution parameter s.

Under assumptions (i)–(iii), any pair of inertial observers reconstructing intervals from clock-and-signal procedures must recover the same limiting signal cone and the same quadratic invariant compatible with it. The unique local quadratic form satisfying this requirement — homogeneous, isotropic, compatible with a finite invariant signal speed — is the Minkowski quadratic form. The committed-record sector is therefore locally Minkowskian in the flat limit: not as a postulate, but as the minimal structure forced by the operational procedures already required to define committed records and their causal relations.

The present paper does not derive Minkowski geometry from first principles, nor does it insert it as an unconstrained metric postulate. It adopts Minkowski structure as the minimal flat operational completion of the committed-record sector, conditional on the c-identification (Sections 4.2 and 8.3) and subject to coarse-grained deformation by fold-density gradients, which generate the curved spacetime of the companion gravity derivation [3]. Curved spacetime in VERSF is therefore the deformation of this locally Minkowskian record geometry — not a separate input.

4.2 The Lorentz Emergence Result

Result (Lorentz Invariance from CRE). *All observables defined on $\mathscr{H}$_phys are invariant under transformations that preserve the causal structure and invariant intervals of committed records.*

Derivation. Let O be any physical observable — a quantity measurable by experiments carried out within the universe. By criterion (C1), O cannot depend on the proto-temporal ordering s directly, since s is not in the observable algebra of any subsystem. Therefore O is a function on $\mathscr{H}$_phys = $\mathscr{H}$_proto/∼, not on $\mathscr{H}$_proto itself.

It remains to show that $\mathscr{H}$_phys is Lorentz invariant. Let Λ be a Lorentz transformation. Λ acts on commitment events by relabelling their spacetime coordinates while preserving:

- Causal relations among events (since Lorentz transformations preserve lightcone structure)
- Spacetime intervals between events (by definition of Lorentz invariance)

- The physical content of each commitment event (which depends on local records, not global coordinates)

Therefore $\mathscr{R}(\Lambda \cdot H) = \mathscr{R}(H)$ for any proto-temporal history H, meaning $\Lambda \cdot H \sim H$. The action of Λ on $\mathscr{H}$_proto descends to the identity on $\mathscr{H}$_phys. Since all observables are functions on $\mathscr{H}$_phys, all observables are invariant under Λ. □

Corollary. No experiment can determine the proto-temporal ordering of spacelike-separated commitment events. Proto-time ordering is an unobservable degree of freedom — the temporal analogue of gauge redundancy.

On what the result establishes and what it leaves open. The Lorentz Emergence Result is exact at the level of committed records: all observables defined on $\mathscr{H}$_phys are exactly Lorentz invariant as a consequence of the CRE quotient construction. A separate, open question — addressed in Section 8.3 — is whether the underlying fold-void interface dynamics that generates commitment events is itself exactly Lorentz covariant, or whether it introduces corrections at the commitment scale that are invisible at the level of committed records but could, in principle, generate observable departures at extreme energies. These two questions are logically independent. The result does not require fold-level covariance to hold; it establishes Lorentz invariance as a property of the observable sector regardless of the fold-level structure. Resolution of this open problem could generate additional predictions — commitment-scale Lorentz-departing effects at extreme energies — as a further discriminator between VERSF and classical relativity, rather than threatening the result's conclusions. The relevant energy scale is not the Planck scale of quantum gravity programmes but the characteristic scale at which fold-density fluctuations become non-negligible, which is a derived quantity of the VERSF fold dynamics.

On circularity. A careful reader will notice that the proof invokes Lorentz transformations as structure-preserving maps. Does this assume Lorentz invariance in order to derive it?

The answer is no, but the distinction requires precision. The committed record structure $\mathscr{R}(H)$ is defined entirely in terms of: (i) the identity and physical content of commitment events; (ii) their causal precedence relations — which events can send signals to which; and (iii) the invariant proper intervals between events. None of these presupposes a specific symmetry group. They are operationally defined: determined by local measurements using local clocks and local signal exchanges, without any prior commitment to global spacetime symmetry.

Given this operationally grounded definition, we ask: what is the maximal group of bijections on the set of commitment events that preserves $\mathscr{R}$? This is a question about the automorphism group of the causal and metrical structure of committed records. A careful statement is needed here. Zeeman [22] shows that the causal automorphisms of Minkowski spacetime — bijections preserving lightcone precedence — form the orthochronous Lorentz group augmented by translations and dilations, i.e. a subgroup of the conformal group. Alexandrov [9] establishes the related result that any bijection of $\mathbb{R}^4$ preserving the causal order is in this group. Neither result by itself isolates the Lorentz group alone. The additional ingredient that fixes the scale and excludes dilations is the invariant interval structure: the committed record map $\mathscr{R}$ encodes not only causal precedence but also definite invariant spacetime intervals $d(C_i, C_j) = \Delta s^2_{ij}$ between each pair of events. A dilation scales all intervals by a constant factor, which changes Δs^2 and therefore does not preserve $\mathscr{R}$. Translations shift the coordinates of all events but, in the absence of a fixed origin, do not affect the relative intervals or causal relations — they descend to the identity on the equivalence classes in $\mathscr{H}$_phys. The automorphism group of the combined structure — causal precedence together with fixed invariant intervals — is therefore the Lorentz group (or more precisely the Poincaré group, once translations are accounted for), not the full conformal group. The committed record structure, defined by both causal precedence and invariant intervals, inherits this restricted automorphism group. One further precision: Zeeman's and Alexandrov's theorems are theorems about bijections of Minkowski spacetime — they work by taking Minkowski causal structure as their domain. Invoking them here therefore imports Minkowski structure as a prerequisite rather than deriving it. The correct reading of this argument is that the Zeeman-Alexandrov results confirm consistency: given the (assumed) Minkowskian causal structure of committed records, the symmetry group of the observable quotient is the Lorentz group. They do not constitute an independent derivation of Minkowski structure from commitment events alone; that derivation is the subject of the c-identification open problem in Section 8.3. The explicit conditional statement of this formal result is given in Appendix A, in the Remark following Definition A.4.

The remaining implicit step is the identification of the causal structure of committed records with that of Minkowski spacetime. This is not stipulated. In VERSF, commitment events are generated by irreversible interactions that propagate at or below the speed of light; their causal structure is therefore Minkowskian in the flat limit, as established in Section 4.1A. There is a further identification that the present argument relies upon but does not derive: the c that bounds commitment-event signal propagation and the c that appears in the Minkowski interval must be the same constant. The identification of the signal-propagation bound with the metric light speed is not an arbitrary assumption, nor merely a consistency condition — it is forced by operational closure. If the signal cone defined by commitment interactions and the interval cone defined by metrical reconstruction from clock-and-signal procedures differed, the same pair of record events would receive incompatible causal classifications under the two procedures: timelike under one c, spacelike under the other. The record sector would then fail to admit a single observer-independent event geometry — different inertial observers using signal-propagation procedures versus metrical procedures would disagree on which events can causally influence which, making the committed record structure $\mathscr{R}$ ill-defined as a consistent triple. Conversely, when $c\_signal = c\_metric$, the causal partial order and the interval function are jointly defined on the same event set without contradiction, and $\mathscr{R}$ is well-defined by Definition A.4. The equality $c\_signal = c\_metric$ is therefore both necessary and sufficient for $\mathscr{R}$ to be well-defined as a consistent triple — required by operational coherence of the record sector, not imposed from outside it. Its deeper derivation from VERSF fold-dynamics remains open and is treated as a named open problem in Section 8.3. This is how analogous identifications are treated in GR (inertial and gravitational mass, required by the equivalence principle before being derived from the field equations) and QFT (the c in the Dirac equation and the c in the Minkowski metric, identified by Lorentz covariance of the formalism). Lorentz invariance thus emerges as the maximal symmetry of the committed record structure under this operationally required identification.

What is assumed and what is derived. It is important to be precise about the logical structure. The construction does not assume Lorentz symmetry as a starting point. It assumes only that commitment events are linked by finite-speed signal propagation and that invariant intervals are operationally definable through local measurement procedures — exchanging signals, comparing clock readings, identifying which events lie inside and outside each other's lightcones. These are empirical procedures that require no prior symmetry group. The symmetry group of the framework is then identified as the automorphism group of this operationally defined causal and metrical structure. The emergence of the Lorentz group follows from that identification — a consequence of the structure, not a premise of it. The logical direction is: operational invariants first, symmetry group second. This is precisely the direction of Zeeman's and Alexandrov's results, and it is the direction of the present construction. No assumption of global Lorentz covariance is made at the level of proto-time; covariance arises only after quotienting by CRE, and therefore cannot be used to define or constrain the underlying ordering structure.

4.3 The Gauge Analogy

The structural parallel with gauge theory is illuminating. In electrodynamics, the gauge potential A_μ is unobservable; only gauge-invariant quantities (field strengths, Wilson loops) are physical. Physical electrodynamics is defined on the quotient of the space of gauge potentials by gauge equivalence. Analogously, in VERSF, proto-temporal ordering is unobservable; only commitment-reordering-invariant quantities are physical. Physical spacetime is defined on the quotient $\mathscr{H}_{phys}$.

This analogy is not merely aesthetic. In gauge theory, gauge invariance is associated with conservation of charge via Noether's theorem — a generative relationship between redundancy and conservation. In VERSF, CRE is a quotient construction that yields Lorentz invariance as the automorphism group of the observable quotient; the structural parallel is that redundancy at the fundamental level produces symmetry at the physical level, though the precise formal relationship differs. The gauge analogy is illuminating as structural motivation rather than as a claim of formal equivalence. The deeper structure in both cases is more constrained than it appears.

4.4 CRE as the Observable-Sector Expression of PAR

A natural concern is whether CRE is an ad hoc principle introduced to rescue Lorentz invariance rather than a consequence of the framework. It is neither. CRE is the observable-sector expression of the **Pre-Factual Algebraic Reversibility (PAR)** condition: PAR governs the pre-factual domain $\mathcal{P}$, and the argument of this section shows that PAR's symmetry requirement carries forward to the observable sector, where it forces the CRE equivalence. The formal proof that this carry-forward is rigorous — establishing the exact mapping from pre-factual path orderings in $\mathcal{P}$ to proto-temporal histories in $\mathscr{H}_{proto}$ — is developed in [5, §9]. What is established here is the structural correspondence and the physical argument that motivates it; the paper claims this correspondence is correct, not that the companion paper's formal proof is reproduced within these pages.

4.4.1 PAR and Its Motivation

PAR formulation. PAR is stated in terms of observational restorability rather than group-theoretic inversion — a distinction that is both technically important and physically significant. For every non-trivial decomposition $r = r_1 \circ r_2$ in the pre-factual domain $\mathcal{P}$, where r_2 represents the prior contribution and r_1 the subsequent compositional step, PAR asserts the existence of a restoring operation $\psi_{\{r_1,r_2\}} : \mathcal{P} \rightharpoonup \mathcal{P}$ such that:

$$F((r_1 \circ r_2) \circ \psi_{\{r_1,r_2\}}) = F(r_2)$$

where $F : \mathcal{P} \to \Sigma$ is the fact map carrying pre-factual elements to their observable fact outcomes. This asserts that the observable configuration associated with r_2 — the state prior to r_1's compositional step — is recoverable after that step. It does not assume global group inversion, a global identity element, or any prior algebraic structure; it requires only the existence of a process restoring observable distinguishability. This formulation is strictly weaker than demanding that r_1 possess a specific algebraic inverse, and is the physically appropriate condition at the level of pre-factual alternatives.

The motivation is structural: an irreversible compositional step in $\mathcal{P}$ would introduce a preferred direction of transition before any physical fact has been formed. Since physical time is generated only by commitment events, a directional asymmetry in $\mathcal{P}$ prior to commitment would constitute irreversibility without a fact to ground it — a violation of the Separation of Levels principle that confines irreversibility exclusively to the commitment map φ.

The role of CC. PAR alone governs only non-trivial decompositions — those in which both components are genuinely structured pre-factual states. A second condition, **Compositional Completeness (CC)**, is required to ensure that every non-null element of $\mathcal{P}$ participates in at least one such non-trivial decomposition, bringing it within PAR's scope. CC is independently motivated as a coherence condition on domain membership: a pre-factual element that never appears in any non-trivial composition would be formally admitted yet compositionally inert — idle structure that the pre-factual sector, as the domain of fact-producing alternatives, cannot coherently contain. Together, PAR and CC imply Internal Admissible Closure (IAC): no internally realised pre-factual contribution is observationally irrecoverable. The logical structure is: CC places every non-null element within PAR's scope; PAR then guarantees the restoring operation; IAC follows universally [5, Theorems 3.1, 6.1, 7.1].

Three convergent routes to PAR. The companion papers establish PAR as a necessary consequence from three independent derivation strategies, each invoking a distinct physical concept. First, from the admissibility condition A2 together with the Pre-Factual Symmetry Condition (PFS): any irreversible transition introduces a preferred compositional direction not grounded in any formed fact, which A2 forbids. Second, from A2 together with Observational Faithfulness: any information-destroying transition eliminates observable distinguishability without a fact to ground the elimination, which A2 again forbids. Third, from the Landauer-type

Pre-Factual No-Erasure Condition (PNEC): genuine erasure of distinguishability requires a cost-bearing outlet, and no such outlet exists in $\mathcal{P}$ prior to fact formation. All three routes converge independently on the same conclusion — PAR is not an arbitrary postulate but the point at which admissibility logic, information-theoretic logic, and thermodynamic logic arrive together.

4.4.2 From PAR to CRE: Structural Correspondence

PAR's consequence for the observable sector. Suppose a physical observable O depended on the proto-temporal ordering of commitment events — say, O took different values for H_1: $A \prec B$ and H_2: $B \prec A$ where A and B are spacelike-separated. A clarification is needed here about the domain of PAR. PAR governs non-trivial compositional steps in the pre-factual domain $\mathcal{P}$ — the space of unrealised alternatives prior to commitment. Commitment events A and B are post-commitment objects in $\mathscr{H}$_proto, not elements of $\mathcal{P}$. The bridging step is the following: prior to either A or B occurring, the two proto-temporal trajectories leading to H_1 and to H_2 correspond to distinct pre-factual paths through $\mathcal{P}$ — sequences of compositional steps $r_1 \circ r_2 \circ \cdots$ leading to the respective commitment outcomes, differing only in the order of the steps corresponding to the A- and B-type interactions. These pre-factual paths are within $\mathcal{P}$'s domain by CC: every non-null pre-factual element participates in the compositional structure, so the steps leading to A and B are genuine elements of $\mathcal{P}$ and are within PAR's scope. Given this, PAR requires that the ordering difference between the two paths — the swap of which interaction step comes first — be observationally restorable, since any irrecoverable directional asymmetry between them would constitute a preferred direction prior to fact formation. The full formal argument mapping pre-factual path orderings onto $\mathscr{H}$_proto histories is developed in [5, §9]; what is established here is the structural correspondence that motivates it. An observable O distinguishing H_1 from H_2 would detect a directional asymmetry that PAR forbids at the pre-factual level. Therefore no physical observable can depend on the proto-temporal ordering of spacelike-separated commitment events.

This is the structural correspondence that motivates CRE. The equivalence relation $H_1 \sim H_2$ whenever $\mathscr{R}(H_1) = \mathscr{R}(H_2)$ is the observable-sector expression of the PAR symmetry of the pre-factual domain — the formal proof that this correspondence is rigorous, establishing the exact mapping between pre-factual path orderings and $\mathscr{H}$_proto histories, is developed in [5, §9]. Within this paper, the claim is that PAR motivates and structurally necessitates CRE; Lorentz invariance is its geometrical corollary. CRE therefore constitutes the maximal observational equivalence relation compatible with PAR and compositional completeness.

4.5 Proto-Time Is Not Hidden-Variable Time

A further objection must be addressed directly. One might suppose that proto-time is simply a hidden-variable time — an absolute Newtonian background time reintroduced under a different name, in violation of the spirit if not the letter of relativistic physics.

This objection fails for a structural reason. Hidden-variable time constructions — such as those in Bohmian mechanics or certain preferred-foliation interpretations — assign a physically meaningful, in-principle-measurable ordering to spacelike-separated events. The preferred foliation is an additional physical structure, and its existence has observable consequences: different foliations give different predictions for measurement outcomes in certain scenarios.

Proto-time in VERSF is categorically different: it does not assign observable simultaneity or measurable ordering to events. Its role is purely generative — it parametrises the unitary evolution of the global wavefunction — and all measurable temporal structure arises only after quotienting by commitment reordering equivalence. The quotient operation removes proto-time from the observable sector entirely.

A hidden-variable time is a frame; proto-time is a bookkeeping parameter that disappears from physics by construction. The difference is not semantic: hidden-variable time generates predictions that differ from standard relativity; proto-time, by the CRE result, generates predictions that are exactly Lorentz invariant.

4.6 Worked Example: Two Spacelike Detectors

Consider two spacelike-separated commitment events A and B — two detectors in distant labs each recording a particle arrival. Because A and B are spacelike separated, no physical signal connects them; neither is in the causal past of the other.

In proto-time, two distinct histories are possible:

- **History H_1:** $A \prec B$ — detector A fires first in the global proto-temporal ordering
- **History H_2:** $B \prec A$ — detector B fires first

These are genuinely different points in $\mathscr{H}$_proto. The formal record map $\mathscr{R}$ returns, for both histories, the same triple (Definition A.4): the event set {A, B} with identical physical content, the causal partial order $A \nprec B$ and $B \nprec A$ (both spacelike), and the invariant interval Δs^2_AB (spacelike and identical). Explicitly:

- **Same committed records.** Both histories contain exactly the same commitment events with identical physical content.
- **Same causal partial order.** $A \nprec B$, $B \nprec A$ — neither lies in the causal past of the other.
- **Same invariant interval.** Δs^2_AB is spacelike and numerically identical in H_1 and H_2.

Therefore $\mathscr{R}(H_1) = \mathscr{R}(H_2)$, establishing $H_1 \sim H_2$ under CRE. Now consider any observable $O \in \mathcal{O}$ — formally, any function $O : \mathscr{H}$_proto $\to \mathbb{R}$ constant on CRE equivalence classes. For this system the only physically available observable is the joint outcome:

$O(H_1) = (A=1, B=1) = O(H_2)$

The observable takes the same value on both histories, as required by CRE membership. No extension of the observable algebra — no additional local measurement, no third observer, no timing instrument — can distinguish H_1 from H_2, because any such instrument would itself be a commitment event whose record structure is covered by $\mathscr{R}$. The proto-temporal ordering is not a limitation of measurement precision; it is outside the observable algebra by construction.

Different inertial observers will assign different coordinate orderings to A and B; each accesses the same equivalence class $[H] \in \mathscr{H}$_phys via a different coordinate chart. The structural analogy with gauge theory noted in Section 4.3 is here made concrete: the redundancy is the proto-temporal ordering; the invariant structure is the Lorentz-covariant committed record; the automorphism group of that structure is the Lorentz group.

---

5. Physical Consequences

5.1 Commitment Thresholds and Classical Emergence

Classical spacetime structure — with well-defined temporal ordering, stable causal relations, and Lorentz-covariant field dynamics — emerges only in regions where the commitment density $\lambda(s)$ exceeds a threshold λ_c sufficient for stable record formation. Below this threshold:

- Systems exhibit reversible or quasi-reversible dynamics
- No stable local time ordering exists
- The classical spacetime description breaks down

This does not imply that Lorentz invariance is violated below λ_c: the CRE theorem remains valid at all scales, since it concerns the structure of committed records, and below λ_c there are simply too few committed records to define a classical spacetime geometry. The breakdown is one of classical *description*, not of the underlying equivalence principle. This has a specific observational consequence: near the quantum-to-classical transition — realised in superconducting qubit systems at millikelvin temperatures — deviations from standard quantum dynamics should appear as the commitment density approaches λ_c from below. This is the basis for the quantitative prediction developed in Section 7.

5.2 Measurement as Commitment

In standard quantum mechanics, measurement is treated as a primitive that collapses the wavefunction, or as decoherence into a preferred basis, or (in relational QM) as the formation of a relational fact. In VERSF, measurement is precisely defined: it is the event at which λ transitions from zero to positive — a commitment event at proto-time s^* that generates a quantum of physical time at a specific location.

This is not merely a restatement of decoherence. Decoherence is a dynamical process by which quantum coherence is transferred to environmental degrees of freedom; it is reversible in principle and does not in itself produce an irreversible record. Commitment is the event at which a physical fact is irreversibly formed — it requires that Σ cross the threshold λ_c, stabilising the record against all future interactions. Decoherence transfers coherence to environmental degrees of freedom without guarantee that a permanent record is formed; it is a necessary condition for commitment but not a sufficient one. The distinction generates different predictions for the quantum-to-classical transition (Section 7).

5.3 Gravitational Time Dilation as Commitment Rate Variation

In VERSF, the commitment-capacity density κ(x) encodes the local rate at which irreversible facts can form. In the ambient field-free vacuum, κ takes the background value κ_0. Regions near massive objects have a reduced κ relative to κ_0 — the curvature of the fold-density field suppresses the rate of irreversible record formation — and physical time therefore accumulates more slowly in those regions relative to the global proto-time substrate. This is gravitational time dilation, derived from first principles: clocks near a massive object run slower because the local commitment rate is suppressed, not because of an externally imposed metric. The statement that high mass-energy regions "support more routes for commitment" applies to the internal microscopic structure of dense matter (many environmental degrees of freedom for record formation) rather than to the ambient gravitational field surrounding that matter, where the fold-density field is depressed rather than enhanced.

Terminology note. The "fold" in fold-density field refers to the void-energy fold structure of the VERSF ontology — the pre-geometric substratum from which commitment events and spacetime geometry emerge. κ(x) is the local density of fold-structure that supports irreversible fact formation; the fold-void interface is the boundary between regions of active fold-density and the void. These constructs are defined in detail in [3]; within the present paper they function as background motivation for the commitment-rate field. The formal results of this paper do not depend on their detailed structure.

In the companion paper [3], this is made precise: the metric $g_{\mu\nu}(x)$ is derived as the coarse-grained encoding of the local commitment-rate field $\kappa(x)/\kappa_0$, with the weak-field limit recovering the Schwarzschild solution and Newton's law of gravitation. General relativity is the effective field theory of the commitment-density field in the regime where commitment density varies smoothly above threshold.

---

6. Systematic Comparison with Competing Frameworks

We compare VERSF with loop quantum gravity (LQG), causal set theory (CST), and relational quantum mechanics (RQM) at three levels: ontological commitments, mechanism for temporal emergence, and experimental discriminability.

6.1 Loop Quantum Gravity

Ontological commitments. LQG quantises the gravitational field by replacing the smooth spacetime manifold with a discrete quantum geometry described by spin networks (states) and spin foams (transition amplitudes). The fundamental entities are quantum states of geometry; matter fields propagate on this quantum geometric background. Space is discrete at the Planck scale; time, in the covariant (spin foam) formulation, emerges from the transition amplitudes between spin network states.

Temporal emergence. LQG's approach to the problem of time is relational: physical time is defined by correlations between the gravitational degrees of freedom and matter degrees of freedom. In the spin foam formalism, no preferred time exists; the path integral sums over geometries without a background time. The recovery of classical time in the semiclassical limit is technically challenging and remains an open problem.

Where LQG excels. The mathematical framework of LQG is highly developed. The kinematic Hilbert space is rigorously defined, the spectrum of area and volume operators is derived, and Planck-scale discreteness is a theorem rather than an assumption. The connection to general relativity in the classical limit is well-understood at the kinematic level.

Where VERSF differs. LQG quantises geometry and treats time as secondary to geometric structure. VERSF inverts this: commitment is primary, and geometry — including the metric structure of spacetime — is emergent from the density field of commitment events. The two frameworks agree that time is not fundamental but disagree about what it is emergent from. More specifically, LQG has no mechanism for irreversible fact formation: in the standard time-symmetric formulation, spin network transitions are in principle reversible (the spin foam amplitude is symmetric under time reversal), and there is no account of why any particular outcome is the actual one rather than a superposition. KMS-condition extensions and thermodynamic formulations of spin foam models complicate this picture — in particular, the EPRL model's thermal partition function derivations partially address the arrow of time within the spin foam amplitude — but irreversibility in those treatments is typically imported via boundary conditions or KMS state choices rather than derived from the geometric structure itself. VERSF's commitment mechanism provides an explicit account: the transition from reversible quantum superposition to irreversible physical record is primitive, not derived.

Experimental discriminability. LQG predicts Planck-scale modifications to dispersion relations for high-energy photons, potentially observable in gamma-ray bursts. VERSF does not predict this modification in its standard form. Conversely, VERSF predicts the millikelvin decoherence anomaly (Section 7.1) and the cross-channel scale ratio (Section 7.2), which LQG does not. The two frameworks are discriminable in principle from both directions.

6.2 Causal Set Theory

Ontological commitments. CST takes the discrete partial order of spacetime events as primitive. The fundamental structure is a locally finite partial order (causet) on a set of elements, where each element represents an elementary spacetime event. Lorentz invariance is recovered statistically via the random sprinkling of causet elements into a Lorentzian manifold. The continuum spacetime of GR emerges in the appropriate approximation.

Temporal emergence. In CST, temporal order is not emergent — it is primitive. The causal structure (the partial order) is the fundamental object; time is simply the directed aspect of this order. There is no pre-temporal substrate from which causal order emerges. The arrow of time is addressed by Sorkin's sequential growth dynamics, in which the causet grows by accretion of new elements, providing a natural arrow from smaller to larger causets.

Where CST excels. CST has a technically clean ontology: a single primitive from which spacetime geometry emerges. The Lorentz invariance recovery via random sprinkling is elegant and mathematically rigorous. The entropy of causal horizons can be derived from the causet structure.

Where VERSF differs. The critical distinction is between causal order and irreversible commitment. In CST, every event in the causet is equally real — the partial order is a static structure in which past, present, and future all exist on equal footing. There is no intrinsic account of why facts are formed rather than merely ordered. VERSF insists that order alone is insufficient: a causal order can be defined on a space of reversible quantum possibilities without any fact having been established. Irreversibility — the production of a permanent record — is a further condition, and it is this condition that generates physical time in VERSF.

CST's sequential growth dynamics provides an arrow of time by stipulation (new elements are always added, never removed); VERSF derives the arrow from the thermodynamic irreversibility of commitment events ($\Sigma \geq 0$). Additionally, CST does not currently provide a quantum dynamics that recovers the Standard Model matter content or coupling constants. VERSF proposes derivations of the Standard Model gauge group, the fine-structure constant, and particle mass ratios within the broader programme — derivations that are detailed in companion papers and constitute a separate evaluable claim from the results of the present work.

Experimental discriminability. CST predicts a minimum spacetime interval near the Planck length with potentially observable effects in cosmic ray physics and quantum gravity phenomenology. VERSF shares the prediction of a minimum coherence scale $\xi \approx 88$ μm but at a vastly larger scale — a distinct and independently testable prediction. Both frameworks predict departures from classical spacetime at short distances, but at very different scales and through different mechanisms.

6.3 Relational Quantum Mechanics

Ontological commitments. RQM [18] holds that quantum states and physical properties are always relative to an observer system. There are no absolute facts about physical systems; every fact is a fact relative to some other physical system. The wavefunction encodes not absolute properties but the information available to a particular observer. Different observers may assign different states to the same system, and this is not a contradiction but a feature of the relational ontology.

Temporal emergence. In RQM, time is not fundamentally problematic because there is no absolute temporal ordering to recover. Facts are relational — they exist relative to an observer — and temporal ordering is similarly relational. The measurement problem is dissolved: collapse is observer-relative.

Where RQM excels. RQM has the most parsimonious ontology of the three competing frameworks: it requires no modifications to the Hilbert space formalism, no discrete fundamental structures, and no new mechanisms. It dissolves rather than solves several standard puzzles by reframing what is required of a physical theory.

Where VERSF differs. The core issue is global consistency. In RQM, mutually inconsistent facts among observers are treated as acceptable — the relational framework explicitly does not require different observers' accounts to be globally reconcilable. Frauchiger and Renner [15] formalised this tension, showing that under certain experimental conditions, different observers applying quantum theory consistently reach logically contradictory conclusions. Defenders of RQM — including Rovelli, Di Biagio, and Adlam — have argued in response that this does not constitute a genuine inconsistency within RQM but rather reflects that joint knowledge across observers is not required by the framework, since facts are relative and not globally aggregable.

In the Frauchiger-Renner scenario specifically, VERSF's response is structural rather than interpretive: the scenario involves an experimenter (Wigner) treating a measurement performed by a colleague (Friend) as a reversible quantum operation, which in VERSF is inadmissible. Once a commitment event has occurred — once Σ has crossed the threshold at the Friend's measurement device — the resulting fact is thermodynamically irreversible and cannot be coherently superposed or reversed by any subsequent operation. Wigner's unitary treatment of Friend's measurement presupposes that the measurement produced no irreversible record, which contradicts the VERSF definition of measurement. It should be noted that this response depends on the Friend's measurement device having crossed the Σ threshold — in a carefully isolated quantum system of the kind the FR setup requires, this is not automatically guaranteed. VERSF therefore makes a testable structural claim about the FR scenario: the scenario can only be realised in a system where the Friend's measurement device operates sub-threshold, i.e. where commitment has not occurred in the VERSF sense. A sophisticated objection is that this structure makes VERSF's response unfalsifiable in any particular FR experiment — VERSF could always classify a device as sub-threshold if Wigner's reversal succeeds, and above-threshold if it does not. VERSF pre-empts this by specifying an observable distinguisher: the success probability of Wigner's reversal should correlate systematically with independently measurable thermodynamic quantities of the Friend's device — specifically, its entropy production rate Σ, operating temperature, and environmental coupling strength. If Wigner's reversal succeeds while the Friend's device has $\Sigma > \Sigma\_threshold$ (determined independently), VERSF is falsified. A natural regress concern arises: measuring Σ of the Friend's device is itself a physical interaction that could constitute a commitment event, potentially collapsing the superposition. VERSF's response is that Σ can be characterised in a separate calibration run — measuring the entropy production of the same device under identical operating conditions but outside the FR setup — before the superposition is established. This calibration is not part of the FR superposition and does not collapse it. The prediction is then: devices with $\Sigma\_cal \geq \Sigma\_threshold$ should never permit successful FR reversal; devices with $\Sigma\_cal < \Sigma\_threshold$ should permit it. This is a falsifiable cross-experiment prediction. This response shares structural territory with spontaneous collapse models (GRW, CSL) and certain many-worlds variants, which also declare Wigner's reversal inadmissible. What is distinctive about the VERSF response is its grounding: the inadmissibility is not a new postulate or a stochastic collapse law but a consequence of the thermodynamic definition of commitment — the Σ threshold criterion that is independently motivated throughout this framework.

This response is coherent within RQM's own terms, but it comes at a cost: it forecloses the possibility of a global, observer-independent account of physical reality. VERSF's commitment mechanism is designed precisely to provide this: a fact becomes physically real not merely relative to an observer but through the irreversible production of a record that is thermodynamically stable against all future interactions. Commitment is not observer-relative — it is defined by Σ crossing a threshold, which holds regardless of who subsequently interacts with the system. Where RQM accepts observer-relativity as constitutive of physical reality, VERSF treats it as a deficiency to be resolved. Whether this is a virtue or a burden depends on one's prior commitments about what a physical theory must deliver; both positions are internally coherent.

The second distinction is explanatory scope. RQM has little to say about the values of physical constants, the particle content of the Standard Model, or the mass spectrum — these are accepted as empirical inputs. VERSF proposes derivations of these quantities within the broader programme. The two programmes address different questions, but the question of why physical constants have the values they do is not one RQM is structured to answer.

Experimental discriminability. On RQM's own terms, it makes no quantitative predictions that differ from standard quantum mechanics in any currently accessible regime. VERSF predicts the millikelvin decoherence anomaly (Section 7.1) and the cross-channel scale ratio (Section 7.2) — experiments that could in principle distinguish VERSF from both RQM and standard QM simultaneously.

6.4 Relation to the Connes-Rovelli Thermal Time Hypothesis

The commitment density $\lambda = \Sigma/(k_B \ln 2)$ is structurally reminiscent of the Connes-Rovelli thermal time hypothesis (Connes and Rovelli 1994), which also derives a notion of physical time from thermodynamic quantities — specifically, from the modular flow of a KMS state. Given this resemblance, a reader familiar with the thermal time programme will ask: how does VERSF's physical time differ from Connes-Rovelli thermal time?

The distinctions are real and important. In the Connes-Rovelli framework, thermal time is defined by the modular automorphism group of a KMS state — it is a global, state-dependent time that arises from equilibrium thermodynamics and is associated with the entire system. It does not require irreversible record formation, and it is defined even for time-reversal-invariant dynamics. In VERSF, physical time is generated locally and incrementally by commitment events: specific spacetime-localised transitions in which Σ crosses the threshold λ_c, producing an irreversible record. VERSF time is not defined by modular flow and does not require a global equilibrium state. It is generated wherever irreversible fact formation occurs, at a rate that varies from region to region.

Three structural differences follow from this. First, VERSF time has a definite arrow by construction: commitment events are thermodynamically irreversible ($\Sigma \geq 0$), so physical time is unidirectional. In the Connes-Rovelli framework, the modular flow can run in either direction (Δ^{it} vs Δ^{-it} as a mathematical choice), though in a given physical situation the KMS condition fixes the physical direction; this mathematical degree of freedom is absent in VERSF where the arrow is derived rather than input. Second, VERSF time is locally defined and can vary across spacetime, generating gravitational time dilation as a variation in the local commitment rate field $\kappa(x)$; Connes-Rovelli thermal time is a global parameter of the state. Third, VERSF's physical time is tied to the existence of irreversible records — it accumulates only when commitment events occur — whereas Connes-Rovelli thermal time flows continuously in an equilibrium state even without record formation. These are not merely differences in formalism; they generate different experimental predictions, different accounts of the arrow of time, and different treatments of the measurement problem.

6.5 Summary Comparison Table

The table below covers the three main competing quantum-spacetime frameworks — loop quantum gravity, causal set theory, and relational quantum mechanics. The Connes-Rovelli thermal time hypothesis, which is structurally distinct from these programmes (being a hypothesis about the origin of time rather than a full quantum-spacetime framework), is discussed separately in Section 6.4.

	VERSF	Loop QG	Causal Sets	Relational QM
Fundamental ontology	Irreversible commitment events	Quantum geometry (spin networks)	Discrete causal order	Observer-relative facts
Time primitive?	No — emergent from commitment	No — relational/emergent	Yes — partial order is primitive	No — observer-relative
Mechanism for time emergence	Local entropy production rate $\lambda = \Sigma/k_B \ln 2$	Correlations among geometric DOF	Sequential growth dynamics (stipulated)	No mechanism required
Lorentz invariance	Derived from CRE result (exact at record level)	Background independent; classical limit open	Statistical via random sprinkling	Standard QM formalism
Measurement mechanism	Commitment: irreversible record formation	Not specified	Not specified	Observer-relative fact formation
Global consistency of facts	Yes — commitment is thermodynamic, not relational	Not addressed	Yes — causet is globally defined	Not required; observer-relative
Coupling constants derived?	Yes — α, $\sin^2\theta_W$, m_p, m_e	No	No	No
Arrow of time	Derived: $\Sigma \geq 0$	Not explained	Stipulated by sequential growth	Observer-relative; not addressed globally
Quantitative novel prediction	Millikelvin decoherence anomaly; $\xi \approx 88$ μm scale; $\sigma_\tau/\sigma_{peak} = \sqrt{2\ln 2}$	Planck-scale dispersion	Planck-scale discreteness	None

---

7. Quantitative Discriminator Predictions

7.1 Decoherence Rate Anomaly at Millikelvin Temperatures

In the VERSF framework, the physical time rate $\lambda(s)$ is proportional to the local entropy production rate Σ. In a quantum system at temperature T in contact with an environment, Σ depends on the rate of irreversible energy exchange. At millikelvin temperatures — accessible with current dilution refrigerators and superconducting qubit technology — the thermal contribution to Σ becomes comparable to the zero-point contribution, and the commitment density λ deviates from its high-temperature value.

The standard energy-relaxation rate for a superconducting qubit (the T_1 channel) is:

$$\Gamma\_std = \gamma_0 \coth(\hbar\omega / 2k\_BT)$$

where γ_0 is the zero-temperature limit and ω is the qubit transition frequency. In VERSF, this rate acquires a correction from the finite commitment density. The following formula is presented as schematic — it captures the commitment suppression in the T_1 channel, while the full treatment in [8] separates the T_1 and T_2 (pure dephasing) contributions, which the TLS mechanism of §7.1 primarily affects:

$$\Gamma\_VERSF = \gamma_0 \coth(\hbar\omega / 2k\_BT) \cdot (1 - e^{-\lambda(T)/\lambda\_c})$$

The displayed formula applies cleanly to the T_1 channel. The T-linear thermal component $\lambda\_th \propto T$ derives from TLS-induced pure dephasing (see below), which contributes to T_2 rather than T_1 directly — the VERSF suppression in the T_2 channel takes the same functional form with the same threshold but with a different underlying rate formula. In both channels, the commitment-suppression factor $(1 - e^{-\lambda(T)/\lambda\_c})$ applies, and the direction of the prediction (enhanced coherence times T_1 and T_2 relative to Lindblad extrapolation) holds in both. The T_2 enhancement is expected to be more pronounced in the sub-10 mK regime because it is directly driven by the TLS mechanism; the T_1 enhancement depends on whether $\lambda\_ZP < \lambda\_c$. This channel separation is a further discriminating signature of the prediction.

Here λ_c is the commitment threshold for classical record formation. The commitment rate $\lambda(T)$ has two contributions: a zero-point floor λ_ZP arising from vacuum fluctuations at $T = 0$, and a thermal component $\lambda\_th(T) \propto T$ at low temperatures, so that $\lambda(T) = \lambda\_ZP + \lambda\_th(T)$. This decomposition is physically necessary: zero-point electromagnetic vacuum fluctuations are a well-established source of decoherence in superconducting qubits even at absolute zero — they are responsible for the finite T_1 floor γ_0 that appears in the standard formula. In VERSF, vacuum fluctuations couple to the qubit and produce entropy, so they must be assessed against the commitment threshold. Two regimes result. If $\lambda\_ZP \geq \lambda\_c$ — vacuum fluctuations provide sufficient entropy production to cross the commitment threshold — then $\Gamma\_VERSF \to \gamma_0$ as $T \to 0$, recovering the standard zero-temperature floor. If $\lambda\_ZP < \lambda\_c$ — vacuum fluctuations do not cross the commitment threshold and therefore do not produce irreversible records in the VERSF sense — then $\Gamma\_VERSF(T{=}0) = \gamma_0(1 - e^{-\lambda\_ZP/\lambda\_c}) < \gamma_0$, predicting a further suppression of the zero-temperature decoherence floor below the standard value. This is an additional VERSF prediction: the T_1 floor at the lowest accessible temperatures should be higher than the standard Lindblad formula extrapolates, by a factor that depends on the ratio λ_ZP/λ_c. Distinguishing these two regimes experimentally provides a route to measuring λ_ZP independently of $\lambda\_th(T)$. Crucially, this measurement is self-contained: it does not require prior knowledge of which regime holds. The T_1 floor at the lowest accessible temperatures is directly observable, and its deviation from the standard Lindblad extrapolation (or lack thereof) immediately reveals whether $\lambda\_ZP \geq \lambda\_c$ or $\lambda\_ZP < \lambda\_c$, without needing to know the value of λ_ZP in advance. The derivation of the thermal component $\lambda\_th(T) \propto T$ from the low-temperature phonon density of states is given in [8]; the quantitative derivation of λ_ZP from vacuum fluctuation entropy requires further development within the fold-dynamics framework.

It should be noted that λ_c is a framework parameter — a threshold whose value is not derived from first principles within the present paper but must be constrained empirically. Operationally,

λ_c is defined as the minimum entropy production rate required for records to remain stable under all subsequent environmental interactions. An empirical upper bound follows directly from existing transmon qubit data: current T_1 and T_2 measurements at 10–20 mK are consistent with standard Lindblad dynamics, which requires that $\lambda_c \gtrsim \lambda(20\ \text{mK})$ — the commitment threshold lies above the thermal commitment rate at 20 mK. This bound places the VERSF anomaly in the sub-10 mK regime (Section 7.1 below). The decoherence prediction is therefore parameter-free in its *direction* (suppression of Γ relative to Lindblad) but not its *scale* (the temperature of onset and the magnitude of the zero-temperature floor suppression both depend on λ_c and λ_{ZP}). In this sense, λ_c plays a role analogous to a renormalisation scale, marking the transition between reversible and irreversible regimes of the effective description.

Derivation of the suppression form. The Boltzmann activation argument motivates the exponential suppression factor as follows. The commitment rate $\lambda(T) = \lambda_{ZP} + \lambda_{th}(T)$ has two components: a zero-point floor λ_{ZP} from vacuum fluctuations and a thermal component $\lambda_{th}(T) \propto T$. The transition from quantum superposition to irreversible record requires Σ to cross the commitment threshold λ_c. The fraction of environmental interactions providing sufficient entropy production takes a Boltzmann activation form: the activation probability scales as $e^{-\lambda_c/\lambda(T)}$, giving a suppression factor $(1 - e^{-\lambda(T)/\lambda_c})$ representing the fraction of attempted decoherence events that succeed in forming irreversible records.

It should be stated clearly that the exponential functional form is a hypothesis consistent with the commitment architecture rather than a uniquely derived prediction. A proper activation-rate derivation would specify a barrier height (in units of entropy production), an attempt frequency, and a temperature — and the current framework does not yet provide all three. The three constraints the exponential satisfies — vanishing at $\lambda \to 0$, recovering Lindblad at $\lambda \gg \lambda_c$, and activation-type threshold behaviour — are satisfied by a universality class of smooth monotone functions, of which the exponential is the simplest representative. The exponential is therefore a natural and minimal hypothesis, not a derivation. The prediction that is robustly made — independent of the specific functional form — is the existence and direction of the suppression: any activation-type function satisfying these constraints yields a decoherence rate below the Lindblad value in the sub-threshold regime.

On $\lambda_{th}(T) \propto T$. The thermal component of the commitment rate at low temperatures is claimed to satisfy $\lambda_{th}(T) \propto T$, with the full derivation deferred to [8]. Since [8] is a companion paper whose content is not reproduced here, a sketch of the physical argument is provided for referee verification. The correct mechanism at sub-10 mK must be identified carefully. At these temperatures ($T \sim 10$ mK, with qubit frequencies $\omega_0/2\pi \sim 5$–7 GHz giving $\hbar\omega_0/k_B \sim 0.24$–$0.34$ K so that $k_BT \ll \hbar\omega_0$), phonon-phonon Umklapp processes are entirely frozen — Umklapp rates are exponentially suppressed as $e^{-\Theta_D/bT}$ with $\Theta_D \sim$ hundreds of Kelvin — and the thermal bath contribution to qubit relaxation via the microwave environment scales as $e^{-\hbar\omega_0/k_BT}$, not as T. The relevant mechanism for a T-linear thermal commitment rate is instead coupling to two-level system (TLS) defects in the substrate and oxide layers of superconducting qubits. TLS ensembles in the low-temperature limit produce broadband noise with power spectral density $S(\omega) \propto T/\omega$ (the experimentally observed 1/f-like noise in superconducting circuits), arising from TLS relaxation and resonant absorption processes. Integrating this over the relevant bandwidth gives a thermal dephasing contribution $\Gamma_\varphi \propto T$, hence $\lambda_{th} \propto T$ at leading order. This mechanism operates via pure dephasing (T_2 contribution) rather than energy relaxation (T_1), so the predicted suppression enhancement is more pronounced in T_2 than T_1 in this regime — a distinguishing signature relative to non-commitment-based explanations. The full derivation, including the crossover between TLS-dominated and vacuum-fluctuation-dominated regimes, is given in [8].

Division of labour with [8]. Reference [8] ("Two Kinds of Time: Proto-Time and Physical Time in the VERSF Framework") is a companion paper to the present work covering overlapping conceptual territory. The division is as follows. [8] develops the two-layer temporal ontology as a standalone framework, including the derivation of λ(T) from phonon physics, the WdW Hamiltonian constraint analysis, and the connection to thermodynamic time. The present paper takes the two-layer ontology as given and addresses the specific question of Lorentz covariance — showing that proto-time ordering is outside the observable algebra and that Lorentz invariance emerges at the record level — and provides the systematic comparison with competing frameworks and quantitative experimental predictions. The two papers are designed to be read independently, with the present paper providing the covariance and comparison content and [8] providing the deeper derivation of the temporal ontology itself.

Relation to existing data. Transmon qubits operating at 10–20 mK with transition frequencies ω/2π ~ 5–7,GHz have been characterised extensively. Current T_1 and T_2 measurements at these temperatures are consistent with standard Lindblad dynamics within experimental precision. This constrains the commitment threshold to λ_c ≳ λ(20,mK), placing the VERSF anomaly at temperatures below roughly 10,mK for typical qubit parameters — a regime accessible with state-of-the-art dilution refrigerators but not yet systematically explored for the specific purpose of testing commitment-threshold physics. The prediction is therefore not in conflict with existing data but directs experimental attention to the sub-10 mK regime, where enhanced T_1 and T_2 values relative to Lindblad extrapolation constitute the signature.

Experimental signature.

Prediction summary. VERSF makes two distinct predictions in the decoherence channel. First, a suppression of the decoherence rate — an enhancement of T_1 and T_2 — relative to the standard Lindblad prediction in the sub-10 mK regime, where λ_th(T) < λ_c; the existence and direction of this anomaly are parameter-free. Second, a further suppression of the zero-temperature T_1 floor if the zero-point commitment rate λ_ZP < λ_c — meaning vacuum fluctuations do not cross the commitment threshold — yielding Γ_VERSF(T=0) < γ_0. The first prediction constrains λ_c from the thermal onset temperature; the second constrains λ_ZP from the T=0 floor independently. Neither LQG, CST, nor RQM predicts any deviation from standard decoherence in either regime.

The exponential form (1 - e^-λ/λ_c) is the minimal form consistent with activation-threshold reasoning; other smooth functions satisfying the same boundary conditions are asymptotically equivalent near threshold, and a quantitative fit to data would constrain or falsify the functional form independently of the existence prediction.

7.2 Cross-Channel Scale Ratio for Localised Commitment-Density Anomalies

The VERSF framework predicts a specific, parameter-free ratio between the spatial scales of observable signatures of a localised commitment-capacity depression. A region of suppressed commitment-capacity density $\kappa(x) < \kappa_0$ produces simultaneous signatures in three channels:

- **Temporal channel:** physical processes within the region slow, with half-depth radius σ_τ
- **Optical channel:** light paths are deflected (enhanced effective refractive index), with characteristic radius σ_{peak}
- **Electromagnetic channel:** a corona of enhanced commitment activity at the field boundary, with peak radius σ_{em}

From the VERSF optical propagation model — in which the effective refractive index $n_{eff} = n_0(\kappa_0/\kappa)^{1/2}$ is fixed by metric-conformal consistency and null-cone compatibility — the relationship between the temporal and optical scales is:

$$\sigma_\tau / \sigma_{peak} = \sqrt{2 \ln 2} \approx 1.18$$

Profile class and the role of the Gaussian. The spatial profile $\kappa(x) = \kappa_0(1 - A \cdot e^{-r^2/2\sigma^2})$ is used here not as a uniquely privileged physical form, but as the minimal smooth, isotropic, rapidly decaying profile representing a localised commitment-capacity depression. Any sufficiently regular, radially symmetric profile with a single characteristic scale σ admits a local expansion of the form $\kappa(r) = \kappa_0(1 - A\, f(r/\sigma))$, where f is smooth, $f(0) = 1$, and $f''(0) < 0$. The Gaussian corresponds to the case in which f has minimal higher-order structure beyond the quadratic term. The ratio derived below is therefore expected to hold at leading order for the entire class of such profiles, with corrections controlled by deviations from Gaussianity. Here A satisfies $0 < A < 1$, which is required for $\kappa(x) > 0$ throughout — the commitment-capacity density must remain positive everywhere for the optical propagation model to be well-defined.

Physical origin of the single scale. The emergence of a single characteristic scale σ reflects the locality of the commitment-capacity depression: $\kappa(x)$ is determined by the spatial distribution of fold-density suppression, which is governed by localised energy or entropy constraints. In the absence of additional structure — multiple sources, anisotropy, or sharp boundaries — this produces a single-scale, isotropic profile. The ratio $\sigma_\tau/\sigma_{peak}$ is therefore not a property of the specific functional form chosen, but of the underlying single-scale geometry of the commitment-density field.

The exponent 1/2 in $n_{eff} \propto (\kappa_0/\kappa)^{1/2}$ is not a free parameter: it is required by null-cone compatibility. For the optical path length to transform covariantly under the same conformal factor as the spacetime metric — which in VERSF is set by $\kappa(x)/\kappa_0$ — the refractive index must scale as the square root of the metric conformal factor. A different exponent would break null-cone compatibility, causing photon propagation to violate the causal structure defined by the commitment-density field. This derivation is detailed in [4].

Derivation of the ratio. For a general single-scale profile $f(r/\sigma)$, the ratio takes the form:

$$\sigma_\tau / \sigma_{peak} = r_{half} / r_{peak}$$

where r_half solves $f(r\_half/\sigma) = 1/2$ and r_peak solves $d/dr\,[r\, f'(r/\sigma)] = 0$. For the Gaussian, $f(u) = e^{-u^2/2}$, these yield $r\_half = \sigma\sqrt{2 \ln 2}$ and $r\_peak = \sigma$, giving the ratio $\sqrt{2 \ln 2}$. For general f, the ratio depends only on these two dimensionless solutions, which are controlled by the quadratic curvature of f near the origin: the dominant terms are determined by $f''(0)$, and corrections enter at order $(f''''(0)/f''(0))\sigma^2$. The ratio is therefore fixed to leading order by the local curvature structure of the profile — not by the specific functional form.

For the Gaussian profile specifically: σ_τ is the half-depth radius — the radial distance where $\kappa(\sigma\_\tau) = \kappa_0(1 - A/2)$, giving $\sigma\_\tau = \sigma\sqrt{2 \ln 2}$. σ_peak is the peak-deflection impact parameter — the value of b maximising the deflection angle $\alpha(b) \propto b \cdot e^{-b^2/2\sigma^2}$, which peaks at $b = \sigma$. Therefore $\sigma\_peak = \sigma$, and:

$$\sigma\_\tau / \sigma\_peak = \sigma\sqrt{2 \ln 2} / \sigma = \sqrt{2 \ln 2} \approx 1.18$$

The electromagnetic channel scale σ_em equals σ because the gradient $\partial\kappa/\partial r$ is maximal at $r = \sigma$ for the Gaussian profile — the EM corona arises from enhanced commitment activity at the boundary of the depression, placing it at $\sigma\_em = \sigma$. The full derivation is in [4]. The three-channel consistency relation is:

$$\sigma\_\tau / \sqrt{2 \ln 2} = \sigma\_peak = \sigma\_em \equiv \sigma$$

where σ is the characteristic scale of the underlying commitment-density profile.

Robustness. The ratio $\sqrt{2 \ln 2}$ is exact for the Gaussian and is the leading-order prediction for any smooth, isotropic, single-scale commitment-density profile. Deviations arise only from higher-order shape corrections — specifically from $f^{(4)}(0)$ and beyond — and the leading correction to the ratio is of order $(f^{(4)}(0)/f''(0))\sigma^2$, which is dimensionless and small whenever the profile has no sharp features on scales below σ. The prediction is therefore robust to moderate deviations from Gaussianity, with the correction scale explicitly controlled by the ratio of fourth to second profile derivatives evaluated at the origin. A systematic expansion in profile moments, treating deviations from the Gaussian as perturbations, provides a well-defined route to quantifying these corrections and constitutes an open but tractable theoretical extension.

Any experimental realisation of a localised commitment anomaly should exhibit this cross-channel consistency. Deviation of the measured ratio from $\sqrt{2 \ln 2}$ constitutes an immediate falsification of the VERSF optical model for single-scale profiles, or evidence of multi-scale or anisotropic structure in the commitment-density field.

---

8. Discussion

8.0 Relation to the VERSF Programme

The present work forms part of a broader research programme within the VERSF framework, in which the emergence of physical time, irreversible record formation, and quantum dynamics are treated as aspects of a unified underlying structure. Three components of that programme are particularly relevant to the results established here.

First, the commitment ontology — the identification of physical facts with irreversible record formation in finite-capacity systems — is developed in *Two Descriptions of Reality: The Coherence Scale as a Commitment Threshold* [23]. That work establishes, under minimal operational constraints, that any universe supporting both reversible dynamics and stable records must contain a threshold separating proto-factual evolution from irreversible fact formation. The present paper takes this structure as given and explores its consequences for observable symmetry and spacetime structure.

Second, the two-layer temporal framework — distinguishing proto-time (the ordering parameter of reversible quantum evolution) from physical time (the accumulation of irreversible commitment events) — is developed in *Two Kinds of Time: Proto-Time and Physical Time in the VERSF Framework* [8]. The proto-temporal ordering employed in Sections 2–4 is directly grounded in that construction. The CRE result of the present paper can be understood as the statement that proto-time ordering is eliminated from the observable sector by the same mechanism that generates physical time.

Third, the distinction between coherence and irreversible commitment is examined in *The Quantum Zeno Effect as Suppression of Irreversible Commitment* [24]. That work shows that standard quantum dynamics are recovered at leading order, while higher-order deviations arise when measurement suppresses commitment rather than coherence. The decoherence prediction in Section 7.1 and the threshold structure assumed throughout the present paper are consistent with, and partially motivated by, those results.

The role of the present paper within this programme is therefore specific: it establishes that, given the commitment ontology and the two-layer temporal structure, relativistic covariance is preserved at the level of observables through the Commitment Reordering Equivalence construction. In this sense, it provides the spacetime-consistency component of a broader framework in which quantum dynamics, thermodynamics, and time are derived from a common structural foundation.

8.1 VERSF's Structural Position

The comparison in Section 6 reveals a consistent pattern. VERSF shares with LQG the view that spacetime geometry is not fundamental. It shares with CST the view that discrete events, not a continuous manifold, are the primitive objects. It shares with RQM the view that facts are not absolute properties of systems in isolation. But it differs from all three in a specific and consequential way: it identifies **irreversibility** as the further condition needed to transform a causal order, or a geometric state, or a relational fact, into a genuine physical reality.

This is not a minor amendment. It changes what the theory has to explain and what it predicts. A theory in which all events in a partial order are equally real — past, present, and future — has no mechanism for why the present moment is experienced as special, why memory is of the past rather than the future, or why measurement outcomes are permanent. VERSF addresses all three:

the present is the frontier of commitment accumulation; memory is the accumulated record of past commitment events; measurement outcomes are permanent because commitment is thermodynamically irreversible.

It is not claimed that the present framework constitutes a complete derivation of spacetime structure from first principles. The result established here is conditional: given a commitment-based ontology satisfying the PAR condition, Lorentz invariance emerges at the level of observable records through CRE, and two quantitative predictions follow that distinguish the framework from all current competitors. The purpose of this paper is to establish the internal consistency and empirical testability of this structure — not to claim closure of all foundational questions. Several significant open problems are acknowledged in Section 8.3, and their presence does not undermine the conditional result; it defines the research programme's forward direction. The present work isolates a single structural question — the compatibility of proto-temporal ordering with relativistic covariance — and resolves it within the VERSF framework.

8.2 The Identification Question

Scope of the present result. The present work establishes a conditional result: given a commitment-based ontology satisfying PAR and compositional completeness, and given the Minkowskian causal structure of committed records, CRE is the observable-sector expression of PAR's symmetry requirement and yields Lorentz invariance at the level of committed records. It does not claim a full derivation of spacetime structure from first principles, nor a complete dynamical theory of commitment formation. Readers wishing to verify the formal underpinnings of $\mathscr{H}$_proto, $\mathscr{R}$, and $\mathcal{O}$ are directed to Appendix A, which provides the precise definitions as a supplement to the body treatment in Section 4.1. The formal proof that CRE is PAR's observable-sector expression is established in the companion programme [5].

The most important open question for VERSF — as for any theory that derives standard physics from a novel substrate — is whether the identification steps are fully earned. The framework derives mathematical structures that match measured physical constants; whether those structures are the physical objects they are identified with requires either a more complete derivation chain or experimental confirmation.

The two predictions in Section 7 are the programme's primary near-term answer to this question. They are novel, quantitative, and near-term accessible. If the decoherence rate anomaly is observed at millikelvin temperatures, or the cross-channel scale ratio is confirmed in a controlled commitment anomaly experiment, the identification steps are supported by evidence that no other framework predicts. If either prediction is falsified, specific sectors of the VERSF framework fail in identifiable ways, and the programme can be updated accordingly.

8.3 Open Problems

Electroweak symmetry breaking. The VERSF programme derives the Standard Model gauge group and constrains the Higgs mass but does not yet derive the Higgs mechanism — the dynamical breaking of $SU(2) \times U(1)_Y \rightarrow U(1)_{EM}$ — as a consequence of the commitment architecture. This is a significant open problem.

The signal-propagation speed identification. The Lorentz emergence argument of Section 4.2 depends on a load-bearing assumption: that the c bounding commitment-event signal propagation and the c appearing in the Minkowski metric are the same constant. This identification is not derived within the present paper; it is flagged in Section 4.2 as a gap whose resolution is the subject of ongoing work in the fold-dynamics programme [3]. The importance of this open problem deserves emphasis: if the two speeds differ by any amount, the causal structure of committed records is not exactly Minkowskian, and the Zeeman–Alexandrov automorphism argument yields a different symmetry group. The conditional result of this paper — Lorentz invariance at the level of committed records — is therefore conditional not only on PAR and CC but also on this identification. It is treated as a structural assumption of the commitment interaction model, justified by the expectation that VERSF fold-dynamics will derive it, but the derivation has not yet been achieved.

Strong commitment-field regime. The gravity derivation in [3] works in the weak-field and semiclassical regimes, where the commitment-capacity density κ(x) varies smoothly and remains well above threshold. What VERSF predicts when κ(x) varies strongly — in the vicinity of extremely dense matter, collapsed objects, or regions where fold-density gradients become large — is not yet fully worked out. VERSF does not quantise gravity and does not invoke a Planck scale as a fundamental input; the natural question in VERSF's own terms is what the commitment architecture implies when the fold-density field enters the strong-variation regime. One natural expectation is that λ_c behaviour changes qualitatively in this regime, but this is a conjecture rather than a derived result; the precise strong-field commitment structure requires further development.

Fold-level covariance. The CRE result establishes exact Lorentz invariance at the level of committed records (Section 4.2). A separate open question is whether the fold-void interface dynamics that generates commitment events is itself exactly Lorentz covariant, or whether it introduces corrections at the commitment scale — invisible at the level of committed records but potentially observable at extreme energies — that depart from exact covariance. As noted in Section 4.2, these two questions are logically independent: the result does not require fold-level covariance, and resolution of this open problem could generate additional experimental signatures in VERSF's own framework rather than threatening the result's conclusions. Framing this in terms of a "Planck scale" would import assumptions from quantum gravity that VERSF does not make; the relevant scale is instead the characteristic energy at which fold-density fluctuations become non-negligible, which is a derived quantity of the VERSF fold dynamics.

8.4 Anticipated Objections and Scope Clarifications

The VERSF framework introduces structural elements — in particular the commitment ontology and the PAR → CRE correspondence — that differ from standard approaches. Several natural objections arise; we address them here to clarify the scope and status of the present results.

(1) Is PAR an ad hoc assumption? A primary concern is whether the Pre-Factual Algebraic Reversibility condition is introduced to enforce the desired outcome — eliminating proto-temporal ordering from the observable sector — rather than being independently motivated. Within the VERSF programme, PAR is not postulated arbitrarily but arises as the common consequence of three independent constraints: admissibility (A2), observational faithfulness, and thermodynamic no-erasure conditions (Section 4.4.1 and [5]). Each of these excludes irreversible transitions in the pre-factual domain on different grounds. The convergence of these arguments provides the structural motivation for PAR. The present paper uses PAR as an input condition whose consequences for observable structure are explored. The claim is therefore conditional: given PAR and compositional completeness, CRE follows and yields Lorentz invariance at the level of committed records.

(2) Does the CRE construction assume Lorentz invariance? A second concern is that the use of invariant intervals and causal structure presupposes the Lorentz symmetry that the construction claims to recover. This concern is addressed by distinguishing between the operational structure — defined by signal propagation and interval measurement procedures — and the symmetry group, identified as the automorphism group of that structure. The construction assumes only that commitment events admit a finite signal-propagation bound and that invariant intervals are operationally definable. Given these invariants, the symmetry group is identified via the Zeeman–Alexandrov results. Lorentz invariance is therefore derived as the symmetry of the invariant structure, not imposed as an initial axiom. However, the identification of the signal-propagation bound and the metric speed of light as the same constant c remains an explicit assumption at the present stage, and the Lorentz result is correctly interpreted as conditional on this identification (Sections 4.2 and 8.3).

(3) Is λ_c arbitrary? The appearance of a threshold parameter λ_c raises the concern that the framework introduces an unconstrained phenomenological constant. This is not the intended role of λ_c. Operationally, λ_c is defined as the minimum entropy production rate required for records to remain stable under all subsequent environmental interactions — it is an emergent stability threshold determined by the interaction structure of the system, not a fundamental constant. The framework makes predictions that are parameter-free in direction (suppression of decoherence below the Lindblad rate) and parameter-dependent in scale (the temperature of onset depends on λ_c). This is analogous to the role of a renormalisation scale in effective field theory: λ_c marks a transition between regimes of validity, and its value is empirically constrained and in principle measurable from the existing transmon data bound $\lambda_c \gtrsim \lambda(20\ \text{mK})$.

(4) Does the framework provide a complete dynamical theory? A further objection is that VERSF specifies structural constraints — PAR, CRE, the commitment ontology — but does not yet provide a complete dynamical theory of commitment formation or fold dynamics. This is correct and is explicitly acknowledged. The present paper isolates a specific structural question — the compatibility of proto-temporal ordering with relativistic covariance — and resolves it within the VERSF framework. The relationship to the broader programme is: this paper addresses structural consistency and observable consequences; companion work addresses the derivation of PAR, fold dynamics, and the gravitational sector. The framework is a research programme in development rather than a completed theory.

(5) Are the experimental predictions robust? The decoherence prediction involves a functional form $(1 - e^{-\lambda/\lambda_c})$ that is not uniquely derived. The framework distinguishes between robust predictions — the existence and direction of the anomaly (suppression of decoherence in the sub-threshold regime) — and model-dependent details (the precise functional form of the transition). The existence of a systematic deviation from Lindblad behaviour in the sub-10 mK regime is the decisive test, and any activation-type suppression mechanism consistent with the commitment architecture yields this qualitative prediction independently of the specific functional form.

(6) Is VERSF simply a reformulation of existing approaches? One may ask whether VERSF reduces to thermal time, collapse models, or relational interpretations. The distinctions outlined in Section 6 are structural: commitment vs geometry (LQG), commitment vs causal order (CST), commitment vs observer-relativity (RQM), commitment vs equilibrium modular flow (Connes-Rovelli thermal time). While partial overlaps exist, the defining feature of VERSF — irreversible record formation as the generator of physical time — is not present as a primitive in these frameworks. Spontaneous collapse models (GRW, CSL) are a partial exception: they do treat irreversible record formation as primitive, and the paper acknowledges structural overlap with them in §6.3. The distinction from collapse models is not ontological but derivational and structural: in GRW and CSL, collapse is introduced as a new stochastic postulate modifying the Schrödinger equation; in VERSF, irreversible commitment is derived from the thermodynamic Σ threshold criterion, which is independently motivated and does not require modifying quantum dynamics. More importantly, collapse models make no Lorentz emergence argument of the kind developed here — the connection between irreversibility and the CRE quotient structure is specific to VERSF. The differences from all frameworks addressed in Section 6 are therefore not merely interpretational but structural and, via the predictions of Section 7, testable.

Summary of scope. The results of this paper are best understood as a conditional structural statement: given a commitment-based ontology satisfying PAR and compositional completeness, and given the Minkowskian causal structure of committed records, Lorentz invariance emerges at the level of observable physics through the CRE quotient. The purpose of the present work is to establish the internal consistency and empirical accessibility of this structure. Questions of deeper derivation and dynamical completion are explicitly identified as open problems within the programme.

---

9. Conclusion

We have established three results.

First, the apparent conflict between proto-time ordering and relativistic covariance is addressed by the Commitment Reordering Equivalence principle — shown to be the observable-sector expression of Pre-Factual Algebraic Reversibility, not an independent postulate. Physical history is the quotient of proto-temporal histories under the equivalence of identical committed record structures, and all observables on this quotient are Lorentz invariant, given the Minkowskian causal structure of committed records. This result is conditional on that identification — treated as a structural assumption of the commitment interaction model and flagged as the primary open problem in Section 8.3. Proto-time ordering is an unobservable degree of freedom, structurally analogous to gauge redundancy, and Lorentz invariance is its geometrical corollary rather than an independent axiom.

Second, VERSF is structurally distinct from loop quantum gravity, causal set theory, and relational quantum mechanics in a specific and consequential way: it identifies irreversible commitment — the thermodynamic event of permanent record formation — as the primitive bridge from quantum possibility to physical reality. This provides a global consistency of facts (lacking in RQM on its standard reading), a derived arrow of time (lacking in LQG and stipulated rather than derived in CST), and an explicit measurement mechanism (absent in all three).

Third, VERSF makes quantitative predictions that distinguish it from all competing frameworks. In the decoherence channel: a suppression of decoherence rates in superconducting qubits in the sub-10 mK regime (existence and direction are parameter-free; scale depends on λ_c), and a further suppression of the zero-temperature T_1 floor if vacuum fluctuations do not cross the commitment threshold (constraining λ_{ZP} independently). In the optical channel: a parameter-free cross-channel scale ratio $\sigma_\tau/\sigma_{peak} = \sqrt{2 \ln 2}$ for localised commitment-density anomalies, exact at leading order for Gaussian profiles. All predictions are near-term accessible with current experimental technology. The result does not claim to complete the derivation of spacetime structure from irreversible commitment, but establishes a consistent and testable structural bridge between reversible dynamics and relativistic observables.

Time is not a background. It is not a dimension. It is not a relational fact. It is a cumulative product — built, moment by moment, one irreversible commitment at a time. The framework therefore identifies irreversibility — not geometry or causal order — as the structural origin of physical time.

References

[1] Taylor, K. (2025). Introducing the Ticks-Per-Bit Framework. VERSF Theoretical Physics Programme.

[2] Taylor, K. (2025). Why the Universe Chose SU(3) × SU(2) × U(1). VERSF Theoretical Physics Programme.

[3] Taylor, K. (2026). Gravity from Fold-Density Gradients: Deriving Newton's Constant in the VERSF Framework. VERSF Theoretical Physics Programme.

[4] Taylor, K. (2026). Coupled Temporal, Optical, and Electromagnetic Signatures of a Localised Fold-Density Anomaly. VERSF Theoretical Physics Programme.

[5] Taylor, K. (2026). On the Structural Status of Algebraic Reversibility and Compositional Completeness in Fact-Producing Universes. VERSF Theoretical Physics Programme. [Primary formal treatment: establishes PAR + CC ⟹ IAC; proves A2 ⊬ PAR; develops three convergent routes to PAR; demonstrates PAR and CC formalise conditions implicit in standard quantum theory and relativity.]

[6] Taylor, K. (2026). Internal Admissible Closure from Pre-Factual Reversibility in Fact-Producing Universes. VERSF Theoretical Physics Programme. [Weight-level framework derivation: introduces the Separation of Levels principle; proves IAC via irreducible-residual argument; establishes that IAC implies condition (S2) closing the gap to complex Hilbert space.]

[7] Taylor, K. (2026). K = 7: The Constraint Dimensionality of Stable Physical Reality. VERSF Theoretical Physics Programme.

[8] Taylor, K. (2026). Two Kinds of Time: Proto-Time and Physical Time in the VERSF Framework. VERSF Theoretical Physics Programme.

[9] Alexandrov, A. D. (1953). On Lorentz transformations. *Uspekhi Matematicheskikh Nauk*, 8(1), 129–130.

[10] Barbour, J. (1999). *The End of Time*. Oxford University Press.

[11] Bombelli, L., Lee, J., Meyer, D., & Sorkin, R. (1987). Space-time as a causal set. *Physical Review Letters*, 59(5), 521.

[12] Brukner, Č. (2018). A no-go theorem for observer-independent facts. *Entropy*, 20(5), 350.

[13] Connes, A., & Rovelli, C. (1994). Von Neumann algebra automorphisms and time-thermodynamics relation in generally covariant quantum theories. *Classical and Quantum Gravity*, 11(12), 2899.

[14] Di Biagio, A., & Rovelli, C. (2021). Stable facts, relative facts. *Foundations of Physics*, 51, 30.

[15] Frauchiger, D., & Renner, R. (2018). Quantum theory cannot consistently describe the use of itself. *Nature Communications*, 9, 3711.

[16] Page, D. N., & Wootters, W. K. (1983). Evolution without evolution: Dynamics described by stationary observables. *Physical Review D*, 27(12), 2885.

[17] Penrose, R. (1989). *The Emperor's New Mind*. Oxford University Press.

[18] Rovelli, C. (1996). Relational quantum mechanics. *International Journal of Theoretical Physics*, 35(8), 1637.

[19] Sorkin, R. D. (1997). Forks in the road, on the way to quantum gravity. *International Journal of Theoretical Physics*, 36(12), 2759.

[20] Thiemann, T. (2007). *Modern Canonical Quantum General Relativity*. Cambridge University Press.

[21] Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In *Complexity, Entropy, and the Physics of Information*, ed. W. Zurek. Addison-Wesley.

[22] Zeeman, E. C. (1964). Causality implies the Lorentz group. *Journal of Mathematical Physics*, 5(4), 490–493.

[23] Taylor, K. (2026). Two Descriptions of Reality: The Coherence Scale as a Commitment Threshold. VERSF Theoretical Physics Programme.

[24] Taylor, K. (2026). The Quantum Zeno Effect as Suppression of Irreversible Commitment. VERSF Theoretical Physics Programme.

Appendix A: Formal Structure of Histories, Records, and Observables

This appendix provides the minimal formal definitions underlying the Commitment Reordering Equivalence construction of Section 4. The purpose is not to provide a full axiomatic reconstruction, but to demonstrate that the CRE construction rests on a well-defined quotient structure and observable algebra — establishing that the main results are grounded in precise mathematical objects rather than conceptual language alone.

A.1 Proto-History Space

Definition A.1 (Commitment event). A commitment event C is a triple $C = (x, \varphi, \sigma)$ where $x \in \mathbb{R}^4$ is a spacetime coordinate, φ is the physical content of the irreversible record formed (outcome value, detector state, particle type, etc.), and $\sigma > 0$ is the local entropy production associated with the commitment.

Definition A.2 (Proto-temporal history). A proto-temporal history is a pair $H = (\{C_i\}_{i \in I}, \prec)$ where $\{C_i\}$ is a finite or countably infinite set of commitment events and $\prec$ is a total order on I — the proto-temporal ordering. The space of all proto-temporal histories is denoted $\mathscr{H}_{proto}$.

The restriction to countable index sets is motivated by the commitment threshold λ_c: genuine commitment requires entropy production to cross a definite threshold, which discretises commitment events relative to the continuous background of environmental interactions. In the regime $\lambda \gg \lambda_c$ where classical spacetime is well-defined, the number of commitment events in any bounded spacetime region is therefore at most countable. In principle, a measure-theoretic generalisation allowing uncountable index sets with an appropriate σ-algebra on I is possible; the countable case is treated here as physically motivated and technically sufficient for the results of Appendix A.

The total order $\prec$ is the proto-time ordering of Section 2.1. It is not assumed to coincide with any coordinate time ordering on the spacetime positions x_i.

A.2 Committed Record Structure

Definition A.3 (Record space). The record space $\mathcal{S}$ is the set of triples $S = (E, \leq_c, d)$ where:

- E is a finite or countably infinite set of abstract event labels with associated physical content
- $\leq_c$ is a partial order on E representing causal precedence (the light-cone partial order)
- $d : E \times E \to \mathbb{R}$ is a function assigning invariant spacetime intervals $d(e_i, e_j) = \Delta s^2_{ij}$

Definition A.4 (Committed record map). The record map $\mathscr{R} : \mathscr{H}_{proto} \to \mathcal{S}$ sends each proto-temporal history $H = (\{C_i\}, \prec)$ to the triple:

$\mathscr{R}(H) = (\{C_i\}, \leq_c, d)$

where $\leq_c$ is determined by which events lie within each other's past light-cones (independently of $\prec$), and $d(C_i, C_j) = \eta_{\mu\nu}(x_i - x_j)^\mu (x_i - x_j)^\nu$ is the Minkowski interval between the spacetime positions of C_i and C_j. The proto-temporal ordering $\prec$ does not enter the definition of $\mathscr{R}$.

Remark on the structure of the Minkowski input. The appearance of the Minkowski metric $\eta_{\mu\nu}$ in Definition A.4 makes explicit what is assumed versus what is derived. The entry of $\eta_{\mu\nu}$ can be understood through a four-level hierarchy:

Level 1 — Primitive operational data. Commitment events generate locally observable facts: signal reachability between event pairs (defining causal precedence), local clock comparisons at coincident events (defining proper time increments), and radar-distance assignments from exchanged signals. These are the raw material of the record sector and require no prior metric.

Level 2 — Derived local structure. From the signal reachability relation, a causal partial order $\leq_c$ is defined. From local clock-and-signal procedures, an equivalence class of interval assignments is derived — observers disagree on coordinate distances but agree on the quadratic combination that yields proper time. The invariant signal cone and an interval class thereby emerge from Level 1 data.

Level 3 — Flat-limit completion. In the flat limit (vanishing fold-density gradients, homogeneous and isotropic inertial frames), the minimal local quadratic form compatible with Level 2 data and a single invariant signal speed is the Minkowski quadratic form. This is the minimal flat operational completion of the committed-record sector, as established in Section 4.1A. $\eta_{\mu\nu}$ encodes this completion — it is not an unconstrained metric postulate but the unique local quadratic form forced by the operational structure of committed records in the flat limit.

Level 4 — Current status. The exact derivation of $\eta_{\mu\nu}$ from VERSF fold-dynamics — including the identification $c_{signal} = c_{metric}$, which is required by operational closure of the record sector (Section 4.2) — remains an open derivation target addressed in [3]. The Corollary A.2 result is therefore conditional on Level 3 being the correct flat-limit completion, which is assumed here and justified by the argument of Section 4.1A.

The formal result is correctly stated as conditional: *given the minimal flat operational completion of Section 4.1A and the c-identification of Section 4.2, CRE yields exact Lorentz invariance at* *the level of observable records.* The value of the construction is not that it independently derives Minkowski structure, but that it demonstrates — rigorously, given these inputs — that proto-temporal ordering is outside the observable algebra and that the Lorentz group is the symmetry of the quotient.

A.3 Observable Algebra

Definition A.5 (Physical observable). A physical observable is a function $O : \mathscr{H}\_proto \to \mathbb{R}$ satisfying:

$$O(H_1) = O(H_2) \text{ whenever } \mathscr{R}(H_1) = \mathscr{R}(H_2)$$

The algebra of all physical observables is:

$$\mathcal{O} = \{ O : \mathscr{H}\_proto \to \mathbb{R} \mid \mathscr{R}(H_1) = \mathscr{R}(H_2) \Longrightarrow O(H_1) = O(H_2) \}$$

Elements of $\mathcal{O}$ are precisely the functions on $\mathscr{H}$_proto that factor through the quotient $\mathscr{H}$_phys = $\mathscr{H}$_proto/~. This is the same structure as the algebra of gauge-invariant observables in field theory: functions on the full space that are constant on equivalence classes. For the Lorentz invariance result of this paper — that all $O \in \mathcal{O}$ are invariant under Λ — the extensional definition is sufficient: the result holds for every function on $\mathscr{H}$_phys, whether or not it corresponds to a physically realisable measurement. (In a full dynamical treatment, the physical observable algebra would be a proper subalgebra of $\mathcal{O}$ carrying additional structure — a C*-algebra or Poisson algebra — and characterising it is an open problem for future work.)

A.4 The CRE Equivalence Relation

Definition A.6 (Commitment Reordering Equivalence). Two proto-temporal histories are equivalent under CRE if and only if their committed record structures coincide:

$$H_1 \sim H_2 \Longleftrightarrow \mathscr{R}(H_1) = \mathscr{R}(H_2)$$

This defines an equivalence relation on $\mathscr{H}$_proto. The physical history space is the quotient:

$$\mathscr{H}\_phys = \mathscr{H}\_proto / \sim$$

A.5 Main Structural Statement

Proposition A.1 (CRE as maximal equivalence). CRE is the maximal equivalence relation on $\mathscr{H}$_proto under which all elements of $\mathcal{O}$ are well-defined as functions on the quotient.

Proof sketch. Any equivalence relation $\approx$ under which all $O \in \mathcal{O}$ are well-defined must satisfy: $H_1 \approx H_2 \Longrightarrow O(H_1) = O(H_2)$ for all $O \in \mathcal{O}$. To show that $\approx$ must be at least as coarse as CRE, suppose $H_1 \approx H_2$. For any real-valued observable $O \in \mathcal{O}$, O is constant on $\approx$-classes, so $O(H_1) = O(H_2)$. Now consider observables of the form $O = f \circ \mathscr{R}$, where $f : \mathcal{S} \to \mathbb{R}$ is injective. Such f exists concretely: $\mathcal{S}$ is a set of triples (E, ≤_c, d), and one can take f to be projection onto any scalar component of d — for instance, $f(S) = d(e_1, e_2)$ for a fixed pair of event labels, which maps $\mathcal{S}$ into $\mathbb{R}$. This f is well-defined on $\mathcal{S}$ and injective on the relevant component. Since $f \circ \mathscr{R} \in \mathcal{O}$ (it is constant on CRE classes by construction), $H_1 \approx H_2$ forces $f(\mathscr{R}(H_1)) = f(\mathscr{R}(H_2))$ for every such f. Injectivity of f then implies $\mathscr{R}(H_1) = \mathscr{R}(H_2)$, which is precisely CRE. Therefore every ≈-class is contained in a CRE class, meaning CRE is at least as coarse as ≈ and is therefore the maximal such relation. □

Corollary A.2 (Lorentz invariance of $\mathcal{O}$). For any Lorentz transformation Λ acting on the spacetime coordinates of commitment events: $\mathscr{R}(\Lambda \cdot H) = \mathscr{R}(H)$ for all $H \in \mathscr{H}$_proto. This holds because Λ, identified in Section 4.2 as an element of the automorphism group of the committed record structure — the group of bijections preserving causal partial orders and invariant intervals — necessarily preserves the causal and metrical structure that defines $\mathscr{R}$. The identification of this automorphism group with the Lorentz group is established in Section 4.2 via the Zeeman–Alexandrov argument applied to the operationally defined invariant structure of committed records; it is not presupposed here. Therefore $\Lambda \cdot H \sim H$, and Λ acts as the identity on $\mathscr{H}$_phys. All $O \in \mathcal{O}$ are Lorentz invariant.

This is the formal content of the Lorentz Emergence Result of Section 4.2, stated at the level of the defined objects.