Discrete Commitment and the Necessity of Tick–Bit Structure

A Kernel-Level Derivation within the VERSF Framework

Keith Taylor *VERSF Theoretical Physics Programme*

Plain-Language Summary

One of the deepest questions in physics is what *time* actually is. In Newton's picture, time is a fundamental dimension — a kind of universal stage on which events play out, ticking forward at a fixed rate everywhere whether anything happens or not. Einstein replaced this with a more flexible picture in which time is bound up with space, but it remained a *dimension* — part of the fabric of reality, something that exists prior to and independently of events.

This paper argues for a different position: **time is not a fundamental dimension at all. It is something the universe constructs out of events.**

In the VERSF framework, the world is not built from continuous fields evolving smoothly through continuous time. It is built from **discrete events** — moments at which one possibility resolves into one definite outcome — and the structure of these events is captured by what we call *ticks* (the discrete moments themselves) and *bits* (the finite information each moment carries). Together, ticks and bits give the framework its name: TPB, *ticks per bit*.

Earlier VERSF papers introduced this Tick–Bit structure and showed that it does a lot of explanatory work: it accounts for the arrow of time, for the irreversibility of measurement, for thermodynamic entropy, and for the discrete character of quantum events. But those earlier papers had to *postulate* the Tick–Bit structure as an additional assumption layered on top of the framework's deeper foundations. A reader could reasonably ask: is this postulate really necessary, or is it a convenience smuggled in to make the rest of the theory work?

This paper closes that gap. We start from only the three deepest commitments of the VERSF framework — observer-invariant distinguishability, finite admissibility, and irreversible commitment — and we prove that the Tick–Bit structure is not a separate postulate at all. It is *forced* by those three commitments alone. A continuous substrate is incompatible with them. Reversible time is incompatible with them. Random or unstructured time is incompatible with them. The only structure compatible with all three is discrete ticks carrying finite bits of distinguishable content, ordered along worldlines.

The result has a striking philosophical consequence. Once you accept the three foundational principles, you cannot also accept time as a pre-existing dimension. The ordering of events comes *from* the irreversibility of those events, not from a clock running underneath them. The "tick rate" of any worldline is set by how often facts are produced along it, not by an external timekeeper. There is no universal "now" stretching across the cosmos — only the local order of facts along each worldline, with relativistic causal structure linking them.

This places VERSF within the *relational tradition* in the philosophy of time — the view, going back to Leibniz and developed in modern form by Mach, Reichenbach, Whitrow, Julian Barbour, Carlo Rovelli, and the causal-set programme, that time is constituted by relations between events rather than being the stage on which events occur. What is novel here is that VERSF doesn't *postulate* relationalism — it *derives* it. The relational view of time is not a philosophical preference of the framework; it is what the framework forces.

In the language of the broader programme: this paper turns Tick–Bit structure from an *axiom* into a *theorem*, and in doing so turns emergent time from a philosophical position into a structural consequence. Combined with our earlier overdetermination result for quantum mechanics, this means VERSF now derives, rather than assumes, both the discrete temporal substrate and the standard quantum formalism that lives on it. The framework's foundational footprint is correspondingly smaller — and its claim to inevitability correspondingly stronger.

Abstract

We address a remaining foundational gap in the VERSF programme: whether the discrete Tick–Bit (TPB) structure must be assumed, or whether it follows necessarily from the core kernel of the theory. Prior work established TPB as a consistent and uniquely constrained framework under operational finiteness axioms (A1–A4 in the original presentation), and demonstrated its explanatory power across quantum measurement, entropy, and emergent time. However, those results depended on an axiom set extending beyond the minimal VERSF kernel.

In this paper, we show that the TPB structure is not an independent assumption. Starting from the kernel principles — observer-invariant distinguishability (A0), finite admissibility (A1), and irreversible commitment (A2) — we derive the necessity of (i) discrete commitment events, (ii) finite distinguishability capacity per event, and (iii) a countable successor structure along worldlines. These jointly imply a Tick–Bit ontology with a well-defined ticks-per-bit ratio. We further establish a no-go result against continuous-time substrates — including the more sophisticated proposal of a continuous substrate equipped with a discrete commitment map — and a classification result showing that all alternative discrete ontologies either violate at least one kernel principle or reduce operationally to TPB.

The proof rests on a bridging principle we make explicit and defend at length: *operational distinguishability*. A distinction is operational if it makes a difference to some committed fact; A1's scope is operational distinguishability, not mathematical state-space cardinality. This principle links A0, A1, and A2 into a single coherent scope condition under which the no-go arguments operate. In one line: **ODP operationally identifies physical content with distinctions that affect commitment outcomes.**

A direct corollary is that **time is not a fundamental dimension within the VERSF framework**. The ordinal structure of events is induced by the irreversibility of commitment, not by an underlying time parameter; ticks are relational increments rather than metric absolutes; and no global simultaneity relation across worldlines is admitted by the kernel. The result therefore places VERSF within the relational tradition in the philosophy of time (Leibniz, Mach, Reichenbach, Whitrow, Barbour, Rovelli, causal-set theory) — but as a *derived* commitment, not a philosophical preference. Where prior relational accounts postulate that time is constituted by event-relations, the present result proves that any framework satisfying A0–A2 must take this view.

Combined with the overdetermination result for quantum mechanics, this leaves the kernel A0–A2 as the sole foundational input from which both the discrete temporal substrate and the quantum formalism follow.

Table of Contents

---

1. Introduction

The VERSF programme rests on a minimal kernel of three principles:

- **A0** — Observer-invariant distinguishability
- **A1** — Finite admissibility
- **A2** — Irreversible commitment

From this kernel, prior VERSF work has derived a substantial fraction of physical structure: thermodynamic irreversibility, the complex Hilbert state space, the Born probability rule, and the structural overdetermination of the quantum formalism (companion paper, *Kernel Minimality and Representation Overdetermination*).

A separate strand of the programme has developed the **Tick–Bit (TPB) framework**: a discrete-event ontology in which physical reality is structured as a sequence of irreducible commitment events ("ticks"), each of which encodes a finite quantum of distinguishable content ("bits"). The TPB framework gives an account of emergent time, the arrow of entropy, the locality of quantum measurement, and the discrete character of fact-formation.

A structural gap, however, has remained between these two strands. Previous TPB derivations relied on an extended axiom set — operational finiteness, locality of distinguishability, information bounds — that went *beyond* the kernel A0–A2. The natural question is whether this extension is necessary, or whether the TPB structure is already implicit in the kernel:

Is the Tick–Bit ontology a postulate added to the VERSF kernel, or a theorem derivable from it?

The present paper resolves this question. We prove that the kernel A0–A2 alone forces the existence of discrete commitment events, finite per-event distinguishability, successor-based temporal ordering, and a well-defined ticks-per-bit ratio. Continuous-time substrates are ruled out by a no-go theorem; alternative discrete ontologies are shown either to violate the kernel or to reduce operationally to TPB.

The result has two consequences. First, it shrinks the foundational footprint of VERSF: the kernel A0–A2 alone — three principles — generates both the discrete substrate (this paper) and, in conjunction with auxiliary route premises, the standard quantum formalism (companion paper). Second, it strengthens the claim to *inevitability* that the programme has been pursuing: the discrete Tick–Bit substrate is not one ontological choice among many compatible with the framework. It is the unique operational structure that survives the kernel's constraints.

The paper proceeds as follows. §2 restates the kernel and identifies what it does and does not specify, including a definitional preliminary on worldlines that is used throughout. §3 develops the scope condition (operational distinguishability) and the no-continuous-substrate theorem. §§4–8 give the constructive derivation: commitment events must be discrete (§4), each event must carry finite bit capacity (§5), events must form a successor structure (§6), ticks emerge as the minimal increments (§7), and a well-defined ticks-per-bit ratio follows (§8). §9 eliminates the alternative discrete and continuous ontologies. §10 states the main theorem. §11 ties the result to the broader programme. §12 makes the result empirically committed by stating the conditions under which it would be falsified. §13 states scope and limitations. §14 anticipates the most pointed referee objections and responds to each substantively.

---

2. Kernel Principles

2.1 The kernel A0–A2

We restate the three kernel principles in the form used throughout the VERSF programme.

A0 — Observer-Invariant Distinguishability. Physical content consists only of distinctions invariant under the group $\mathcal{G}$ of admissible observer transformations. Formally, the set of physical alternatives $\mathcal{A}$ is acted on by $\mathcal{G}$, and physical content is identified with the quotient $\mathcal{A} / \mathcal{G}$.

A1 — Finite Admissibility. Any bounded subsystem $\mathcal{S}$ admits a finite set of admissible alternatives:

$|\mathcal{A}(\mathcal{S})| < \infty$.

A2 — Irreversible Commitment. Physical facts are produced by an irreversible map

$\mathcal{C} : \mathcal{A} \to \mathcal{O}$

from the pre-commitment alternative space $\mathcal{A}$ to the outcome space $\mathcal{O}$, satisfying:

(i) $\mathcal{C}$ is many-to-one (resolution of alternatives), (ii) $\mathcal{C}$ is non-invertible at the level of physical content (irreversibility), (iii) $\mathcal{C}$ incurs a positive entropy cost $\Delta S \geq k_B \ln 2$ per bit of resolved distinction.

2.2 Worldlines and bounded subsystems (definitional preliminaries)

The arguments to follow refer repeatedly to "worldlines" and "bounded subsystems." We give operational definitions sufficient for the kernel-level derivation, deferring relativistic refinements to companion work.

Definition (Bounded subsystem $\mathcal{S}$). A bounded subsystem is a delimited collection of admissible alternatives — a region of the alternative space $\mathcal{A}$ to which A1 applies as a finite-cardinality condition. Boundedness here is logical (finite alternative-set cardinality after observer-invariant identification) rather than metrical; the kernel does not yet equip $\mathcal{A}$ with a metric.

Definition (Worldline). A *worldline* is a maximal chain of commitment events linked by the irreversible-precedence relation $\prec$ induced by A2(ii). Two events lie on the same worldline iff they are connected by $\prec$-comparability. The kernel forces $\prec$ to be a partial order on the global event set; worldlines are its maximal totally ordered subchains.

These definitions are minimal. They commit only to what A0–A2 force: subsystems are individuated by their finite alternative-content, and worldlines are individuated by $\prec$-connectivity. No metric, no embedding into a manifold, and no notion of "spatial location" are presupposed.

2.3 Immediate consequences

From A0–A2, the following are immediate:

- **(C1)** Physical reality consists of *resolved* distinctions, not pre-commitment superpositions.
- **(C2)** Per-event distinguishability is finite (from A1 applied to bounded events; formalised in §5).
- **(C3)** Resolution is irreversible and induces an ordering structure (from A2(ii)).

2.4 What the kernel does not yet specify

Crucially, A0–A2 do not yet specify:

- whether the commitment map $\mathcal{C}$ acts at *discrete* moments or as a *continuous* limit;
- whether per-event bit capacity is uniform across events;
- whether the order induced by A2 is total, partial, or branching;
- whether ticks are absolute or relational;
- the specific structure of $\mathcal{G}$ (which symmetries count as admissible observer transformations).

The work of §§3–8 is to show that the first four questions have *forced answers* once the kernel is taken seriously. The fifth (the structure of $\mathcal{G}$) is genuinely additional input — we will note where its specification matters and where the kernel-level result is independent of it.

---

3. The No Continuous Substrate Theorem

We first rule out the most common alternative ontology: a continuously resolved substrate.

3.0 The scope of A1 — operational distinguishability

Before stating the theorem, we clarify the scope at which A1 (finite admissibility) operates. This clarification is essential to what follows because A1 will do significant work in eliminating continuum-based alternatives, and a natural objection is that A1 is being applied beyond its proper domain — to *mathematical* state spaces rather than to *physical* distinctions.

We address this with an explicit bridging principle that makes the conjunction of A0, A1, and A2 operationally sharp.

Operational Distinguishability Principle (ODP). *A distinction is physical content iff it makes a difference to some committed fact. Equivalently: a distinction* $d \in \mathcal{A}$ *has physical content iff there exists some commitment process under which the outcome* $\mathcal{C}(\alpha)$ *depends, in an observer-invariant way, on whether the alternative* α *realises d.*

This principle is not an additional axiom; it is what A0, A1, and A2 jointly enforce when read together:

- **A0** confines physical content to $\mathcal{G}$-invariants;
- **A2** specifies that physical content is what gets *committed* (the image of $\mathcal{C}$);
- **A1** bounds the cardinality of admissible alternatives.

A distinction that never makes a difference to any committed fact is not in the image of $\mathcal{C}$ (by A2's many-to-one structure: it falls within a $\mathcal{C}$-fibre and is washed out at commitment), and is therefore not part of physical content. A distinction that only makes a difference under observer-dependent classifications fails A0. A distinction that makes an observer-invariant difference to some commitment is, by definition, an admissible alternative — and by A1, the set of such distinctions is finite per bounded subsystem.

ODP is therefore the operative reading of A1 throughout this paper. **A1 applies at the level of operational distinguishability — physically realisable distinctions that make a difference to some committed fact — not at the level of mathematical state-space cardinality.**

This yields the operative trichotomy: any ontology positing structure finer than the admissibility bound must either be

(a) **operationally meaningless** — the fine structure does not correspond to any observer-invariant distinction in the image of $\mathcal{C}$, and is therefore not physical content. This is consistent with A1 but classifies the fine structure as representational (mathematical scaffolding) rather than ontological.

(b) **physically accessible** — the fine structure does correspond to observer-invariant distinctions in the image of $\mathcal{C}$. In this case the structure falls under A1's scope and must be finite.

(c) **observer-dependent** — the fine structure corresponds to distinctions that depend on the choice of observer transformation in $\mathcal{G}$. This violates A0 directly.

A1 cannot be evaded by the manoeuvre "the substrate is continuous but only finite distinctions are physical." If the continuous substrate has no physical content beyond its finite distinctions, then the substrate itself is mathematical scaffolding rather than ontology — the ontology is the finite distinction set (case (a)). If the continuous substrate does have additional physical content, A1 forbids it (case (b)). If the additional content is observer-dependent, A0 forbids it (case (c)).

This scope clarification, together with the trichotomy, is what licenses the use of A1 in Proposition 1 below.

Proposition 1 (Continuous Substrate Failure)

No ontology in which physical content is carried by a continuous substrate is compatible with the kernel A0 ∩ A1 ∩ A2 — including the apparently weaker proposal of a continuous substrate equipped with a discrete commitment map.

Proof. The proposition has four cases. The first three address the direct continuous-resolution proposal; the fourth addresses the more sophisticated proposal that continuum and discreteness can be combined.

Case (i): infinite distinguishability is physically meaningful. Then for any bounded system $\mathcal{S}$, the set of admissible distinctions is uncountable. By the trichotomy of §3.0, this either falls into case (b) — fine structure with physical content, in which case A1 is directly violated (uncountable ≫ finite) — or case (a), in which the operative ontology is the finite distinction set rather than the continuum. In the latter sub-case the continuum is mathematical scaffolding, not ontology, and we are not actually committed to a continuous substrate at the level of physical content.

Case (ii): partial-resolution states are not physically meaningful but are mathematically representable. Then the only way the continuous substrate enters physical predictions is through a chosen representation — a smooth interpolation, a parametrisation, a threshold value, a coarse-graining scheme. If physical outcomes depend on such choices, then the map from substrate to facts is not invariant under $\mathcal{G}$: a different admissible representation (e.g. a different reparametrisation, a different smoothing kernel, a different threshold) yields different "physical" facts from the same underlying state. The substrate-to-fact map therefore fails the A0 invariance requirement — it falls into case (c) of the trichotomy.

Case (iii): commitment is approached only as a limit. Then for any time t at which one would say "commitment has occurred," there exists $t' < t$ at which commitment had not yet occurred and $t'' > t$ at which it had. The transition has no unique resolution point. A Cauchy-style sequence of partial-commitment states converging to an idealised "completion" does not, on its own, define a commitment event at any t: the limit specifies an asymptote, not a definite occurrence. This violates A2(ii), which requires a definite resolved state with a well-defined occurrence. (The proposal "the limit *is* the event" reduces to case (i) or case (iv): if the limit is taken to be ontologically real and finer than any finite stage, it adds infinite physical structure (case (i)); if it is taken to define a unique discrete event, the substrate has reduced to a discrete one (case (iv.b) below).)

Case (iv): continuous substrate with discrete commitment map (the *coarse-graining loophole*). This is the most sophisticated alternative, and the one a careful referee will press hardest. The proposal is: let the underlying state space be continuous (e.g. a smooth manifold of pre-commitment states), but let the commitment map $\mathcal{C}$ act discretely — perhaps via a threshold function, a measurement-induced collapse, a basin-of-attraction partition, or a coarse-graining scheme — so that the *facts produced* are discrete even though the *substrate* is continuous. This appears to satisfy A1 (only finitely many facts per bounded region) while preserving a continuum ontology underneath. Concrete instances of this strategy include Bohmian mechanics (continuous configuration space with effectively discrete record-formation through environmental decoherence), GRW-style spontaneous-collapse models (continuous Hilbert evolution with discrete collapse events), and many ψ-ontic interpretations.

The proposal fails for the following reason. A discrete commitment map acting on a continuous substrate must specify *which* continuous states map to *which* discrete outcomes. This specification requires a partition of the continuous substrate into preimage-classes, one per discrete outcome. The partition is defined by some scheme — threshold values, basin-of-attraction boundaries, measurement-operator eigenspaces, etc.

We now ask: is this partition observer-invariant?

(iv.a) **If the partition is not observer-invariant**, then different observers (in $\mathcal{G}$) compute different discrete outcomes from the same continuous substrate state. The discreteness of *facts* would be observer-dependent — falling into case (c) of the trichotomy and contradicting A0. The map $\mathcal{C}$ would be a representational artefact, not a feature of physical reality.

(iv.b) **If the partition is observer-invariant**, then by ODP the only physically meaningful structure on the continuous substrate is the discrete partition itself — the equivalence classes induced by $\mathcal{C}^{-1}$. By A0+A2, distinctions within a single equivalence class are not physical content (they make no difference to any committed fact, since all states in a class commit to the same outcome; they are therefore not in the image of $\mathcal{C}$ and fall under ODP case (a)). The continuous substrate is therefore *physically equivalent* to the discrete quotient space.

In case (iv.a), the proposal violates A0. In case (iv.b), the proposal is operationally identical to a discrete ontology — the continuous substrate is mathematical scaffolding above a discrete physical reality, in the same sense that a manifold's coordinate chart is scaffolding above a coordinate-free geometric object. The continuum is then not the *ontology* but a *redundant representation* of an underlying discrete structure.

The coarse-graining loophole therefore does not produce a third option between continuous and discrete ontologies. It collapses, under A0 ∩ A1 ∩ A2, into either a kernel-violating proposal (iv.a) or a discrete ontology in continuous clothing (iv.b).

Conclusion. All four cases lead either to direct contradiction with at least one kernel principle or to operational reduction to a discrete ontology. Hence no continuous substrate carries physical content under A0 ∩ A1 ∩ A2. ■

Continuous descriptions are admissible representations; discrete commitment structure is the invariant ontology.

Remark on smooth approximations. This result does *not* preclude continuous *descriptions* of physical processes — Schrödinger evolution, hydrodynamic flow, classical field theories. It precludes a continuous *ontology of fact-formation*. Smooth descriptions are admissible as coarse-grained approximations to the underlying discrete substrate. Operationally, "coarse-grained approximation" means: the smooth description reproduces the predictions of the discrete underlying theory to within experimental resolution at the scales of interest, in the same sense that hydrodynamics reproduces the bulk predictions of molecular dynamics. The case (iv) argument makes the ontological direction precise: the smooth description is the redundant representation; the discrete commitment structure is the ontology.

Remark on ψ-ontic interpretations. The case (iv) argument is also relevant to interpretations of quantum mechanics that treat the wavefunction ψ as a continuous physical entity. Such interpretations face exactly the dilemma of (iv.a) vs (iv.b): either the continuous wavefunction carries physical content beyond what is recovered by discrete measurement outcomes — in which case observer-dependence of that surplus content must be addressed — or it does not, in which case the wavefunction is a representational tool over a discrete fact-structure. VERSF takes the latter horn explicitly.

---

4. Discreteness of Commitment Events

Theorem 1 (Discrete Commitment Necessity)

Under A0 ∩ A1 ∩ A2, the commitment map $\mathcal{C}$ acts at discrete moments. Every physical fact corresponds to a *distinct, identifiable commitment event*, and along any worldline the events are isolated — there exists, at each event, a neighbourhood of the order $\prec$ containing no other event.

Proof. By Proposition 1, commitment cannot be realised continuously. The negation of "continuous" in this setting — given the structure of $\mathcal{C}$ as a many-to-one map from alternatives to outcomes — is that the map acts at isolated points. We must establish that these isolated points are well-defined as *events*: that they have identity, are countable, and admit no infinitesimal accumulation along any worldline.

(a) **Identity.** Each commitment event must produce a distinguishable record (A2(i)) and that distinguishability must be observer-invariant (A0). Hence each event has a well-defined identity in the quotient $\mathcal{A} / \mathcal{G}$.

(b) **Countability.** By A1, each bounded subsystem admits only finitely many alternatives; by A2(ii), each commitment is irreversible, so an infinite sequence of commitments in a bounded region would require infinite distinguishability accumulated within finite admissible space, contradicting A1. Hence within any bounded region, commitment events are at most finite — and across the union of bounded regions, at most countable.

(c) **No accumulation.** This is where we apply the bridge principle developed in §4.1. Suppose, toward contradiction, an accumulation point t^* of commitment events along some worldline w. Then for any neighbourhood $(t^* - \varepsilon, t^* + \varepsilon)$ in the $\prec$ order, arbitrarily many commitment events occur within it. By the argument of §4.1 below, accumulation requires arbitrarily fine *temporal distinguishability* — which is operationally distinguishable content, hence within A1's scope by ODP — and is therefore forbidden by A1.

Accumulation is therefore impossible: events along any worldline are isolated, in the sense that each has a neighbourhood (in the $\prec$ order, restricted to w) containing no other event. ■

4.1 The bridge from A1 to temporal distinguishability

The no-accumulation step in (c) requires an explicit bridge. The objection it must answer is: A1 bounds the cardinality of admissible alternatives in a bounded subsystem; why should this bound the *temporal density* of events along a worldline?

The answer rests on the observation that **the temporal positions of events along a worldline are themselves physical content within ODP's scope**. We argue this in two stages.

Stage 1: temporal positions are facts. The position of an event in the $\prec$ order along a worldline is, by A2(ii), an irreversible feature of that event: once event e_i has occurred at its position in the order, no observer transformation can move it. By A0, the order along a worldline is observer-invariant (a different observer cannot disagree about which of two facts on the same worldline came first; if they did, A0's invariance requirement would fail). Therefore the temporal position of each event, considered as a feature of the event-content, is observer-invariant content — a fact in the strict sense of A2.

Stage 2: distinguishing events at temporal positions is operational distinguishability. Distinguishing event e_i from event e_j on the same worldline is, by Stage 1, distinguishing two pieces of observer-invariant content — two facts about the order of commitments. Such a distinction makes a difference to the committed record (the record carries the order; reordering the events would change the record). By ODP, this distinction is in the scope of A1.

Now the no-accumulation argument is straightforward. If an accumulation point exists, then there are infinitely many events in some bounded neighbourhood, and distinguishing them all as separate events requires distinguishing infinitely many observer-invariant order-positions within that neighbourhood. By Stages 1–2, this is operational distinguishability within A1's scope, so the cardinality must be finite — contradicting the supposition of accumulation.

Equivalently, the contrapositive: if temporal distinguishability is *not* arbitrarily fine, then events at the supposed accumulation point cannot all be distinguished, and the supposed accumulation collapses to a finite set of events.

Hence accumulation is impossible.

The bridge can be summarised: **A1 governs operational distinguishability (ODP); the temporal positions of events along a worldline are operational distinguishabilities (Stages 1–2); therefore A1 bounds temporal density.** This is the rigorous form of the argument that L5 of the prior version flagged as in need of sharpening.

---

5. Finite Bit Capacity per Event

5.1 Per-event localisation of A1

A1 is stated for bounded subsystems $\mathcal{S}$: $|\mathcal{A}(\mathcal{S})| < \infty$. To apply it to events, we need a per-event localisation.

Definition (Per-event alternative space). Given a commitment event e occurring within a bounded subsystem $\mathcal{S}(e)$ (the smallest bounded subsystem containing the event in the sense of §2.2), the *per-event alternative space* is

$\mathcal{A}\_event(e) := \{\alpha \in \mathcal{A}(\mathcal{S}(e)) : \mathcal{C}(\alpha)$ is determined at e$\}$,

i.e. the set of admissible alternatives that the event e resolves between.

By construction $\mathcal{A}\_event(e) \subseteq \mathcal{A}(\mathcal{S}(e))$; by A1, $\mathcal{A}(\mathcal{S}(e))$ is finite; therefore $\mathcal{A}\_event(e)$ is finite. This finiteness is what the next theorem leverages.

Theorem 2 (Finite Bit Capacity)

Each commitment event encodes a finite number of distinguishable alternatives. The information per event is

$I = \log_2(N\_dist)$ bits,

where $N\_dist = |\mathcal{A}\_event| < \infty$ is the number of alternatives resolved at that event.

Proof. By §5.1, $\mathcal{A}\_event(e)$ is finite for any event e. Let $N\_dist := |\mathcal{A}\_event(e)|$. The Shannon information content of resolving among N_dist equiprobable alternatives is $\log_2(N\_dist)$ bits; for non-equiprobable distributions, the bound holds with equality replaced by $\leq$. In either case the per-event information is finite. ■

Interpretation. Each commitment event defines a *bit-level distinguishability unit*: an irreducible quantum of resolved information, with finite capacity bounded by the local admissibility cardinality.

---

6. Successor Structure of Events

Theorem 3 (Successor Ordering)

The set of commitment events along any worldline forms a *discrete totally ordered set with a successor function*: a countable set on which the irreversibility-induced order ≺ is total, asymmetric, transitive, observer-invariant, and discrete (each non-maximal element has an immediate successor; each non-minimal element has an immediate predecessor).

Proof. By Theorem 1, events are discrete and countable, and isolated within each worldline. By A2(ii), each event is irreversible — once event e_i has occurred, it cannot be undone — which induces an asymmetric relation ≺ on the event set: e_i ≺ e_j iff e_i must occur before e_j can. By A0 (and the argument of §4.1, Stage 1), this ordering is observer-invariant along any single worldline.

The relation ≺ restricted to a single worldline w is therefore:

- **asymmetric** (irreversibility, A2(ii));
- **transitive** (composition of irreversible steps is irreversible);
- **total** on w (any two events on the same worldline are ≺-comparable, by definition of worldline in §2.2);
- **observer-invariant** (A0 + Stage 1 of §4.1);
- **discrete** (Theorem 1: each event is isolated, so each non-maximal element has an immediate successor in the order, and each non-minimal element has an immediate predecessor).

A discrete asymmetric transitive total observer-invariant order on a countable set, with successors and predecessors as just described, is a *discrete totally ordered set with a successor function*. ■

6.1 Order type: (ℕ, S), (ℤ, S), or discrete totally ordered with successor

The proof of Theorem 3 establishes the structural properties listed but does not, on its own, fix the order type up to isomorphism. Three sub-cases are compatible with the kernel:

(i) **(ℕ, S):** the worldline has a least element (a "first event") and is unbounded above. This is the standard finite-past picture.

(ii) **(ℤ, S):** the worldline is bi-infinite: no first or last event. This is compatible with the kernel since A2(ii) does not, on its own, require an absolute beginning.

(iii) **Bounded discrete:** the worldline is finite, with both a first and a last event.

The kernel does not uniquely determine which sub-case obtains. Cosmological boundary conditions, additional thermodynamic-arrow assumptions, or substrate-specific results from the broader VERSF programme (e.g. the cosmological-constant derivation or commitment-capacity initial conditions) may select among them. The relevant kernel-level statement is the weaker one: *the order along each worldline is a discrete totally ordered set with successor function*.

When we use the shorthand "(ℕ, S)" elsewhere in this paper and the prior literature, we mean it as the *generic structural case* — the case in which a least element exists. Where the difference between (ℕ, S) and (ℤ, S) matters, we will be explicit.

6.2 Note on partial vs. total order across worldlines

The argument above gives a *total* order along a single worldline. Across distinct worldlines — spacelike-separated commitment events in the relativistic limit — the order is *partial*, not total: events may be unordered if their light-cones do not overlap. The kernel A0–A2 forces a total successor order *along each worldline*, but does not force a global total order across all worldlines. The transition from per-worldline successor structure to relativistic partial-order causal structure is treated in companion VERSF work on emergent Lorentz invariance and is not derived here.

---

7. Emergence of Ticks

7.1 Definition and well-formedness

Definition (Tick)

A **tick** is the minimal increment between successive commitment events along a single worldline:

$\tau := (e_i, e_{i+1})$,

where e_{i+1} is the immediate successor of e_i in the $\prec$ order restricted to the worldline.

Equivalently, a tick is a *successor edge* in the discrete totally ordered structure of Theorem 3.

Well-formedness

The definition is well-formed because: (i) by Theorem 1, events are discrete and isolated; (ii) by Theorem 3, each non-maximal event has a unique immediate successor along its worldline; (iii) by Theorem 2, the increment between successive events corresponds to the resolution of a finite quantum of distinguishable content. Each tick is therefore a well-defined, finite, observer-invariant increment.

7.2 Ticks as relational — the kernel underdetermines absolute scale

Ticks are *not* absolute temporal units. They are relational increments defined by the succession of commitment events. Two distinct worldlines may exhibit different tick rates — the number of ticks per unit of any external time parameter — and this rate-difference is precisely what becomes time dilation in the emergent relativistic regime.

We are careful here about the modal force of this claim. The earlier draft of this paper argued that A0 *forbids* an absolute tick scale on the grounds that different observers would have to agree on what a "unit tick" is in a non-invariant way. This formulation is too strong: whether an absolute scale violates A0 depends entirely on the structure of $\mathcal{G}$, which the kernel does not specify (cf. §2.4). If $\mathcal{G}$ = {identity}, an absolute scale is trivially invariant; if $\mathcal{G}$ includes Lorentz-type transformations (as required to recover special relativity in the broader programme), an absolute tick rate is not invariant and is forbidden by A0.

The honest kernel-level statement is therefore: **the kernel underdetermines tick scale**. The kernel forces ticks to be *ordinally meaningful* (Theorem 3) but does not, by itself, force them to be *metrically absolute* or *metrically relational*. The metric status is determined by $\mathcal{G}$. Insofar as the broader VERSF programme commits to a $\mathcal{G}$ containing relativistic transformations — as it must to recover empirical relativity — ticks are forbidden absolute scale and are thereby relational. But this is a programme-level commitment, not a kernel theorem.

The substantive kernel claim survives unweakened: ticks exist, are ordinally meaningful, and have finite information content per increment. Their metric status is settled at the next level of the framework's commitments, not here.

7.3 Time as emergent, not fundamental

The results of §§3–7 add up to a substantive position on the nature of time, which we now state explicitly.

Thesis. *Within the VERSF framework, time is not a fundamental dimension. The temporal ordering of events is induced by the irreversibility of commitment, not by an underlying time parameter against which events are indexed.*

The supporting structure is distributed across the preceding theorems: by Proposition 1, no continuous time-dimension carries physical content; by Theorem 1, events are discrete and constitute the only ontological substrate; by Theorem 3, the temporal order $\prec$ is defined by the irreversibility relation A2(ii), not by an antecedent time parameter; by §6.2, the order is total along each worldline but only partial across worldlines (no global simultaneity); by §7.2, ticks are at most relational increments, not metric absolutes.

The position is that events do not occur *in time*; rather, they *constitute* time. There is no clock running underneath the world ticking forward whether or not facts are produced; the production of facts *is* what ticking consists of.

Relation to the relational tradition. This places VERSF within the relational tradition in the philosophy of time, with antecedents in Leibniz's relational ontology, Mach's critique of

Newtonian absolute time, Reichenbach's analysis of temporal asymmetry as derived rather than primitive (Reichenbach 1956), Whitrow's relational treatment of temporal succession (Whitrow 1961, 1980), and modern formulations including Julian Barbour's *Platonia* (Barbour 1999), Carlo Rovelli's relational quantum mechanics and timeless framework (Rovelli 1996, 2018), and the causal-set programme of Bombelli, Lee, Meyer, and Sorkin (1987). Each of these traditions takes the position that time is constituted by relations between events rather than being a stage on which events occur.

What is novel in the present result is the *modal status* of the position. In the prior relational accounts, the relational view of time is postulated — taken as a foundational commitment of the framework. In VERSF, the relational view is *derived*: it follows from the kernel principles A0–A2 by Theorems 1–3, and any framework that satisfies those principles must take this view. The kernel does not permit a substantivalist time-as-dimension reading; such a reading would violate at least one of A0, A1, or A2. The relational view is therefore not a philosophical preference of VERSF but a structural consequence of its foundations.

Relation to spacetime. The present paper derives only the temporal structure along a single worldline. The transition to a relativistic spacetime — in which spacelike-separated events are causally unordered, light-cone structure determines simultaneity surfaces, and Lorentz invariance emerges — requires additional structure beyond A0–A2 and is treated in companion VERSF work on emergent Lorentz invariance and the proto-time / physical-time distinction. The present result is consistent with that programme but does not, by itself, derive spacetime; it derives the per-worldline temporal substrate on which spacetime is built.

A philosophical caution. The thesis is about the *ontological* status of time, not about the practical use of time-coordinates in physics. Smooth time-coordinates remain admissible as coarse-grained representations (cf. the Remark on smooth approximations after Proposition 1), in the same way that fluid dynamics remains admissible as a coarse-grained representation of molecular motion. The position is that the ontology underneath is discrete event-relations; the smooth time-coordinate is the calculational tool on top.

---

8. The Tick–Bit Ratio (TPB)

Theorem 4 (TPB Necessity)

Any physical system satisfying A0 ∩ A1 ∩ A2 admits:

(i) a discrete sequence of ticks (Theorems 1, 3, §7), (ii) a finite number of bits per tick (Theorem 2), (iii) a well-defined ratio

```
TPB(𝑆) := ticks per bit of resolved distinguishable content
```

within any bounded subsystem S over any finite extent of the worldline ordering.

Proof. (i) is established by Theorems 1 and 3 together with the definition in §7.1. (ii) is Theorem 2. (iii) follows from the conjunction of (i) and (ii): given any bounded subsystem $\mathcal{S}$ over a finite extent of the worldline ordering, both the number of ticks and the total bits resolved are finite (by A1, applied to the bounded subsystem), so their ratio is well-defined and finite. ■

Interpretation. TPB is not a postulated dimensionless ratio; it is a derived structural quantity expressing the "coarseness" of fact-formation in a given subsystem — how many discrete commitment events are required to resolve one bit of distinguishable content. Its variability across subsystems and worldlines is precisely the freedom that allows VERSF to recover relativistic time dilation, thermodynamic temperature gradients, and the variable resolution scales seen across physical regimes.

---

9. Elimination of Alternative Ontologies

We now systematically eliminate the alternatives. Each candidate is shown either to violate at least one kernel principle or to reduce operationally to TPB.

9.1 Continuous time

Already eliminated by Proposition 1.

9.2 Random or irregular event timing

A model in which events occur at random or irregular moments — without successor-ordered structure — would violate A0 if the irregularity is observer-dependent (case (c) of §3.0's trichotomy), or A2(ii) if the irregularity permits non-invariant orderings of facts. Either way, the kernel forces succession.

A model in which events occur irregularly *but with an invariant order* simply *is* TPB with non-uniform tick-spacing (where "tick spacing" is itself relational; cf. §7.2). This collapses to the TPB ontology.

9.3 Reversible time

A reversible event structure would allow commitment events to be undone, directly violating A2(ii). Eliminated.

9.4 Global synchronous updates

A model in which all events across the universe update synchronously — a "universal clock" — requires an observer-invariant simultaneity relation between spacelike-separated events. Such a relation does not exist within the admissible-transformation group $\mathcal{G}$ once $\mathcal{G}$ is committed to including Lorentz-type transformations (as the broader programme requires). Hence global synchronous updates violate A0 in any $\mathcal{G}$ containing such transformations. (At the strict kernel level — without further specifying $\mathcal{G}$ — global synchronous updates are not yet ruled out; but they are ruled out by the minimal extension of the kernel to relativistic $\mathcal{G}$, and that extension is required to recover empirical physics.)

9.5 Fixed-step time with uniform bit cost

A model in which every tick resolves exactly the same number of bits — say, exactly one bit per tick — fixes TPB at a constant value. This is *consistent* with the kernel but is a special case of the general TPB framework; it does not constitute a distinct ontology. Furthermore, fixed-uniform TPB is incompatible with empirical features that VERSF reproduces — variable measurement resolution, regime-dependent thermodynamic structure, relativistic dilation — so while not kernel-violating, it is empirically excluded.

9.6 Branching or non-linear time

A model in which events do not form a successor structure but rather a branching tree — e.g. many-worlds-style ontologies in which every commitment splits the worldline — preserves the per-branch successor structure of Theorem 3 within each branch. Across branches, however, an observer-invariant ordering between events on distinct branches does not exist, since the branches are causally isolated. Branching ontologies therefore *do not violate the kernel within each branch* but require an additional postulate (the branching rule) beyond A0–A2.

The kernel by itself does not adjudicate between single-branch and multi-branch ontologies; it forces the *per-worldline* TPB structure of Theorem 3 in either case. The branching question is orthogonal to the present derivation and is addressed elsewhere in the programme.

This connects to a broad class of contemporary interpretations of quantum mechanics — Everettian many-worlds, consistent-histories formulations with multiple compatible histories, and variants of decoherent-histories. The present result is silent between these and single-branch alternatives; it constrains only the per-branch substrate.

9.7 Classification result

Every alternative to TPB-structured time either: (a) violates at least one of A0, A1, A2 (possibly conditional on a relativistic $\mathcal{G}$ in case 9.4); or (b) reduces operationally to TPB with non-uniform tick spacing; or (c) preserves per-worldline TPB structure and adds an orthogonal postulate (branching).

In particular, no kernel-compatible alternative exists at the level of per-worldline event structure: a structure isomorphic to the TPB ontology is forced.

---

10. Main Result

Theorem 5 (Kernel → TPB Theorem)

A0 ∩ A1 ∩ A2 ⇒ discrete tick–bit structure.

More precisely: any physical system satisfying the VERSF kernel admits, along every worldline, a countable discrete totally ordered sequence (with successor function) of irreversible commitment events, each carrying finite bit capacity, with a well-defined ticks-per-bit ratio. No alternative ontology compatible with the kernel exists at this level of structure.

Proof. Assemble Theorems 1–4. Theorem 1 gives discreteness and isolation; Theorem 2 gives finite bit capacity; Theorem 3 gives the discrete totally ordered successor structure; Theorem 4 gives the well-defined TPB ratio. The classification result of §9 establishes that no kernel-compatible alternative exists at the per-worldline level. ■

Clean form

The result admits a referee-clean restatement, which we record here as the publication-grade form of the theorem:

Any ontology in which physical facts are

- *observer-invariant (A0),*
- *finitely distinguishable (A1, via ODP),*
- *irreversibly produced (A2),*

must reduce to a discrete, successor-ordered event structure with finite per-event information capacity. This structure is isomorphic to the Tick–Bit ontology.

Interpretation

The Tick–Bit ontology is not an additional assumption layered on top of the VERSF kernel. It is the **unique operational structure compatible with the kernel** at the level of per-worldline event organisation.

---

11. Relation to Prior Work

11.1 The previous TPB derivation

The original TPB derivation in *Foundations of Tick–Bit Dynamics* established:

(A1, A2, A3, A4) ⇒ TPB,

where A3 and A4 were operational finiteness and locality of distinguishability, treated as additional axioms beyond the minimal kernel.

11.2 The present result

We have shown:

(A0, A1, A2) ⇒ (A3, A4 as derived) ⇒ TPB.

Specifically:

- A3 (operational finiteness) follows from A1 applied to bounded events under ODP (Theorem 2 with the per-event localisation of §5.1);
- A4 (locality of distinguishability) follows from A0 applied to the per-worldline order via Stage 1 of §4.1 (Theorem 3, §6.2).

Therefore the chain is:

Kernel A0–A2 ⇒ derived A3–A4 ⇒ TPB,

and the original axiom set A1–A4 is now seen to be redundant: A3 and A4 are theorems, not axioms.

11.3 Relation to the overdetermination paper

The companion paper, *Kernel Minimality and Representation Overdetermination*, establishes that the standard quantum formalism — complex Hilbert space and the Born rule — is forced by the kernel A0–A2 conjoined with multiple independent constraint systems acting on it. The two results are complementary:

- the overdetermination paper shows that the kernel forces the *kinematic core* of quantum theory (state space + probability rule);
- the present paper shows that the kernel forces the *temporal substrate* on which that kinematics lives.

Together, they entail:

A0–A2 ⇒ (TPB substrate) + (complex Hilbert space, Born rule).

The kernel A0–A2 is the sole foundational input from which both the discrete event structure and the quantum formalism follow.

11.4 In one line

The earlier TPB derivation showed that ticks and bits suffice for physics; the present result shows that ticks and bits are forced by the deepest principles of physics.

12. Falsification Conditions

A structural-foundations result of this kind must commit itself to empirical conditions that would defeat it. We list three.

(F1) Demonstration of operationally relevant continuous structure. If an experiment showed that physically relevant predictions depend on continuum-fine distinctions in pre-commitment substrate states — distinctions that make an observer-invariant difference to committed records but lie below any finite admissibility cardinality — Proposition 1 would be falsified. This would require the demonstration of an experimentally testable prediction sensitive to substrate distinctions that are uncountable per bounded subsystem.

(F2) Observation of temporal accumulation. If an experiment showed an unboundedly dense sequence of commitment events along a single worldline — measurements producing observer-invariant records at temporal positions arbitrarily close to one another, with all positions distinguishable as separate facts — Theorem 1 would be falsified. (The standard form of such a falsification would be a sequence of measurements with sub-Planckian temporal separations whose order is operationally ascertainable.)

(F3) Operational dependence of facts on a non-invariant smooth-time parameter. If experimental predictions about facts depended ineliminably on a smooth time parameter not reducible to event-relations — i.e. if removing the parameter and using only event-relations changed the predictions — the relational thesis of §7.3 would be falsified.

None of (F1)–(F3) is technically near-feasible at present (all involve sub-Planckian probing or substrate-level access). What matters for the foundational status of the result is that the conditions are *specifiable* and *in-principle testable*: the result is empirically committed, not merely formal.

13. Limitations and Scope

We are explicit about what this paper does and does not establish.

(L1) Per-worldline structure only. The TPB structure is derived for a single worldline. The transition to relativistic causal structure across multiple worldlines (light-cone partial order, emergent Lorentz invariance) requires additional structure beyond A0–A2 and is treated in companion work.

(L2) Branching is orthogonal. §9.6 makes explicit that the kernel does not adjudicate between single-branch and multi-branch ontologies. Both are kernel-compatible at the per-worldline level. The branching question is independent of TPB-necessity.

(L3) Tick scale is not absolute, but only conditional on relativistic $\mathcal{G}$. §7.2 establishes that the kernel does not, on its own, force ticks to be metrically absolute or metrically relational; the metric status is settled by the structure of $\mathcal{G}$. Insofar as the broader programme commits to a relativistic $\mathcal{G}$, ticks are relational; but this is not a strict kernel theorem.

(L4) Bit capacity is not uniform. Theorem 2 establishes that per-event bit capacity is *finite* but not that it is *uniform across events*. Variable bit capacity is in fact required by the broader programme (it is what makes TPB itself a non-trivial dimensionless ratio rather than an absolute constant).

(L5) Order type is not uniquely fixed. §6.1 establishes that the kernel forces a discrete totally ordered set with successor function on each worldline, but does not adjudicate between ($\mathbb{N}$, S), ($\mathbb{Z}$, S), and bounded-discrete sub-cases. Selection among these requires additional cosmological or thermodynamic input from the broader programme.

(L6) Global synchrony is ruled out only in relativistic $\mathcal{G}$. §9.4 establishes that global synchronous updates violate A0 once $\mathcal{G}$ is committed to relativistic transformations. At the strict kernel level (without specifying $\mathcal{G}$), global synchrony is not yet eliminated.

(L7) No empirical predictions are made in this paper beyond the falsification conditions of §12. This is a structural-foundations result. Empirical content arises in the experimental-prediction papers of the broader programme.

---

14. Anticipated Referee Objections and Responses

A foundations paper of this kind invites scrutiny on multiple fronts. We anticipate the most pointed objections and address them directly. Each response is substantive rather than deflective and cross-references the paper's existing arguments where the relevant work has already been done.

Q1. Isn't ODP circular? You define "physical content" via commitment, then use that definition to force a discrete commitment structure.

This is the most dangerous objection and we treat it explicitly. ODP is not a free additional axiom; it is the operative reading of A1 forced by its conjunction with A0 and A2 (§3.0). The structure of the argument is:

- A0 confines physical content to $\mathcal{G}$-invariants;
- A2 says physical content is what gets *committed* (the image of $\mathcal{C}$);
- A1 bounds the cardinality of admissible alternatives.

ODP is the unique reading of A1 consistent with A0 and A2 simultaneously: A1 must apply to invariants (A0) that make a difference to commitment (A2). Any alternative reading either contradicts A0 (by counting non-invariant distinctions), contradicts A2 (by counting distinctions that don't make any difference to facts), or is degenerate (counts representational scaffolding, in which case A1 is trivially violated by any framework rich enough to include redundant representations).

The objection has force only if ODP is taken as an *additional* assumption. As a *derived* reading of A1's scope, it is what the kernel forces — and the no-go arguments then proceed without circularity.

Q2. Why these three principles? Couldn't a weaker kernel work?

The kernel is selected because removing any of the three appears to break essential structure. Specifically:

- Without A0, "physical content" is observer-relative, and no observer-invariant theory exists.
- Without A1, distinguishability is unbounded, and Proposition 1's elimination of continuous substrates fails (the operative ontology can be uncountable).
- Without A2, commitment is reversible, and the entire ordinal structure of Theorem 3 collapses.

We do not claim *strict* minimality. The companion paper (*Kernel Minimality and Representation Overdetermination*, §6) makes the analogous claim about quantum kinematics in terms of *relative minimality*: within the current derivational landscape, A0–A2 is the smallest assumption set under which the present route operates. Future work may identify a logically weaker kernel; this would strengthen the framework rather than undermine the present result.

Q3. A1 looks suspicious. Quantum field theory in a bounded volume has uncountably many states.

A1 governs operational distinguishability, not Hilbert-space dimension (§3.0, ODP). A bounded volume in QFT has uncountably many *mathematical* states, but the number of operationally distinguishable outcomes — distinct measurement records that can be written down with finite precision — is finite for any finite-resolution apparatus. Bekenstein-style holographic bounds and the entropy bounds in algebraic QFT both point in this direction. ODP is the reading consistent with operational physics; the uncountability of the state space is mathematical scaffolding (case (a) of the §3.0 trichotomy).

Q4. The §4.1 bridge argument is suspicious. Why are temporal positions of events along a worldline observer-invariant facts?

Because A0+A2 jointly force this. A2(ii) says facts are irreversible; if event $e_i \prec e_j$ on a worldline, then by A2 that order is fixed once committed. A0 says physical content is observer-invariant; if a different observer transformation could reorder e_i and e_j, that transformation would change which fact came first — making the order observer-dependent and therefore not part of physical content (A0). But the order *is* part of physical content: it is a feature of the committed records (which fact was recorded before which).

The argument is therefore: A2 forces the order to be a feature of the committed content; A0 forces that feature to be observer-invariant. There is no circularity because we are not assuming a temporal ordering exists; we are showing that any kernel-compatible ordering must be observer-invariant.

If the referee pushes further — "what if there's no temporal ordering at all?" — that is a separate question, handled by the no-accumulation argument (Theorem 1) plus Theorem 3, which constructs the order from A2's irreversibility relation rather than from a presumed time parameter.

Q5. What about reversible unitary evolution between measurements? Doesn't quantum mechanics violate A2?

A2 governs commitment, not all dynamics. Unitary evolution is the dynamics of *pre-commitment* alternatives — the structure on $\mathcal{A}$ prior to action by $\mathcal{C}$. A2 specifies the irreversibility of the commitment step itself (the "measurement" or "fact-formation" event), not the reversibility of the evolution between commitments. This distinction matches standard textbook quantum mechanics: between measurements, evolution is reversible; at measurement, the outcome is irreversibly recorded.

The companion paper on quantum-formalism overdetermination gives a more detailed treatment of this division. The present paper inherits it.

Q6. The discrete-vs-continuous distinction is unfalsifiable. If smooth descriptions reproduce all observations, what does discreteness add?

It adds the claim that the smooth description is a coarse-grained representation of underlying discrete structure, not a fundamental ontology. This claim is *empirically committed* in the conditions specified in §12: an observation of operationally relevant continuous structure (F1), or temporal accumulation (F2), or ineliminable smooth-time dependence (F3) would falsify it.

That these conditions involve sub-Planckian probing or substrate-level access does not make them in-principle untestable — it makes them *currently inaccessible*. Foundational results often have this character; the same holds for many predictions of quantum-gravity programmes. The result is committed to specifiable conditions for its defeat; that is what empirical commitment requires.

Q7. The order-type result is incomplete. You don't fix (ℕ, S) vs (ℤ, S).

We acknowledge this explicitly (§6.1, L5). The kernel forces a discrete totally ordered set with successor function on each worldline, but does not adjudicate among the sub-cases. Selection requires additional structure — cosmological boundary conditions, thermodynamic-arrow assumptions, or programme-specific results from elsewhere in VERSF.

This is a feature, not a bug. A kernel that *forced* (ℕ, S) over (ℤ, S) would have to import additional structure (a "first event," or an arrow-of-time direction beyond A2's local irreversibility) that the kernel does not contain. The honest report is "discrete totally ordered with successor function," and we make it.

Q8. How does discrete substrate reconcile with Lorentz invariance?

This is treated in companion work on emergent Lorentz invariance, not the present paper (§13, L1). The relevant result there is that Lorentz invariance arises as an emergent symmetry of the coarse-grained smooth description over discrete event-relations — not as a fundamental symmetry of the substrate. This is structurally similar to the way fluid Galilean invariance emerges from molecular dynamics that lack any sharp Galilean symmetry at the molecular level.

The present paper is silent on this. It derives only the per-worldline structure; the cross-worldline causal structure that supports Lorentz invariance requires additional work.

Q9. Case (iv) of Proposition 1 is too binary. What if the partition is invariant under some subgroup of $\mathcal{G}$ but not all of $\mathcal{G}$?

A partial-invariance proposal sits between cases (iv.a) and (iv.b). Two responses.

First, A0 quantifies over all admissible observer transformations in $\mathcal{G}$ — physical content must be invariant under the *full* group, not a subgroup. Partial invariance under a proper subgroup means non-invariance under the full group, which puts the proposal in case (iv.a) and contradicts A0.

Second, if one wishes to weaken A0 to "invariant under some subgroup" — i.e. modify the kernel — that is a different theory, not a defeater of the present one. The present result is conditional on A0 as stated; weakening A0 produces a different framework whose properties would need separate analysis.

Q10. The relational thesis on time is too strong. Even if the ontology is event-relations, operational physics still uses smooth time coordinates everywhere.

We agree that smooth time coordinates remain operationally indispensable, and we say so explicitly (§7.3, "philosophical caution"). The thesis is about *ontological* status, not *practical use*. Smooth coordinates are admissible as coarse-grained representations in the same sense that fluid dynamics is admissible as a coarse-grained representation of molecular motion.

The thesis becomes substantive — and falsifiable — at the level of foundational predictions where the ontological commitment makes a difference: predictions about Planck-scale temporal structure, the discreteness of fact-formation rates in extreme regimes, and the absence of operational content at sub-event scales. These are the conditions of §12.

Q11. How is this different from causal-set theory?

Causal-set theory postulates a discrete causal order as a foundational ontology. The present result *derives* such an order (along worldlines) from three more abstract principles. Specifically:

- causal sets assume discreteness; we derive it (Theorem 1);
- causal sets assume the order; we derive it (Theorem 3);
- causal sets do not, by themselves, account for finite per-element information content; we derive it (Theorem 2).

The result is therefore complementary to causal-set theory — it supplies a foundational rationale for why a causal-set-style ontology is forced, rather than treating discreteness and order as primitive. The relativistic extension across worldlines connects to the causal-set programme more directly and is treated in companion work.

Q12. What does this add beyond existing QM reconstruction programmes (Hardy, Masanes–Müller, Chiribella et al.)?

Two differences. First, *scope*: existing reconstructions take quantum kinematics (Hilbert space, Born rule) as the target. The present paper targets the *temporal substrate* on which any kinematics lives, which is logically prior. Second, *modal status*: existing reconstructions establish that *some* sufficient axiom set forces quantum mechanics. The companion overdetermination paper establishes that the quantum formalism is forced by *multiple independent* such axiom sets sharing a common kernel. The present paper does the analogous work for the substrate.

Together, the two papers position VERSF as deriving both the substrate (this paper) and the formalism (companion) from a common minimal kernel — a structural unification not available in single-target reconstruction programmes.

Q13. The result feels like setup. Where's the new physics?

The "new physics" here is structural rather than predictive: the result moves Tick–Bit ontology from postulate to theorem (§11), and moves the relational view of time from philosophical preference to derived consequence (§7.3). Predictive consequences appear in the broader VERSF programme — the cosmological-constant derivation, the fine-structure-constant derivation, the no-go theorem on non-simplicial substrates, and so on — and are downstream of the present substrate result rather than contained in it.

A foundational result that turns axioms into theorems shrinks the assumption-base of every downstream prediction. That is the work this paper does.

---

15. Conclusion

We have shown that within the VERSF framework:

1. A continuous-substrate ontology — including the more sophisticated proposal of a continuous substrate equipped with a discrete commitment map — is incompatible with the kernel (Proposition 1, including the case (iv) coarse-graining loophole closure).
2. Commitment events must be discrete and isolated along each worldline (Theorem 1, with the bridging argument of §4.1).
3. Each event carries finite bit capacity (Theorem 2, with the per-event localisation of §5.1).
4. Events along any worldline form a discrete totally ordered set with successor function (Theorem 3); the specific order type — ($\mathbb{N}$, S), ($\mathbb{Z}$, S), or bounded-discrete — is not uniquely fixed by the kernel (§6.1).
5. Ticks emerge as the minimal increments between events; their metric status (absolute vs. relational) is settled by $\mathcal{G}$, not by the kernel itself (§7).
6. The TPB ratio is well-defined and structurally derived (Theorem 4).
7. All alternative ontologies either violate the kernel or reduce operationally to TPB (§9).
8. Therefore the kernel A0–A2 alone forces a structure isomorphic to the Tick–Bit ontology at the per-worldline level (Theorem 5).
9. As a direct corollary, time is not a fundamental dimension within the VERSF framework: it is the relational ordinal structure induced by the irreversibility of commitment events, defined per-worldline, with metric status conditional on $\mathcal{G}$ and no global simultaneity (§7.3).
10. The result is empirically committed: §12 specifies in-principle conditions under which it would be falsified.

The Tick–Bit substrate, previously postulated as part of the extended VERSF axiom set, is now seen to be a *theorem* of the kernel — not an assumption added to it. The relational view of time, previously a philosophical commitment of frameworks like causal-set theory and Rovelli's relational quantum mechanics, is now seen to be a structural consequence of the kernel rather than an additional postulate.

Final statement

Discrete tick–bit structure is not postulated in VERSF. It is forced by the minimal kernel of the theory. The relational view of time is not a philosophical preference of the framework. It is what the framework forces. Combined with the overdetermination of the quantum formalism, this leaves three principles — observer-invariance, finite admissibility, and irreversible commitment — as the sole foundational input from which both the temporal substrate of physics and its quantum kinematics follow.

Notation Summary

Symbol	Meaning
$\mathcal{K}$	the VERSF minimal kernel (A0, A1, A2)
$\mathcal{A}$	space of pre-commitment alternatives
$\mathcal{O}$	space of committed outcomes

Symbol	Meaning
$\mathcal{C}$	commitment map $\mathcal{A} \to \mathcal{O}$
$\mathcal{G}$	group of admissible observer transformations
$\mathcal{S}$	a bounded subsystem
$\mathcal{S}$(e)	smallest bounded subsystem containing event e
$\mathcal{A}(\mathcal{S})$	admissible alternative set for $\mathcal{S}$
$\mathcal{A}$_event(e)	per-event alternative space (§5.1)
ODP	Operational Distinguishability Principle (§3.0)
N_dist	number of distinguishable alternatives at an event
I	per-event information content (bits)
e_i	the i-th commitment event along a worldline
$\prec$	irreversible-precedence relation on events
w	a worldline (maximal $\prec$-totally-ordered chain)
($\mathbb{N}$, S)	natural numbers with successor function (one possible order type for a worldline)
($\mathbb{Z}$, S)	integers with successor function (alternative order type for a bi-infinite worldline)
τ	a tick — successor edge between events
TPB	ticks per bit (derived structural quantity)

Selected References

- Sorkin, R. D. (1994). Quantum mechanics as quantum measure theory. *Mod. Phys. Lett. A* 9, 3119–3127.
- Reichenbach, H. (1956). *The Direction of Time*. University of California Press.
- Whitrow, G. J. (1961, rev. 1980). *The Natural Philosophy of Time*. Oxford University Press.
- Barbour, J. (1999). *The End of Time: The Next Revolution in Physics*. Oxford University Press.
- Rovelli, C. (1996). Relational quantum mechanics. *Int. J. Theor. Phys.* 35, 1637–1678.
- Rovelli, C. (2018). *The Order of Time*. Riverhead.
- Bell, J. S. (1987). *Speakable and Unspeakable in Quantum Mechanics*. Cambridge.
- Penrose, R. (1971). Angular momentum: an approach to combinatorial space-time. In *Quantum Theory and Beyond*, Cambridge.
- Bombelli, L., Lee, J., Meyer, D., Sorkin, R. D. (1987). Space-time as a causal set. *Phys. Rev. Lett.* 59, 521–524.
- Earman, J. (1989). *World Enough and Space-Time: Absolute versus Relational Theories of Space and Time*. MIT Press.
- 't Hooft, G. (2016). *The Cellular Automaton Interpretation of Quantum Mechanics*. Springer.
- Wheeler, J. A. (1990). Information, physics, quantum: the search for links. In *Complexity, Entropy and the Physics of Information*, Addison-Wesley.

- Taylor, K. (VERSF programme): *Foundations of Tick–Bit Dynamics*; *Why a Fact-Producing Universe Must Satisfy Interference*; *Kernel Minimality and Representation Overdetermination*; *No-Go Theorem on Non-Simplicial Relational Substrates*; *The Two-Planck Cosmological Constant Derivation*; *Commitment-Capacity Density and the No Multi-Primitive Occupancy Theorem*. versf-eos.com.