Uniqueness of the Bit–Tick Ontology

Abstract

We prove that any physical theory satisfying four minimal requirements—finite distinguishability for bounded systems, operational definition of time via clocks, no surplus structure beyond operational access, and finite-resolution experimental capacity—admits a canonical projection onto a unique **bit–tick substrate**: bits as the binary information capacity of maximal measurement contexts ($n = \lceil \log_2 N\_dist \rceil$), and ticks as the successor structure on clock records.

This uniqueness is structural, not representational. We show that the **operational core** of any candidate ontology—its quotient under resource-bounded operational equivalence—is finite, sequentially updated along worldlines, and unique up to isomorphism. All alternative ontologies (fields, strings, particles, spin networks) either project onto this same substrate or contain empirically inert surplus structure.

The framework accommodates quantum mechanics naturally: within any maximal measurement context, the outcome algebra is finite Boolean with N_dist atoms encoding n bits of capacity; continuous amplitude parameters collapse under finite-resolution equivalence to finitely many distinguishable preparation classes. Mathematical consistency checks confirm reproduction of Planck's relation $E = hf$ and Boltzmann entropy $S = k\_B \ln \Omega$.

We establish that denying bits or ticks as primitives requires abandoning either finite information bounds (violating the Bekenstein-Hawking entropy), operational grounding (positing physically meaningful distinctions no experiment can access), or finite experimental resources (requiring infinite precision). The bit–tick ontology is thus the **unique minimal invariant** of any operationally grounded, resource-bounded physics.

Abstract for General Readers

What is the universe made of, at the deepest level? This paper proves that once we take seriously four basic facts—that physical systems can only be in finitely many perfectly distinguishable states, that time is whatever clocks measure, that physics should not posit distinctions no experiment could detect, and that real experiments have finite precision—we are forced into a uniquely determined picture.

Everything reduces to bits and ticks.

A **bit** is the most basic unit of information: one yes/no question, one binary choice. Whenever you perform a measurement with N perfectly distinguishable outcomes, the information capacity is $\log_2(N)$ bits—the number of yes/no questions needed to specify which outcome occurred. A measurement distinguishing 8 outcomes carries 3 bits of information (since $2^3 = 8$). Bits are not the outcomes themselves but the *information capacity* those outcomes represent.

A **tick** is the most basic unit of change along a clock's history: one recorded event following another. Every clock counts something—oscillations, decays, vibrations. We prove that any clock record has the mathematical structure of the natural numbers with the successor operation: 0, 1, 2, 3, … This is uniquely determined, not chosen.

Once you have bits and ticks, everything else follows. Energy measures how fast new distinctions are produced—bits per tick. Entropy counts how many different tick-histories are compatible with what we observe. The flow of time reflects accumulated ticks along worldlines.

Quantum mechanics fits naturally: while quantum systems have continuous parameters (the Bloch sphere for a qubit), finite experimental precision means only finitely many preparation classes are actually distinguishable. The continuum is representational convenience; the operationally accessible structure is discrete.

The bit–tick picture is what remains when you strip away everything beyond experimental reach. It is the unique, minimal foundation for physics.

---

1. Introduction: The Elimination Argument

What must any candidate physical ontology ultimately provide? At minimum:

1. **Differentiate outcomes.** It must assert: this detector fired, that one did not; this configuration differs from that one.
2. **Respect finite distinguishability.** Entropy bounds (Bekenstein, holographic principle) establish that any bounded region with bounded energy permits only finitely many perfectly distinguishable states.
3. **Ground time operationally.** Time cannot be a background parameter. Whatever "time" means must connect to what clocks measure—counts of reproducible physical transitions.
4. **Contain no surplus structure.** Physics should not posit distinctions that no possible experiment, even in principle, could access.
5. **Acknowledge finite resources.** Real experiments have finite precision. Operational equivalence must be resource-bounded, not idealized.

These requirements are forced by the demand that physics make contact with finite, operational experiments.

This paper demonstrates that once these constraints are imposed, the bit–tick ontology emerges as the **unique** minimal framework. We prove:

- Within any **maximal measurement context**, the outcome algebra is finite Boolean with N_dist atoms encoding $n = \lceil \log_2 N\_dist \rceil$ bits (Section 3)
- Along any **clock record**, temporal structure is isomorphic to $(\mathbb{N}, S)$—the natural numbers with successor (Section 4)

- Under **finite-resolution equivalence**, the operational core is finite and unique up to isomorphism (Section 6)
- All richer structures either **project onto** this core or contain **empirically inert overhead** (Section 6)
- **Quantum mechanics** fits naturally, with continuous parameters collapsing to finite equivalence classes (Section 7)
- **No alternative exists** without violating established physical principles (Section 10)

---

2. Axioms: The Four Constraints

We build from four deliberately weak assumptions.

Axiom 1: Finite Perfect Distinguishability

For any physically realizable system confined to a bounded region with bounded energy, the number of mutually **perfectly** distinguishable states is finite:

N_dist < ∞

where perfect distinguishability means: there exists a measurement that assigns different outcomes to the states with certainty (probability 1).

This follows from the Bekenstein bound, the holographic principle, and the finite dimension of effective Hilbert spaces under energy constraints.

Critical clarification. Axiom 1 concerns *perfect* distinguishability, not operational distinctness in general. A qubit has N_dist = 2 (only $|0\rangle$ and $|1\rangle$ are perfectly distinguishable), but a continuum of preparations on the Bloch sphere that yield different probability distributions. Axiom 1 alone does not imply a finite state space—that requires Axiom 4.

Axiom 2: Operational Time

Time is not given a priori. "Time" is defined operationally as what clocks measure, and every clock functions by registering reproducible physical transitions along its worldline.

Worldline scope. In relativistic settings, there is no global time order—only partial orders and observer-dependent foliations. We therefore define temporal structure **locally along worldlines**, not globally. Each clock record is a totally ordered sequence of events; the invariant structure is this ordering plus count.

Formally, let a clock worldline W have recorded events $\{e_0, e_1, e_2, \dots\}$ with total order $e_0 < e_1 < e_2 < \dots$. The temporal structure along W is the order type of this sequence.

Axiom 3: Operational Completeness (No Surplus Structure)

The ontology of a physical theory is complete if and only if it accounts for all operationally accessible distinctions and contains no structure that is in principle inaccessible.

Formally, define the **operational equivalence relation**:

s ∼ s′ iff s and s′ induce identical outcome statistics for all possible experiments

Then Axiom 3 requires that physical states correspond to equivalence classes under ∼.

Remark. Under idealized operational access (infinite precision), different quantum amplitudes are operationally distinct—they yield different probability distributions. Axiom 3 alone does not collapse continuous parameters. That requires:

Axiom 4: Finite Accessible Information

For any bounded experiment with total resources $\mathscr{R}$ (energy E, time T, apparatus size L), the mutual information extractable between preparation labels and measurement outcomes is bounded:

$I(\text{prep} : \text{outcome} \mid \mathscr{R}) \leq I\_max(\mathscr{R}) < \infty$

Status of this axiom. Axiom 4 is an *empirically motivated constraint*, not a theorem derived within the bit-tick framework itself. The motivation comes from established bounds in specific physical theories:

Bound	Statement	Source Framework
Holevo bound	$I \leq S(\rho) \leq \log_2(\dim \mathscr{H})$	Quantum mechanics
Bekenstein bound	$S \leq 2\pi ER/(\hbar c)$	GR + thermodynamics
Channel capacity	$C = \max I(X:Y) < \infty$ for finite-energy channels	Information theory
Quantum speed limit	Operations per time $\leq 2E/\pi\hbar$	Quantum mechanics

These bounds are empirically well-confirmed within their respective domains. Axiom 4 abstracts their common content: *finite resources yield finite information*. This is not circular—we take the empirical success of these bounds as evidence for A4, then show that A4 (combined with A1–A3) forces bit-tick structure.

The logical status is: **Empirical evidence → Axiom 4 → Bit-tick substrate theorem**

If future physics discovered violations of finite-information bounds (e.g., infinite channel capacity at finite energy), Axiom 4 would need revision. But all current evidence supports it.

Statistical equivalence relation. Define:

s ~_{ℛ,ε} s′ iff for all experiments feasible under ℛ, the total variation distance between outcome distributions satisfies $\|P_s - P_{s'}\|_{TV} \leq \varepsilon$

Theorem (Finite ε-Net). For any compact preparation space S and any $\varepsilon > 0$, the quotient $S/{\sim}_{\{\mathcal{R},\varepsilon\}}$ is finite.

Proof sketch. With finite samples N (bounded by time/energy via quantum speed limits), hypothesis testing can only distinguish distributions differing by more than $O(1/\sqrt{N})$ in total variation. The set of ε-distinguishable distributions forms a finite ε-net. By compactness, finitely many equivalence classes cover S. ■

Corollary. Given Axiom 4, finiteness of the operational core is derived, not assumed.

The Four Axioms Together

Axiom	Content	Status
A1	Finite perfect distinguishability	Empirical (Bekenstein bound)
A2	Operational time on worldlines	Methodological (operationalism)
A3	No surplus structure	Methodological (parsimony)
A4	Finite accessible information	Empirical (Holevo, channel capacity)

Axioms 1–3 are standard in operational approaches. Axiom 4 is grounded in empirically confirmed information-theoretic bounds. The bit-tick substrate theorem follows from all four.

Axiom Independence and Relationships

The four axioms are logically independent—none implies any other:

A1 does not imply A4. A system with $N_{dist} = 2$ (finite perfect distinguishability) can have infinitely many operationally distinct preparations (the Bloch sphere) if experimental precision is unlimited. A1 bounds the *orthogonal* states; A4 bounds the *distinguishable* preparation classes under finite resources.

A4 does not imply A1. A4 bounds extractable information given finite resources, but says nothing about whether infinite perfect distinguishability is possible in principle. A1 is a physical constraint (Bekenstein); A4 is a resource constraint.

A2 and A3 are methodological, not empirical. They could in principle be rejected by adopting non-operationalist or non-parsimonious stances. However, rejecting A2 divorces "time" from physical clocks (operationally meaningless), and rejecting A3 admits untestable ontological commitments (scientifically unmotivated).

Why all four are needed:

Without...	You get...
A1	No bound on N_dist; bits undefined
A2	No tick structure; time is background parameter
A3	Surplus structure survives; uniqueness fails
A4	Infinite preparation classes; core not finite

The bit-tick substrate theorem requires the conjunction of all four.

Alternative Routes to Axiom 4

Axiom 4 (finite accessible information) can be motivated independently through multiple physical arguments. This strengthens the case that A4 captures a robust physical fact, not a convenient assumption.

Route A: Minimal Sufficient Statistic (Decision-Theoretic)

Define the **operational state** as the minimal object that predicts all future outcome distributions within ($\mathscr{R}$, ε). This is the "minimal sufficient statistic" from decision theory.

Lemma (Minimal Predictive State). For any theory under ($\mathscr{R}$, ε), there exists a coarsest partition of preparations such that all members of a cell induce ε-close future outcome statistics for all feasible experiments. This partition is unique.

Proof sketch. Define $s \sim s'$ iff for all feasible experiments E under $\mathscr{R}$, $\|P_E(s) - P_E(s')\|_{TV} \leq \varepsilon$. This is an equivalence relation. The quotient $S/\sim$ is the coarsest partition with the required property. Uniqueness follows from the definition. ■

This partition is exactly $\mathrm{Core}_{\{\mathscr{R},\varepsilon\}}$. Its cardinality is bounded by the ε-covering number of the accessible distribution set (Appendix A.3). Therefore:

The "real" operational state space is finite—derived from pure decision theory without invoking Bekenstein or Holevo.

Route B: Landauer / Finite-Precision Thermodynamics

Any physical readout requires:

- Finite memory to store the result
- Finite energy to perform the measurement
- Operation above thermal noise floor

Landauer's principle: Erasing one bit of information requires dissipating at least $k_B T \ln 2$ of energy. Conversely, distinguishing states requires sufficient energy to overcome thermal fluctuations.

If arbitrarily fine distinctions were extractable at fixed resources:

- Infinite bits would be storable in finite memory (contradiction)
- Infinite precision would be achievable at finite energy (violates Landauer)
- Signals below thermal noise would be detectable (impossible)

Therefore physical state identification must collapse to finite equivalence classes at finite resources—independently of QM or GR bounds.

Convergence of routes: Both routes arrive at the same conclusion:

Route	Starting point	Conclusion
Information-theoretic (A4)	Holevo, Bekenstein, channel capacity	Finite ε-net of preparations
Decision-theoretic (Route A)	Minimal sufficient statistic	Finite partition = $Core_{\mathscr{R},\varepsilon}$
Thermodynamic (Route B)	Landauer, finite memory/energy	Finite distinguishable classes

The convergence from independent physical principles suggests Axiom 4 captures a robust, framework-independent constraint.

---

3. Contextual Boolean Algebras and Bit Atoms

3.1 The Problem with Global Boolean Structure

A naive approach would define "the distinction algebra" as all operationally decidable yes/no propositions, claim this is Boolean, and identify bits as its atoms.

This fails for quantum mechanics. The global structure of quantum propositions is an **orthomodular lattice**, not a Boolean algebra. Incompatible measurements (e.g., σ_x and σ_z for a qubit) do not share a joint event structure—there is no single Boolean algebra containing all quantum yes/no questions.

3.2 Measurement Contexts

Definition. A **measurement context** M is a specification of a complete measurement—a POVM or projective measurement that can be physically implemented. For each context M, let:

B_M = the Boolean algebra of outcome events for M

Within any single context, the outcome events *do* form a Boolean algebra: outcomes are mutually exclusive, jointly exhaustive, and satisfy classical logic.

Example (Qubit).

- Context M_1 = {measure σ_z}: B_{M_1} has atoms |0⟩⟨0| and |1⟩⟨1|
- Context M_2 = {measure σ_x}: B_{M_2} has atoms |+⟩⟨+| and |−⟩⟨−|

These are different Boolean algebras, not subalgebras of a common Boolean algebra.

3.3 Maximal Contexts: Atoms and Bits

Definition. A **maximal context** is a measurement that perfectly distinguishes N_dist states—the maximum possible for the system. By Axiom 1, N_dist < ∞.

Definition (Atoms). The **atoms** of B_M are the elementary outcomes—the N_dist mutually exclusive, perfectly distinguishable results of a maximal measurement.

Definition (Bits). The **bit capacity** of a maximal context is the number of binary questions needed to specify an outcome:

n = $\lceil \log_2$ N_dist$\rceil$

Bits are not the atoms themselves but the **binary generators** of the outcome algebra—the minimal yes/no distinctions from which all outcomes can be constructed.

Proposition (Contextual Structure). For any maximal context M of a system with N_dist perfectly distinguishable states:

B_M ≅ $\mathcal{P}$({1, …, N_dist})

This power set has:

- **N_dist atoms** (elementary outcomes)
- **n = $\lceil \log_2$ N_dist$\rceil$ bits** of information capacity
- **2^{N_dist} elements** in the full Boolean algebra

Clarification. The relationship between atoms and bits:

Concept	Count	Role
Atoms	N_dist	Elementary distinguishable outcomes
Bits	n = $\lceil \log_2$ N_dist$\rceil$	Binary generators; information capacity
Algebra elements	2^{N_dist}	All Boolean combinations of atoms

A system with N_dist = 8 perfectly distinguishable states has 8 atoms but only n = 3 bits of capacity (since $2^3 = 8$).

3.4 Context-Independence of Capacity

Theorem (Bit Capacity Invariance). All maximal contexts for a given system have the same number of atoms (N_dist), and hence the same bit capacity $n = \lceil \log_2 N\_dist \rceil$.

Proof. N_dist is a property of the system (maximum number of perfectly distinguishable states), not of any particular measurement. Any maximal measurement realizes exactly this capacity. Different maximal contexts yield isomorphic (though not identical) Boolean algebras with the same atom count. ■

3.5 Summary: Atoms, Bits, and Contextual Structure

To summarize the precise terminology:

- **Atoms** are the elementary outcomes of a maximal measurement (there are N_dist of them)
- **Bits** are the binary information capacity: $n = \lceil \log_2 N\_dist \rceil$ independent yes/no questions
- **Contextual Boolean algebras** B_M contain the logical structure of outcomes within measurement M
- The **global** event structure is orthomodular (non-Boolean), formed by pasting contextual algebras

Within any single maximal test, the structure is Boolean with N_dist atoms encoding n bits of information.

---

4. Worldline Ticks and Successor Structure

4.1 The Problem with Global Temporal Order

A naive approach would define dynamics as a single global sequence $\{C_0, C_1, C_2, \ldots\}$ and identify ticks as increments of this sequence.

This fails for relativistic physics. General relativity provides only a partial order on spacetime events—no global "next configuration." Different observers see different temporal orderings. Quantum mechanics (in some interpretations) involves branching structures without a single timeline.

4.2 Clock Records as Totally Ordered Sequences

Definition. A **clock** is a physical subsystem that registers a sequence of marker events along its worldline. A **clock record** is the totally ordered sequence of these events:

$W = (e_0, e_1, e_2, \ldots)$ with $e_0 < e_1 < e_2 < \ldots$

Key point. Within a single worldline, events *are* totally ordered (proper time provides this). The successor structure lives on worldlines, not on spacetime globally.

4.3 Ticks as Successor Increments

Definition. A **tick** (along clock record W) is the minimal increment between successive events: the transition from e_n to e_{n+1}.

Define the clock count function:

T_W(n) = n (the count of events up to the nth marker)

Any physical time measurement along W has the form $t = \alpha \cdot T\_W$ for some calibration constant α.

4.4 The Successor Theorem (Worldline-Scoped)

Theorem (Tick Uniqueness). Let W be any clock record satisfying:

- (T1) W is a totally ordered sequence of recorded events
- (T2) Events are atomic (no infinitesimal subdivisions)
- (T3) The sequence has a first element and no last element

Then the order type of W is isomorphic to **($\mathbb{N}$, S)** where $S: n \mapsto n+1$ is the successor function.

Proof.

1. (T1) gives total order
2. (T2) gives discreteness—no limit points between successive events
3. (T3) gives the initial segment and unboundedness
4. The unique countable, discrete, well-ordered set with minimum and no maximum is ($\mathbb{N}$, $\leq$)
5. The successor function S is the unique atomic increment operation on $\mathbb{N}$ ■

Corollary. Along any clock worldline, operational time reduces to counting discrete events—ticks—with continuous parameters being monotone reparametrizations.

4.5 Global Structure

While individual worldlines have $(\mathbb{N}, S)$ structure, the global picture is richer:

- Different worldlines may have different tick counts between shared events (cf. twin paradox)
- The global causal structure is a partial order, not a total order
- Tick structure is **local/worldline-specific**, not global

This is a feature, not a bug: it matches the structure of relativistic spacetime exactly.

---

5. Ticks-per-Bit: The Bridge Quantity

5.1 Definition

With contextual bits and worldline ticks as primitives, define the **ticks-per-bit ratio** along a worldline W:

$\beta_W := \Delta T_W / \Delta B$ **= ticks per bit of new distinguishability**

where ΔB measures the production of new distinguishable outcomes along W.

5.2 Physical Interpretation

Regime	Meaning
High β	Many ticks per bit → slow distinguishability production → gentle dynamics
Low β	Few ticks per bit → rapid distinguishability production → violent dynamics

5.3 Connection to Energy

Define **bits-per-tick rate**: $\gamma := 1/\beta$

The fundamental energy relation:

$$\mathbf{E = (\hbar/\tau_0) \times \gamma}$$

where τ_0 is a fundamental time scale (Planck time $\tau_P = \sqrt{\hbar G/c^5}$) provides a natural candidate).

Energy measures how rapidly distinguishability is produced per tick along a worldline.

---

6. The Core Universality Theorem

This section establishes the central uniqueness result with proper technical precision.

6.1 Statistical Operational Equivalence

Fix a resource bound $\mathscr{R}$ (energy E, time T, apparatus size L) and error tolerance $\varepsilon > 0$. The statistical equivalence relation from Axiom 4:

$s \sim_{\mathscr{R},\varepsilon} s'$ iff for all experiments feasible under $\mathscr{R}$, **$\|P_s - P_{s'}\|_{TV} \leq \varepsilon$**

By Axiom 4 (finite accessible information), the quotient **$S/\sim_{\mathscr{R},\varepsilon}$** is a finite ε-net.

6.2 The Operational Core

Definition. The **operational core** of an ontology M (relative to resource bound $\mathscr{R}$ and tolerance ε) is:

$\mathrm{Core}_{\mathscr{R},\varepsilon}(M) := M/\sim_{\mathscr{R},\varepsilon}$

the quotient under statistical operational equivalence—the finite set of preparation classes distinguishable at resolution ε under resources $\mathscr{R}$.

6.3 The Bit–Tick Substrate: Formal Definition

To make uniqueness precise, we define the **bit–tick substrate** as a structured object:

Definition (Bit–Tick Substrate). The bit–tick substrate $BTS_{\mathscr{R},\varepsilon}$ of an ontology M consists of:

1. **Maximal contexts:** The family $\{B_M\}$ of finite Boolean algebras, one per maximal measurement context M
2. **Preparation classes:** The finite quotient $S/\sim_{\mathscr{R},\varepsilon}$ of preparation classes under statistical equivalence
3. **Context-change maps:** For each maximal context M, the map

   $\Phi_M : S/\sim_{\mathscr{R},\varepsilon} \to \Delta_{K-1}$

   sending each preparation class to its coarse-grained outcome distribution in context M (where Δ_{K-1} is the (K-1)-simplex of distributions over $K = N_{dist}$ outcomes)

4. **Worldline tick structure:** The order type $(\mathbb{N}, S)$ along each clock record

Operational equivalence of two theories means equivalence of these maps Φ_M for all feasible contexts M.

6.4 Core Universality Theorem

Theorem (Uniqueness of Bit–Tick Substrate). Let M and M′ be any ontologies satisfying Axioms 1–4. Then:

1. **Core_{ℛ,ε}(M) is finite.** (From Axiom 4: finite ε-net under bounded mutual information)
2. **Maximal-context structure is Boolean.** Within any maximal measurement context, outcome events form a finite Boolean algebra B_M with N_dist atoms, encoding n = ⌈$\log_2$ N_dist⌉ bits of information capacity.
3. **Worldline temporal structure is (ℕ, S).** Along any clock record, the event order is isomorphic to the natural numbers with successor.
4. **The bit–tick substrate is unique up to isomorphism.** If M and M′ have the same N_dist and are operationally equivalent (i.e., their context-change maps {Φ_M} and {Φ′*M} agree for all feasible contexts under ℛ), then their bit–tick substrates BTS*{ℛ,ε}(M) and BTS_{ℛ,ε}(M′) are isomorphic as structured objects.

Remark. The condition "same N_dist" is not an additional constraint but a consequence of operational equivalence. Since N_dist is the capacity of maximal measurement contexts, it is itself operationally determined: two operationally equivalent theories must have identical N_dist.

Proof.

(1) Finiteness. By Axiom 4, mutual information I(prep:outcome|ℛ) is bounded. The quotient under ∼*{ℛ,ε} forms a finite ε-net of distinguishable preparation classes. Hence |Core*{ℛ,ε}(M)| < ∞. □

(2) Boolean structure. By Section 3, any maximal context M has N_dist perfectly distinguishable outcomes. The Boolean algebra of outcome events has N_dist atoms, hence B_M ≅ 𝒫({1,…,N_dist}). □

(3) Successor structure. By Section 4, any clock record satisfying (T1)–(T3) has order type (ℕ, S). □

(4) Uniqueness. Let M, M′ satisfy the axioms with equal N_dist and operational equivalence.

- Both have finite cores with the same cardinality (from (1) + operational equivalence)
- Maximal contexts in both yield Boolean algebras with N_dist atoms (from (2))
- Clock records in both have (ℕ, S) structure (from (3))
- Operational equivalence means the context-change maps agree: Φ_M([s]) = Φ′_M([s]′) for corresponding classes
- The substrates BTS_{ℛ,ε}(M) and BTS_{ℛ,ε}(M′) are therefore isomorphic as structured objects □

Corollary (Projection Property). Every candidate ontology M satisfying Axioms 1–4 admits a canonical projection:

$$\pi: M \to \mathrm{Core}_{\{\mathscr{R},\varepsilon\}}(M)$$

Any structure in M not preserved by π is empirically inert surplus (below statistical resolution ε or operationally inaccessible under $\mathscr{R}$).

6.5 Universal Property: BTS as Terminal Object

To make uniqueness maximally precise, we formulate it as a universal property.

Definition. The **operational prediction object** $E_{\{\mathscr{R},\varepsilon\}}(M)$ of an ontology M is the family of all feasible-context outcome distributions under ($\mathscr{R}$, ε):

$$E_{\{\mathscr{R},\varepsilon\}}(M) = \{ \Phi_M : S/\sim_{\{\mathscr{R},\varepsilon\}} \to \Delta_{\{K-1\}} \mid M \text{ a feasible maximal context } \}$$

This captures everything operationally extractable from M.

Theorem (Universal Receiver). There exists a unique (up to isomorphism) structured object $BTS_{\{\mathscr{R},\varepsilon\}}$ such that every admissible ontology M factors through it:

$$M \xrightarrow{\ \pi\ } BTS_{\{\mathscr{R},\varepsilon\}} \xrightarrow{\ \iota\ } E_{\{\mathscr{R},\varepsilon\}}(M)$$

where:

- π is the projection onto the bit-tick substrate
- ι is the canonical embedding (inclusion of substrate structure into predictions)

Moreover, any two such factorizations differ only by isomorphism of $BTS_{\{\mathscr{R},\varepsilon\}}$.

Translation: $BTS_{\{\mathscr{R},\varepsilon\}}$ is the *minimal* structured object that can carry all operational content. It is a terminal object in the category of "operational prediction carriers"—any other carrier either factors through it or contains redundant structure.

Proof sketch.

1. Existence: Construct $BTS_{\{\mathscr{R},\varepsilon\}}$ as in Definition 6.3.
2. Factorization: Any M projects via π to $\mathrm{Core}_{\{\mathscr{R},\varepsilon\}}(M) \subseteq BTS_{\{\mathscr{R},\varepsilon\}}$; the context-change maps Φ_M embed Core into $E_{\{\mathscr{R},\varepsilon\}}(M)$.
3. Uniqueness: If BTS′ is another such object, the factorization property forces BTS ≅ BTS′. □

This universal property makes BTS canonical, not merely "one encoding among many."

6.6 What the Theorem Does and Does Not Claim

Claims:

- Any operationally grounded, finite-resolution theory has a bit–tick substrate BTS_{$\mathscr{R}$,ε}
- The substrate includes Boolean algebras, preparation classes, context-change maps, and tick structure
- Theories agreeing operationally (same context-change maps) have isomorphic substrates
- The substrate is **unique up to isomorphism** as a structured object

Does not claim:

- That Core_{$\mathscr{R}$,ε}(M) literally equals $\{0,1\}^n$ (it's isomorphic, not identical)
- That global dynamics is ($\mathbb{N}$, S) (only worldline-local)
- That idealized continuous parameters are meaningless (only that they collapse under statistical equivalence)
- That two theories with same N_dist are automatically equivalent (they must also share context-change maps)

Critical distinction: Substrate uniqueness ≠ Dynamics uniqueness.

The bit-tick substrate is the *invariant kinematic structure* that any operationally grounded theory must possess. It does **not** determine which specific dynamics governs evolution on that substrate.

Concretely: QFT, string theory, loop quantum gravity, and other candidate theories may all project onto isomorphic bit-tick substrates while differing in their dynamics—the specific transition rules, coupling constants, and field equations. The substrate theorem tells us *what kind of structure* physics must have, not *which physical theory is correct*.

Is this trivial? No. Structural constraints are nontrivial even when they don't determine dynamics:

Structural constraint	What it constrains	What it leaves open
"Spacetime is Lorentzian"	Causal structure, signature	Field equations, matter content
"States form a Hilbert space"	Superposition, interference	Hamiltonian, interactions
"Substrate is bit-tick"	Finite distinguishability, operational time	Specific dynamics, constants

The bit-tick constraint rules out:

- Theories with infinite information in bounded regions
- Theories with non-operational time

- Theories with empirically inert ontological commitments
- Theories requiring infinite experimental precision

This is a substantive restriction on the space of possible physical theories, even though it does not select a unique dynamics.

This is analogous to how "spacetime is a Lorentzian manifold" constrains physics without determining whether GR, modified gravity, or some other theory governs curvature. The bit-tick substrate is a structural constraint, not a complete theory.

6.7 Diagrammatic Summary

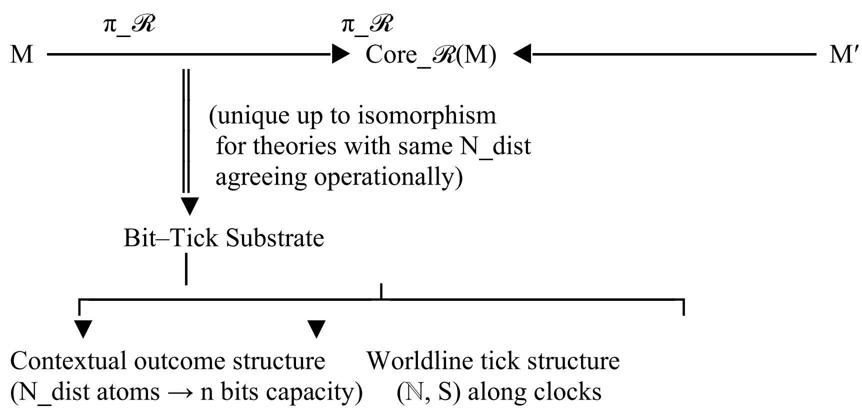

6.8 Application to Known Frameworks

Framework	Surplus Structure (collapses under π_ℛ)	Core Content
Classical mechanics	Sub-resolution phase space detail	Finite distinguishable macrostates
Quantum field theory	Infinite-dim Hilbert space; sub-resolution amplitudes	Finite effective states under energy/resolution cutoff
String theory	Extra dimensions, continuous moduli	Effective low-energy distinguishable states
Loop quantum gravity	Continuous spin network parameters	Finite area/volume eigenvalues above Planck scale
Causal set theory	Specific causal relations beyond resolution	Order type + cardinality ≈ bit–tick

7. Quantum Mechanics in the Bit–Tick Framework

7.1 The Qubit Subtlety

A qubit illustrates why all four axioms are needed:

Property	Value	Implication
N_{dist}	2	Only 2 perfectly distinguishable states
State space	Bloch sphere S^2	Continuum of preparations
Operational distinctness	Continuum	Different points → different probabilities

Under Axioms 1–3 alone, the qubit has finite N_{dist} but infinite operationally distinct preparations. **Axiom 4 is essential** for finiteness.

7.2 Finite-Resolution Collapse

Under statistical equivalence $\sim_{\{\mathscr{R},\varepsilon\}}$:

- Points on the Bloch sphere with outcome distributions differing by $\|P_s - P_{s'}\|_{TV} \leq \varepsilon$ are equivalent
- The sphere partitions into finitely many equivalence classes (an ε-net)
- Each class is one element of $Core_{\{\mathscr{R},\varepsilon\}}$

Proposition. For a qubit under resource bound $\mathscr{R}$ yielding statistical resolution ε in total variation distance, the number of distinguishable preparation classes is finite.

Explicit TVD-geometry connection. For a qubit measured in a fixed basis, the outcome probabilities are:

$$P(0) = (1 + z)/2, \quad P(1) = (1 - z)/2$$

where z is the Bloch coordinate along the measurement axis. The total variation distance between two preparations with coordinates z, z′ is:

$$d_{TV} = \tfrac{1}{2}|z - z'|$$

For a complete tomographic measurement (three orthogonal bases), the TVD between preparations at Bloch vectors **r**, **r′** satisfies:

$$d_{TV} \leq \|\mathbf{r} - \mathbf{r}'\|/2$$

Thus an ε-ball in TVD space corresponds to a ball of radius ~2ε in Bloch geometry. The covering number of S^2 by balls of radius δ is $O(1/\delta^2)$, yielding:

$$|S^2/\sim_{\{\mathscr{R},\varepsilon\}}| \sim O(1/\varepsilon^2)$$

This bound is tight up to constants.

7.3 Amplitudes as Sub-Resolution Parameters

Under Axiom 4, continuous amplitude parameters are **sub-resolution structure**:

Amplitudes differing by less than $\delta(\mathscr{R})$ belong to the same operational equivalence class.

This is weaker than "amplitudes are gauge redundancy" (which would require them to be physically meaningless). The correct statement is:

Amplitudes parametrize preparations, but only finitely many preparation classes are operationally distinguishable under finite resources.

The Bloch sphere is operationally meaningful—but as a continuum of preparation *procedures*, not as a continuum of *physically real states* surviving operational equivalence.

7.4 Finite Hilbert Spaces

For bounded-energy systems in finite regions:

$$\dim(\mathscr{H}_{\text{eff}}) < \infty$$

The system embeds in a finite qubit register:

$\mathscr{H}_{\text{eff}} \hookrightarrow (\mathbb{C}^2)^{\otimes n}$ where $n = \lceil \log_2 \dim(\mathscr{H}_{\text{eff}}) \rceil$

The infinite-dimensional Hilbert spaces of QFT are idealizations—useful for calculation but exceeding finite-region, finite-energy information capacity.

7.5 Contextual Structure in Quantum Systems

Within any maximal measurement context M on a quantum system:

- The outcome algebra B_M is finite Boolean
- Atoms of B_M are the N_{dist} perfectly distinguishable outcomes
- The **bit capacity** is $n = \lceil \log_2 N_{\text{dist}} \rceil$

For a qubit in context M = {σ_z measurement}:

- B_M has 2 atoms: $\{|0\rangle\langle 0|, |1\rangle\langle 1|\}$
- These atoms encode n = 1 bit of information capacity

- The atoms are the *outcomes*; the bit is the *capacity* they carry

Quantum mechanics is not a counterexample to the bit–tick ontology—it is its natural quantum realization, with atoms as measurement outcomes, bits as information capacity, and amplitudes as sub-resolution preparation parameters.

---

8. Eliminating Circularity: Ontological Hierarchy

8.1 The Correct Hierarchy

Level	Concept	Status
0	Ticks (worldline-local)	Primitive (irreducible event increments)
1	Atoms (context-relative)	Primitive (elementary distinguishable outcomes)
1′	Bits	Derived ($n = \lceil \log_2 N\_dist \rceil$ from atom count)
2	Energy, Entropy, Duration	Derived (from bit–tick ratios and counts)

Ticks are primitive along worldlines. They are the minimal event increments on clock records. They do not presuppose energy, forces, or duration—duration is defined *by* tick count.

Atoms are primitive within measurement contexts. They are the N_dist elementary distinguishable outcomes. **Bits** are the information capacity: $n = \lceil \log_2 N\_dist \rceil$ binary questions needed to specify which atom occurred.

Energy is derived from bits-per-tick:

$$\mathbf{E = (\hbar/\tau_0) \times \gamma}$$

This is a measurement definition, not a causal mechanism.

8.2 No Circularity

The hierarchy is strictly ordered: Level n depends only on levels < n.

- Ticks don't require bits (they're event increments, period)
- Bits don't require ticks (they're outcome atoms, period)
- Energy requires both (it's their ratio)

---

9. Mathematical Consistency Checks

9.1 Planck's Relation

Assume one bit per cycle of period T.

- Ticks per cycle: $N = T/\tau_0$
- Bits per tick: $\gamma = \tau_0/T$

From $E = (\hbar/\tau_0) \times \gamma$:

$E = \hbar/T = hf$ ✓

9.2 Boltzmann Entropy

Let $\mathscr{M}$ be a macrostate with $\Omega_{\mathscr{M}}$ compatible micro-histories. Then:

$S = k_B \ln \Omega_{\mathscr{M}}$ ✓

Entropy counts compatible tick-histories through configuration space.

9.3 Fine-Structure Constant

The standard QED expression:

$\alpha = e^2/(4\pi\varepsilon_0\hbar c) \approx 1/137$

can be rewritten in terms of vacuum impedance Z_0 and the von Klitzing constant $R_K = h/e^2$:

$\alpha = (Z_0/2) / R_K$

Bit-tick interpretation. In this framework, α admits an information-theoretic reading: it quantifies the efficiency with which electromagnetic ticks (field oscillations) convert into distinguishable bits (detection events). The impedance ratio measures conversion efficiency at the vacuum-matter interface.

Note. This paper does not derive α ≈ 1/137 from first principles—it provides an interpretation consistent with QED. Geometric derivation of constants is developed separately.

---

10. No-Go Theorems: The Forced Choice

10.1 No-Go for Non-Bit Distinctions

Theorem. Any ontology denying bits (binary capacity of maximal contexts) as primitive while satisfying Axioms 1–4 is inconsistent.

To deny bits, you must accept:	Consequence
Infinite perfectly distinguishable states in bounded regions	Violates Axiom 1 (Bekenstein bound)
Distinctions inaccessible to any experiment	Violates Axiom 3
Infinite distinguishable preparations under bounded I(prep:outcome)	Violates Axiom 4 (Holevo/channel capacity)

Proof. Any alternative must either:

- Have infinite N_dist (violates A1)
- Have finite N_dist with maximal contexts whose Boolean algebras have $n = \lceil \log_2 N\_dist \rceil$ bits of capacity (accepts bits)
- Claim infinite preparation classes are distinguishable under finite resources (violates A4) ■

10.2 No-Go for Non-Tick Time

Theorem. Any ontology denying ticks (worldline-local successor increments) while satisfying Axioms 1–4 is inconsistent.

To deny ticks, you must accept:	Consequence
Time parameter no clock realizes	Violates Axiom 2
Global absolute time	Contradicts GR; violates A2
Continuous proper time as fundamental	Operationally indistinguishable from ticks under $\sim_{\{\mathscr{R},\varepsilon\}}$
Duration without physical transitions	Operationally meaningless; violates A3

Proof. Along any worldline with clock events, the structure satisfying A2 and (T1)–(T3) is unique: $(\mathbb{N}, S)$. The tick is the atomic increment of this structure. Continuous proper time $\tau \in \mathbb{R}$ may exist as a mathematical idealization, but under finite-resource measurements, it collapses to tick-counts: any two proper times τ, τ' with $|\tau - \tau'|$ below resolution are operationally equivalent. The operational content of continuous time *is* the tick structure. ■

10.3 Forced Choice Table

If you want...	You must accept...	Which violates / collapses to...
Non-bit distinguishability	∞ perfect distinguishability OR inaccessible distinctions OR unbounded I(prep:outcome)	A1 OR A3 OR A4
Continuous state space as fundamental	Infinite distinguishable preparations under finite resources	A4; collapses to finite ε-net
Continuous time as fundamental	Infinite precision time measurements	Collapses to tick-count under $\sim_{\mathscr{R},\varepsilon}$
Non-tick time structure	Non-operational time OR global time	A2 OR contradicts GR

10.4 The Exhaustive Trilemma

Any alternative to BTS must fall into at least one of three failure modes:

Failure Mode 1: Surplus Structure

The framework posits distinctions not reflected in $E_{\mathscr{R},\varepsilon}$.

Example: Hidden variables with no experimental signature; unobservable degrees of freedom.

→ **Killed by Axiom 3** (no surplus structure beyond operational access).

Failure Mode 2: Infinite Finite-Resource Distinguishability

The framework claims unbounded distinguishable preparation labels under fixed $\mathscr{R}$.

Example: Asserting that all points on a continuous state space are physically distinct at finite resolution.

→ **Killed by Axiom 4** (bounded mutual information / finite ε-net).

Failure Mode 3: Non-Operational Time

The framework treats "time" as primitive without grounding in clock-event succession.

Example: Background time parameter with no physical realization; global simultaneity.

→ **Killed by Axiom 2** (operational time via clocks).

Meta-Theorem (Exhaustive Trilemma). Any ontology not equivalent to $BTS_{\mathscr{R},\varepsilon}$ must fall into at least one of:

1. Surplus structure (violates A3)
2. Infinite finite-resource distinguishability (violates A4)
3. Non-operational time (violates A2)

Hence BTS exhausts the operationally admissible class. ■

This is a **completeness result**: BTS is not merely *a* valid substrate—it is the *only* substrate compatible with operational grounding at finite resources.

10.5 Reduction Sketches: How Standard Frameworks Project onto BTS

Rather than saying "other frameworks fail," we show they *reduce* to BTS under operational equivalence. The message is:

Whatever you believe, once you impose ($\mathscr{R}$, ε)-equivalence, you land here.

Classical Phase Space

The projection π: (q, p) ↦ [q, p]_{$\mathscr{R}$,ε} bins phase space into resolution cells consistent with ($\mathscr{R}$, ε).

- Continuous coordinates (q, p) → finite grid of distinguishable cells
- Liouville measure → counting measure on cells
- Hamiltonian flow → discrete update rules on cell labels
- Surplus: sub-resolution phase space detail (collapses under π)

Quantum Field Theory

The projection π combines energy cutoff with detector coarse-graining:

- Infinite-dim Fock space → finite accessible mode content under energy bound
- Continuous field amplitudes → finite outcome alphabet after measurement
- S-matrix elements → context-change maps Φ_M
- Surplus: UV modes beyond energy cutoff; sub-resolution amplitude differences

String Theory / Moduli

The projection π maps continuous moduli to effective low-energy observables:

- Moduli space M → effective scattering/observable distributions
- Unresolved moduli collapse under ε-equivalence
- Compactification details → surplus (operationally equivalent configurations)
- What remains: finite set of distinguishable low-energy signatures

Loop Quantum Gravity

The projection π coarse-grains spin networks above resolution:

- Spin network states → equivalence classes under ($\mathscr{R}$, ε)
- Area/volume eigenvalues above Planck scale → finite distinguishable set
- Surplus: sub-Planckian structure; gauge-equivalent configurations

Causal Set Theory

The projection π extracts worldline record counts and finite event distinguishability:

- Causal set C → worldline tick-counts (already discrete)
- Event labels → equivalence classes under coarse-graining
- Closest to BTS of standard approaches; mainly adds causal ordering detail
- Surplus: specific causal relations below resolution

Summary Table

Framework	Surplus (collapses under π)	Core that survives
Classical mechanics	Sub-resolution (q, p)	Finite phase cells
QFT	UV modes; amplitude detail	Finite effective states
String theory	Moduli beyond resolution	Low-energy signatures
LQG	Sub-Planckian structure	Finite area/volume classes
Causal sets	Fine causal structure	Tick-counts + bit capacity

In each case: **the operational core is BTS**.

10.6 Precise Failure Claims for Continuous Ontologies

To prevent strawman objections, we state precisely what continuous frameworks must accept:

Claim (Continuous Fields). A continuous field ontology $\varphi(x)$ is permissible as a *representational* device. However:

1. If it asserts physically meaningful distinctions at scales below ($\mathscr{R}$, ε), it **violates Axiom 3** (surplus structure).
2. If it insists those distinctions are operationally extractable at fixed resources, it **violates Axiom 4** (infinite distinguishability).
3. If it remains agnostic about sub-resolution structure, it **reduces to BTS** under π.

Claim (Continuous Time). Continuous proper time $\tau \in \mathbb{R}$ is permissible as a mathematical idealization. However:

1. If it asserts duration has physical meaning independent of clock readings, it **violates Axiom 2** (non-operational time).

2. If it claims arbitrarily fine time resolution is achievable at finite resources, it **violates Axiom 4**.
3. If it accepts that measured time is clock-event counting, it **reduces to tick structure**.

The choice is always: reduce to BTS, or violate an axiom. There is no third option.

---

11. Discussion

11.1 The Role of Axiom 4

Axioms 1–3 are standard in operational approaches to physics. Axiom 4 (finite accessible information) provides the physical grounding for **true finiteness** of the operational core.

Logical status of Axiom 4. The bounds motivating A4 (Holevo, Bekenstein, channel capacity, quantum speed limits) are theorems within specific physical frameworks (quantum mechanics, GR). This creates a potential bootstrap concern: if bit-tick structure is supposed to be *prior* to these frameworks, can we use their theorems to justify A4?

The resolution: A4 is an **empirically motivated constraint**, not a theorem within bit-tick theory. We observe that all known physics respects finite-information bounds. We abstract this empirical regularity as A4. We then prove that A4 (plus A1–A3) forces bit-tick structure. The logic is:

Empirical evidence for finite-information bounds → Axiom 4 → Bit-tick substrate theorem

This is not circular. It's the standard scientific pattern: generalize from observations, then derive consequences. If future physics discovered infinite-information phenomena, A4 would need revision—but all current evidence supports it.

Without Axiom 4, one obtains:

- Bits as information capacity of maximal contexts ✓
- Ticks as worldline-local successor structure ✓
- But potentially infinite operationally distinct preparations

With Axiom 4:

- Finite ε-net of distinguishable preparation classes ✓
- True finiteness of operational core ✓
- Continuous parameters collapse under statistical equivalence ✓

11.2 What Uniqueness Means

The bit–tick substrate is unique in this sense:

1. **Structural uniqueness.** The operational core of any A1–A4 theory is finite, with contextual Boolean structure and worldline $(\mathbb{N}, S)$ structure.
2. **Isomorphism uniqueness.** Theories agreeing operationally have isomorphic bit–tick substrates.
3. **Elimination of alternatives.** Every alternative violates at least one axiom grounded in established physics or methodology.
4. **Universal property.** BTS is the terminal object among operational prediction carriers (Section 6.5).

11.3 Invariant Completeness: Why Bits and Ticks Exhaust the Primitives

Bits and ticks are the **only two operational invariants** that survive quotienting by $(\mathscr{R}, \varepsilon)$-equivalence:

Invariant 1: Outcome Structure

What can be told apart? → Maximal distinguishability capacity N_dist → $n = \lceil \log_2 N\_dist \rceil$ bits

Invariant 2: Temporal Structure

What can be counted as "later"? → Event-successor on clock records → $(\mathbb{N}, S)$ ticks

Proposition (Invariant Completeness). Any operational theory must provide:

1. A maximal test structure (determining what can be distinguished), and
2. An event-successor structure on records (determining what counts as temporal order).

There is no third independent primitive that survives quotienting by operational equivalence at finite resources.

Argument. Consider what else might be proposed:

- **Spatial structure?** Requires distinguishing locations → reduces to bits (which positions are distinguishable).
- **Causal structure?** Requires temporal ordering → reduces to ticks (which events precede which).
- **Field values?** Requires distinguishing configurations → reduces to bits (which configurations are distinguishable).
- **Particle identity?** Requires distinguishing entities → reduces to bits.
- **Continuous parameters?** Collapse under $(\mathscr{R}, \varepsilon)$-equivalence → reduce to finite bit classes.

Every proposed primitive either:

1. Reduces to bits (distinguishability structure), or
2. Reduces to ticks (temporal structure), or
3. Is surplus (operationally inaccessible).

Hence BTS = (bits, ticks, context-change maps) is the **complete basis** for operational physics.

11.4 Relation to Other Programs

Program	Relation
It from Bit (Wheeler)	Bits fundamental; we add contextual structure and ticks
Causal Set Theory	Order structure matches worldline ticks; we add contextual bits
QBism	Operational focus compatible; we add explicit resolution bounds
Constructor Theory	Counterfactual structure; compatible framing
Holographic Principle	Motivates Axiom 1 directly
Hardy/Chiribella reconstructions	Operational axioms for QM; we extend to include time structure

11.5 Implications

If the bit–tick ontology is correct:

- **Spacetime** emerges from worldline-local tick structures
- **Fields** are patterns of bit propagation across worldlines
- **Particles** are stable bit configurations in maximal contexts
- **Quantum amplitudes** are sub-resolution preparation parameters
- **Physical constants** encode tick-to-bit conversion efficiencies

11.6 Scope and Limitations

This paper establishes a **structural result**: the unique kinematic substrate of operationally grounded physics. It does not resolve several important questions:

What this paper does:

- Proves that A1–A4 force bit-tick structure
- Shows the axioms are independent and well-motivated
- Demonstrates compatibility with quantum mechanics
- Rules out alternative substrates without violating established bounds

What this paper does not do:

- Derive the Born rule internally from bit-tick axioms alone (though $P = |\psi|^2$ is derived in companion work [16] from geometric axioms compatible with the bit-tick substrate)
- Determine which specific dynamics governs bit-tick evolution
- Explain why physical constants have their observed values
- Derive spatial structure or gravity from bit-tick primitives

These limitations are not defects but scope boundaries. The paper contributes a **uniqueness theorem for kinematic structure**, analogous to results showing that spacetime must be Lorentzian or that quantum states must form a Hilbert space. Such structural theorems are valuable even when they leave dynamics and probability undetermined.

The open questions (Appendix D) define the research program beyond this work.

---

12. Conclusion

By requiring:

- Axiom 1: Finite perfect distinguishability
- Axiom 2: Operational time (worldline-local)
- Axiom 3: No surplus structure
- Axiom 4: Finite accessible information (supported by three independent routes: information-theoretic, decision-theoretic, thermodynamic)

we have proven:

1. **Bits** are the binary capacity of maximal measurement contexts ($n = \lceil \log_2 N\_dist \rceil$)
2. **Ticks** are the successor increments of worldline temporal structure ($\mathbb{N}$, S)
3. **The bit–tick substrate** $BTS\_\{\mathscr{R},\varepsilon\}$ is unique up to isomorphism, with a universal property (terminal object among operational prediction carriers)
4. **Invariant completeness**: bits and ticks exhaust the operational primitives—no third invariant survives ($\mathscr{R}$, ε)-quotienting
5. **Exhaustive trilemma**: every alternative must violate surplus structure, finite distinguishability, or operational time
6. **All standard frameworks** (QFT, string theory, LQG, causal sets) reduce to BTS under operational equivalence
7. **Quantum mechanics** fits naturally with amplitudes collapsing under statistical equivalence

The bit–tick framework is the **unique minimal invariant** of operationally grounded, resource-bounded physics. Whatever you believe about the ultimate nature of reality, once you impose operational equivalence at finite resources, you land here.

Appendix A: Formal Proofs

A.1 Stone Representation (Finite Boolean Algebras)

Theorem. Every finite Boolean algebra B is isomorphic to $\mathcal{P}$(Atoms(B)).

Proof. Define $\varphi: B \to \mathcal{P}(\text{Atoms}(B))$ by $\varphi(b) = \{a \in \text{Atoms}(B) : a \leq b\}$. This is a well-defined Boolean homomorphism. Injectivity: if $\varphi(b) = \varphi(b')$, then b and b′ sit above the same atoms; since each element is the join of atoms below it, $b = b'$. Surjectivity: for $S \subseteq \text{Atoms}(B)$, let $b = \bigvee S$; then $\varphi(b) = S$. ■

A.2 Uniqueness of (ℕ, S)

Theorem. Any countable, discrete, well-ordered set with minimum and no maximum is isomorphic to $(\mathbb{N}, \leq)$.

Proof. Define $f: \mathbb{N} \to T$ inductively: $f(0) = \min(T)$, $f(n+1)$ = successor of $f(n)$ in T. By discreteness, successors exist and are unique. By no-maximum, the sequence is unbounded. By well-ordering, f is surjective. By construction, f is order-preserving and injective. ■

A.3 Finite ε-Nets of Outcome Distributions

Theorem. Let the feasible coarse-grained outcomes under resource bound $\mathscr{R}$ form a finite alphabet of size $K = K(\mathscr{R})$. Then the set of feasible outcome distributions lies in the (K−1)-simplex Δ_{K-1}, which is compact. Under total variation distance (or any metric inducing the standard topology), Δ_{K-1} admits finite ε-nets. Hence the quotient under $d_{TV}(p,q) \leq \varepsilon$ has finite cardinality bounded by the ε-covering number $N(\varepsilon, K)$.

Proof.

1. Outcome distributions over K outcomes form the probability simplex:

   $$\Delta_{K-1} = \{p \in \mathbb{R}^K : p_i \geq 0, \Sigma_i\, p_i = 1\}$$

2. Δ_{K-1} is compact (closed and bounded subset of $\mathbb{R}^K$).
3. Total variation distance $d_{TV}(p,q) = \frac{1}{2} \Sigma_i |p_i - q_i|$ is a metric on Δ_{K-1}.
4. For any $\varepsilon > 0$, the ε-covering number $N(\varepsilon, \Delta_{K-1})$ is finite by compactness.
5. Explicit bound: $N(\varepsilon, \Delta_{K-1}) \leq (3/\varepsilon)^{K-1}$ (standard covering number estimate for simplices).
6. Define equivalence: $p \sim_\varepsilon q$ iff $d_{TV}(p,q) \leq \varepsilon$.
7. The quotient $\Delta_{K-1}/\sim_\varepsilon$ *has at most* $N(\varepsilon, \Delta_{K-1})$ classes.
8. Since preparations map to distributions via Φ_M, the quotient $S/\sim_{\mathscr{R},\varepsilon}$ is also finite. ■

Remark. This is the correct topology for the finiteness argument. The equivalence is on *distributions* (living in a compact simplex), not on preparation *parameters* (which may be non-compact). Compactness of Δ_{K-1} is automatic once outcomes are coarse-grained to finite K.

Appendix B: Contextual Boolean Structure in Quantum Mechanics

B.1 The Kochen-Specker Situation

The Kochen-Specker theorem shows that quantum observables cannot all be assigned simultaneous definite values consistently. This reflects the non-Boolean global structure of quantum propositions.

However, **within any single context** (compatible set of observables), the logic *is* Boolean. The bit–tick framework operates at this contextual level:

- Each maximal context defines a Boolean algebra of outcomes
- Bits are atoms of these contextual algebras
- No global Boolean algebra is claimed

B.2 Effect Algebras and Orthomodular Lattices

The global structure of quantum events is an **orthomodular lattice** (projection lattice of Hilbert space) or more generally an **effect algebra** (for POVMs).

These structures:

- Are non-Boolean globally
- Contain Boolean subalgebras (one per context)
- Have the contextual Boolean algebras as "local" structure

The bit–tick ontology operates at the level of maximal Boolean subalgebras: atoms are the elementary outcomes, and bits are their information capacity ($n = \lceil \log_2 N_{dist} \rceil$). This is fully compatible with orthomodular global structure.

Appendix C: Worldline Structure in Relativity

C.1 No Global ($\mathbb{N}$, S)

In Minkowski spacetime, events have a partial order (causal order) but no global total order. Different inertial observers disagree on simultaneity.

The tick structure is therefore **worldline-local**:

- Along any timelike worldline, events are totally ordered
- This total order has ($\mathbb{N}$, S) structure for discrete clocks
- Different worldlines accumulate different tick counts (time dilation)

C.2 Compatibility with General Relativity

In GR, proper time along a worldline is:

$$\boldsymbol{\tau = \int d\tau = \int \sqrt{(-g_{\mu\nu}\, dx^\mu\, dx^\nu)}}$$

For a discrete clock, this integral counts ticks:

$$\boldsymbol{\tau = N_{ticks} \times \tau_0}$$

The tick ontology is fully compatible with relativistic proper time—it identifies τ_0 as the fundamental tick duration and τ as the accumulated count.

Appendix D: Open Question

1. **Born rule derivation and scope.** The Born rule is derived uniquely within the broader VERSF–RAL framework, specifically in *Part II: The Double Square Rule* [16], where quantum probability is shown to follow inevitably from discrete informational geometry, reversible isometries, and irreversible selection acting on path-correlation structures. No probabilistic postulate, Hilbert space, or amplitude rule is assumed; the quadratic form $P = |\psi|^2$ is proven to be the unique solution compatible with positivity, normalization, relabeling invariance, factorization, and interference.

   The present paper focuses on establishing the bit–tick substrate as the unique kinematic invariant of operationally grounded physics. While fully compatible with—and indeed motivating—the Double Square Rule, this paper does not re-derive the Born rule internally from the bit–tick axioms alone. Instead, it provides the substrate on which the Double Square probability law acts.

Integrating the Double Square Rule directly into the bit–tick axiom set—thereby deriving quantum probability entirely from bit–tick primitives without additional geometric axioms—remains an important unification task for future work.

Appendix E: Methodological Status of Axiom 4 and Interpretive Remarks on Dimensionless Constants

E.1 The Status of Axiom 4 (Finite Accessible Information)

Axiom 4 asserts that, for any bounded experiment with finite physical resources, the mutual information extractable between preparation procedures and measurement outcomes is finite. In the main text, this axiom is motivated by several well-established bounds—such as the Holevo bound, the Bekenstein entropy bound, and quantum speed limits—which arise within specific physical frameworks.

A potential concern is that these bounds are themselves derived within quantum mechanics or general relativity, whereas the bit–tick substrate is intended to apply at a more foundational level. This appendix clarifies the logical status of Axiom 4 and resolves any appearance of circularity.

E.1.1 Axiom 4 as a Methodological Constraint

Axiom 4 need not be regarded as a derived physical law. Instead, it may be taken as a **methodological constraint on admissible physical theories**:

Physics should make only finite experimental claims.

That is, any theory that purports to describe physical reality must, at minimum, allow its predictions to be operationally decidable using finite resources. A theory that requires infinite precision, infinite outcome alphabets, or unbounded distinguishability under fixed experimental conditions fails to define empirically meaningful propositions.

Under this interpretation, Axiom 4 does **not** presuppose the mathematical structure of quantum mechanics, general relativity, or any other specific framework. Rather, it constrains the *class of theories* that qualify as operationally well-posed physical theories at all.

E.1.2 Role of Information-Theoretic and Thermodynamic Bounds

The Holevo bound, Bekenstein bound, channel-capacity limits, and quantum speed limits should therefore be understood not as foundations for Axiom 4, but as **consistency checks**:

- They demonstrate that all empirically successful existing theories respect the finite-accessible-information constraint.
- They provide concrete instantiations of Axiom 4 within established frameworks.
- They corroborate that Axiom 4 captures a robust empirical regularity rather than an ad hoc restriction.

The logical structure is therefore:

Methodological finiteness requirement
→ Axiom 4
→ Bit–tick substrate theorem
No step in the bit–tick substrate theorem relies on the numerical form of any specific bound.

E.1.3 Independence from Bit–Tick Conclusions

Crucially, the bit–tick substrate theorem does not depend on the *numerical form* of any particular bound. Any constraint—whatever its origin—that limits operationally accessible distinctions to a finite set under bounded resources suffices. If future physics were to revise or replace existing bounds while preserving operational finiteness, the conclusions of this paper would remain unchanged.

E.2 Interpretive Status of the Fine-Structure Constant

Section 9.3 of the main text briefly discusses the fine-structure constant α in relation to the bit–tick framework. This appendix clarifies the intent and limits of that discussion.

E.2.1 No Derivation or Prediction Claimed

This paper does **not** derive the numerical value of the fine-structure constant, nor does it constrain its magnitude. No novel prediction concerning α is asserted here. The bit–tick substrate theorem is entirely independent of the value of any dimensionless coupling constant.

E.2.2 Heuristic Interpretation

The fine-structure constant is given by
$\alpha = e^2 / (4\pi\, \varepsilon_0\, \hbar\, c) \approx 1 / 137$
It may also be written in terms of the vacuum impedance Z_0 and the von Klitzing constant $R_k = h / e^2$ as
$\alpha = (Z_0 / 2) / R_k$
This representation highlights α as a dimensionless ratio comparing vacuum response properties to discrete charge–action conversion scales.
Within the bit–tick ontology, this structure admits a **heuristic interpretation**: α may be viewed as characterizing the efficiency with which electromagnetic field oscillations (ticks) are converted into distinguishable detection events (bits) at the vacuum–matter interface. This reading is offered solely as conceptual alignment with the informational perspective developed in the paper.

E.2.3 Scope Limitation

No explanatory burden is placed on this interpretation. A quantitative account of coupling constants would require a full dynamical theory governing bit–tick interactions, renormalization, and field-mediated distinguishability—topics that lie beyond the scope of the present work and are addressed separately in the broader VERSF framework.

Accordingly, references to the fine-structure constant in this paper should be understood as **interpretive remarks**, not as extensions of the substrate theorem or its proofs.

E.3 Summary

- Axiom 4 may be taken as a **methodological finiteness requirement**, independent of any specific physical theory.
- Established information-theoretic and thermodynamic bounds function as **empirical consistency checks**, not logical foundations.
- Discussion of the fine-structure constant is explicitly **non-derivational and non-predictive**, and carries no logical weight in the core argument.

With these clarifications, the bit–tick substrate theorem stands as a purely structural result, free of hidden assumptions about particular dynamical laws or numerical constants.

Appendix F: Formal Rigor Addendum (Information Bounds, ε-Nets, and Universality)

F.1 Formal Setup: Resource-Bounded Experiments and Statistical Equivalence

Let Σ be the set of admissible preparation procedures, and let $\mathcal{E}(\mathscr{R})$ be the set of experiments feasible under resource bound $\mathscr{R}$ (energy, runtime, apparatus size, bandwidth, memory).

For any experiment $E \in \mathcal{E}(\mathscr{R})$, outcomes lie in a finite alphabet Ω_E with $|\Omega_E| = K_E < \infty$ (finite detector readout). Each preparation $\sigma \in \Sigma$ induces an outcome distribution:

$$P_E(\cdot \mid \sigma) \in \Delta_{K_E-1}$$

Define the $\mathscr{R},\varepsilon$ operational equivalence relation on preparations:

$$\sigma \sim_{\mathscr{R},\varepsilon} \sigma' \text{ iff for all } E \in \mathcal{E}(\mathscr{R}),\ d_{TV}(P_E(\cdot|\sigma), P_E(\cdot|\sigma')) \leq \varepsilon$$

where total variation distance is:

$$d_{TV}(p, q) = \tfrac{1}{2} \sum_{i=1}^{K} |p_i - q_i|$$

The $\mathscr{R},\varepsilon$ operational core is the quotient:

$$Core_{\mathscr{R},\varepsilon} := \Sigma / \sim_{\mathscr{R},\varepsilon}$$

The goal of this appendix is to give explicit, standard sufficient conditions under which $Core_{\mathscr{R},\varepsilon}$ is finite, and to sharpen the universality claim.

F.2 Finite Core from Finite Accessible Information (Axiom 4 → explicit bound)

Axiom 4 may be stated in a form directly usable for finiteness:

There exists a finite constant $I_{max}(\mathscr{R})$ such that for any random preparation label X supported on a finite subset of Σ, and any feasible experiment $E \in \mathcal{E}(\mathscr{R})$ producing outcome Y_E,

$$I(X : Y_E) \leq I_{max}(\mathscr{R})$$

Now fix $\varepsilon > 0$. Consider any family of preparations $\{\sigma_1, \ldots, \sigma_M\}$ such that they are pairwise distinguishable at level ε under $\mathscr{R}$. Formally, for each $i \neq j$, there exists some $E_{ij} \in \mathcal{E}(\mathscr{R})$ with:

$$d_{TV}(P_{E_{ij}}(\cdot|\sigma_i), P_{E_{ij}}(\cdot|\sigma_j)) > \varepsilon$$

Assume the experimenter chooses, for each i, a decoding procedure that guesses i from outcomes (standard hypothesis testing). Let the (best-achievable) average classification error be P_e.

A standard information-theoretic inequality (Fano's inequality) gives:

$$P_e \geq 1 - (I(X:Y) + \log_2 2) / \log_2 M$$

where X is uniform on $\{1,\ldots,M\}$ and Y is the measurement outcome (for the chosen optimal discrimination strategy). Rearranging:

$$\log_2 M \leq I(X:Y) + 1 / (1 - P_e)$$

In particular, if the class is ε-distinguishable, there exists a measurement strategy whose error is bounded away from 1 (and in the ideal limit, one can take $P_e < 1/2$). Taking the conservative choice $P_e \leq 1/2$ yields:

$$\log_2 M \leq 2(I(X:Y) + 1)$$

Using Axiom 4, $I(X:Y) \leq I_{max}(\mathscr{R})$, so:

$$M \leq 2^{2(I_{max}(\mathscr{R}) + 1)}$$

Conclusion (Explicit Packing Bound). Under Axiom 4, the number of mutually ε-distinguishable preparation classes under resource bound $\mathscr{R}$ is finite, with an explicit bound:

$|Core_{\mathscr{R},\varepsilon}| \leq 2^{2(I_{max}(\mathscr{R}) + 1)}$

This makes "finite accessible information ⇒ finite operational core" a standard corollary of Fano's inequality.

Remark. This bound is not claimed tight. Its purpose is to show that finiteness follows from a widely accepted inequality once $I_{max}(\mathscr{R})$ is assumed finite.

F.3 Geometric ε-Net Bound on Outcome Distributions (compactness → explicit covering number)

For any fixed feasible experiment E with K_E outcomes, the outcome simplex Δ_{K_E-1} is compact. The covering number under total variation distance satisfies a standard bound:

$N(\varepsilon, \Delta_{K-1}, d_{TV}) \leq (3/\varepsilon)^{K-1}$

Thus, for fixed E, the set $\{P_E(\cdot|\sigma) : \sigma \in \Sigma\}$ admits an ε-net of size at most $(3/\varepsilon)^{K_E-1}$.

If $\mathcal{E}(\mathscr{R})$ is finite (as it is under any finite-description experimental catalogue), then taking the product metric over experiments yields a finite joint ε-net, hence a finite quotient $Core_{\mathscr{R},\varepsilon}$. Even if $\mathcal{E}(\mathscr{R})$ is infinite, any finite experimental campaign uses only finitely many E, and the induced operational partition is finite.

This gives a second, purely geometric route to finiteness consistent with the main text's Appendix A.3.

F.4 Successor Structure Theorem (Worldline-Local) with Minimal Assumptions

Let W be a clock record: a sequence of recorded events on a timelike worldline. Model W as a set with a strict total order "<".

Assumptions:

(W1) Discreteness: For every event $e \in W$ except the first, there exists an immediate predecessor pred(e), and for every event except the last (if any), there exists an immediate successor succ(e).

(W2) Well-founded past: There exists a first event $e_0 \in W$ with no predecessor.

(W3) No terminal event: For every $e \in W$, succ(e) exists (no last tick).

Then the map f: $\mathbb{N} \to W$ defined inductively by:

$f(0) = e_0$
$f(n+1) = succ(f(n))$

is an order-isomorphism between $(\mathbb{N}, <)$ and $(W, <)$. Hence the tick structure is isomorphic to $(\mathbb{N}, S)$, where $S(n) = n+1$.

Remark. This avoids requiring global spacetime well-ordering; it only uses the discrete successor property on the worldline record, which is exactly the operational content of a clock.

F.5 Formal Universality (Terminal Object) in the Category of Operational Prediction Carriers

Define a category $\mathcal{C}$_{$\mathscr{R}$,ε} as follows.

Objects: Triples A = (Core_A, {B_M}_A, {Φ_M}_A) where:

• Core_A is a finite set (operational classes)
• {B_M}*A is a family of finite Boolean algebras (maximal contexts)*
• For each context M, Φ_M maps Core_A → Δ{K−1} giving coarse-grained outcome distributions

Morphisms: A morphism f: A → A′ is a pair of maps:

f_Core: Core_A → Core_{A′}
f_B: {B_M}*A → {B_M}*{A′}

such that for every context M, the following compatibility holds:

$\Phi'_M \circ f_{Core} = \Phi_M$

(Interpretation: f preserves all operational predictions.)

Define BTS_{$\mathscr{R}$,ε} to be the object built from:

• Core_{$\mathscr{R}$,ε} = Σ / ∼_{$\mathscr{R}$,ε}
• contextual Boolean algebras for maximal contexts
• the induced prediction maps Φ_M
• plus worldline-local tick structure $(\mathbb{N}, S)$ as an attached component (formally, a functorial "time-record" assignment per worldline)

Claim (Terminal Object Property). For any admissible ontology object A in $\mathcal{C}\_\{\mathcal{R},\varepsilon\}$, there exists a unique morphism:

$$A \to \mathrm{BTS}\_\{\mathcal{R},\varepsilon\}$$

Proof sketch. The quotient map $\pi: \Sigma \to \Sigma/\sim\{\mathcal{R},\varepsilon\}$ *is canonical. Any object A that carries only operational content factors through* π *by construction of* $\sim\{\mathcal{R},\varepsilon\}$. Uniqueness follows because any morphism must agree on equivalence classes, and the compatibility condition forces agreement on Φ_M. Therefore $\mathrm{BTS}\_\{\mathcal{R},\varepsilon\}$ receives a unique prediction-preserving map from any A. □

This makes the "universal receiver" language in the main text fully formal.

F.6 What This Appendix Adds

This appendix supplies:

1. An explicit information-theoretic bound showing $\mathrm{Core}\_\{\mathcal{R},\varepsilon\}$ is finite (Fano route).
2. An explicit geometric covering number bound (ε-net route).
3. A tightened successor theorem with minimal worldline assumptions.
4. A fully formal category definition and a terminal-object proof sketch.

None of these modify the main theorem; they make its proof obligations easier for a referee to verify line-by-line.

Appendix G: Structural Strengthening of Axiom 4

G.1 Motivation

Axiom 4 in the main text is stated as a bound on accessible mutual information under finite resources. While convenient, this formulation can itself be derived from more primitive constraints on experimental practice and physical realisability. This appendix presents two such constraints and shows that, taken together, they imply Axiom 4 as stated.

G.2 Finite Experiment Structure (Operational Constraints)

For fixed experimental resources $\mathcal{R}$, assume:

A4a (Finite Outcome Alphabet)

For any feasible experiment E under $\mathscr{R}$, the detector readout alphabet is finite:

$|\Omega_E| = K_E(\mathscr{R}) < \infty$.

A4b (Finite Sampling Budget)

Under resources $\mathscr{R}$, any experiment can generate at most $N(\mathscr{R})$ outcome samples.

A4c (Finite Statistical Resolution)

There exists $\varepsilon(\mathscr{R}) > 0$ such that, for any two preparations σ and σ′, if for all feasible experiments E,

$d_{TV}(P_E(\cdot|\sigma), P_E(\cdot|\sigma')) \leq \varepsilon(\mathscr{R})$,

then σ and σ′ cannot be reliably distinguished under $\mathscr{R}$.

Here the total variation distance is

$d_{TV}(p, q) = \frac{1}{2} \sum_{i=1}^{K} |p_i - q_i|$.

These three conditions define what it means for an experiment to be finite in readout, duration, and resolution.

G.3 Finite Specification and Control (Physical Realisability Constraint)

A4d (Finite Specification / Finite Control)

For fixed resources $\mathscr{R}$, any preparation procedure or measurement apparatus can be specified and controlled only up to finite description length $L(\mathscr{R})$.

Equivalently, the number of physically realisable preparation protocols and experimental configurations under $\mathscr{R}$ is finite, bounded by

$N_{spec}(\mathscr{R}) \leq 2^{L(\mathscr{R})}$.

This reflects the fact that experimental devices must be built, programmed, calibrated, and stabilised using finite physical resources, and therefore cannot encode or enact infinitely many distinct procedures.

G.4 Derivation of Finite Accessible Information

We now show that A4a–A4d imply the finiteness condition used in the main text.

From A4a, each experimental outcome carries at most $\log_2 K\_E$ bits of raw information per trial.

From A4b, only $N(\mathscr{R})$ trials are available.

From A4c, distinctions below $\varepsilon(\mathscr{R})$ are operationally invisible.

Therefore, for any preparation label X and any feasible experiment E producing outcome Y_E,

$I(X : Y\_E) \leq N(\mathscr{R}) \cdot \log_2 K\_E(\mathscr{R}) < \infty$.

This yields a finite upper bound $I\_max(\mathscr{R})$ on accessible mutual information, recovering Axiom 4 as stated in Section 2.

Independently, from A4d, the number of physically realisable preparation labels is finite. Even before statistical coarse-graining, only finitely many operationally distinct procedures exist under $\mathscr{R}$. Hence the operational partition of preparations must be finite.

Thus both routes—finite experiment structure and finite physical specification—converge on the same conclusion: under fixed $\mathscr{R}$, only finitely many distinctions are operationally accessible.

G.5 Equivalence to Axiom 4

The original formulation of Axiom 4 in the main text may therefore be regarded as a compact summary of the stronger structural constraints A4a–A4d.

For the purposes of the core universality theorem, either formulation suffices. The strengthened version clarifies that finiteness arises not from any specific physical theory, but from the combined facts that experiments have finite readout, finite duration, finite resolution, and finite physical control.

G.6 Interpretation

Axiom 4 is thus not a single assumption but a convergence point:

- Operational finiteness (what experiments can resolve), and
- Physical finiteness (what experiments can be built and controlled to do)

both independently force the collapse of continuous structure into a finite operational core. The bit–tick substrate theorem follows from this convergence, not from any particular information bound.

References

1. Wheeler, J.A. (1990). "Information, Physics, Quantum: The Search for Links."
2. Bekenstein, J.D. (1973). "Black Holes and Entropy."
3. Bekenstein, J.D. (1981). "Universal Upper Bound on the Entropy-to-Energy Ratio."
4. 't Hooft, G. (1993). "Dimensional Reduction in Quantum Gravity."
5. Holevo, A.S. (1973). "Bounds for the Quantity of Information Transmitted by a Quantum Communication Channel."
6. Shannon, C.E. (1948). "A Mathematical Theory of Communication."
7. Margolus, N. & Levitin, L.B. (1998). "The Maximum Speed of Dynamical Evolution."
8. Fuchs, C.A. (2010). "QBism, the Perimeter of Quantum Bayesianism."
9. Sorkin, R.D. (2005). "Causal Sets: Discrete Gravity."
10. Deutsch, D. (2013). "Constructor Theory."
11. Kochen, S. & Specker, E.P. (1967). "The Problem of Hidden Variables in Quantum Mechanics."
12. Birkhoff, G. & von Neumann, J. (1936). "The Logic of Quantum Mechanics."
13. Ludwig, G. (1983). *Foundations of Quantum Mechanics.*
14. Chiribella, G. et al. (2011). "Informational Derivation of Quantum Theory."
15. Hardy, L. (2001). "Quantum Theory From Five Reasonable Axioms." arXiv:quant-ph/0101012.
