Why the Foundational Conditions Represented by VERSF Are Unavoidable: A Structural Justification

Keith Taylor
VERSF Theoretical Physics Program

For the General Reader

The Void Energy-Regulated Space Framework (VERSF) is built on three foundational axioms: that physical distinctions are finite in granularity, that forming a physical fact commits irreversible correlation, and that the substrate of the universe has a finite capacity for localizing facts. These axioms might look like assumptions — choices made to get a particular theory off the ground.

This paper argues they are not choices. They are the only conditions under which physics of any kind is possible at all.

We show, in strict logical order, that any universe capable of supporting physical laws — any universe where something counts as a measurement, a record, or a reproducible result — must satisfy all three conditions. A universe that violates any one of them does not merely lack our physics. It lacks the structure needed for any physical law to hold.

The VERSF axioms are not the starting point of a particular framework. They are the unavoidable prerequisites of physics itself, and VERSF is the framework built on recognizing that.

Abstract

The Void Energy-Regulated Space Framework rests on three foundational axioms: (A1) finite distinguishability — physical distinctions have a minimum resolvable granularity; (A2) irreversible commitment — every act of forming a physical fact exports correlation irreversibly beyond local causal control; (A3) finite localization capacity — there exists a maximum density at which facts can be formed per unit region. These axioms are sometimes presented as motivated postulates of the VERSF framework. This paper proves they are not postulates. Each one is a structural necessity: any universe that supports stable, reproducible physical facts — the minimal precondition for any physical law to exist — must satisfy all three. We prove this through a chain of four theorems. Theorem 1 establishes that physical laws require fact-stable

state classes; this is the single primitive from which everything follows. Theorem 2 shows that fact-stable state classes require finite distinguishability, using the information-cost argument that resolving a region to scale ℓ requires $I(\ell) \sim \log(V/\ell^3) \to \infty$ as $\ell \to 0$, which no finite physical record can sustain. Theorem 3 shows that fact formation requires irreversible commitment, grounded not in thermodynamics but in relativistic causal locality: no finite agent controls degrees of freedom outside its causal past, so local CPTP operations that reduce accessible correlations necessarily produce a monotone Entanglement Ledger that cannot be reversed by any locally-admissible operation. Theorem 4 shows that Theorems 2 and 3 jointly force a finite localization capacity — this is the closure condition that prevents the first two necessities from being mutually self-undermining. None of these are VERSF-specific claims. Each is stated in substrate-neutral language and shown to hold for any universe capable of supporting law-governed physical facts accessible to finite observers. The VERSF framework is the natural theoretical structure that emerges when these three necessities are taken seriously as the foundation of physics.

Result Summary.

	Theorem	Claim
T1	Lawhood requires fact-stable state classes	Physics possible ⟹ stable, recoverable, reproducible facts
T2	Lawhood + realizability ⟹ finite distinguishability	Any fact-supporting state space has a minimum resolvable granularity $\delta_{min} > 0$
T3	Finite observers + causal boundedness ⟹ irreversible commitment	Every local fact-forming operation in a local probabilistic theory deposits an irreversible correlation residue
T4	$A1 \wedge A2$ ⟹ finite localization capacity	Theorems 2 and 3 jointly force a finite upper bound on stabilizable distinction density

T2 and T3 are logically independent consequences of lawhood: finite distinguishability follows from operational realizability of facts, while irreversible commitment follows from causal boundedness of finite observers. Neither theorem presupposes the other. T4 is the closure condition forced when both are jointly imposed.

Table of Contents

---

1. Introduction: The Axioms and the Question {#1}

The Void Energy-Regulated Space Framework is built on three foundational axioms:

A1 — Finite Distinguishability: Physical distinctions have a minimum resolvable granularity. There exists a length scale $\delta_{min} > 0$ below which two physical states are operationally identical. No physical fact can record a distinction finer than δ_{min}.

A2 — Irreversible Commitment: Every act of forming a physical fact commits irreversible correlation structure to the environment. This commitment cannot be undone by any locally-bounded agent. The Entanglement Ledger — the total correlation exported by fact-formation events — is monotonically non-decreasing.

A3 — Finite Localization Capacity: The substrate has a finite maximum response to fact localization. There exists a minimum scale $\ell^* > 0$ below which fact formation is impossible. In VERSF this is represented constitutively as a maximum substrate tension $\mathcal{T}_{max}$.

When presented this way, these axioms can look like the arbitrary starting assumptions of one particular framework. The natural question is: *why these axioms?* Why not other axioms? What compels us to accept them?

This paper answers that question. We prove that A1, A2, and A3 are not choices. They are unavoidable: any universe capable of supporting law-governed physical facts accessible to finite observers must satisfy all three. Denying any one of them does not produce an alternative physics — it produces a universe in which the concept of a physical law is undefined.

The argument proceeds through four theorems in strict logical order. The first establishes the single primitive — physics requires stable, reproducible physical facts — and proves it is not a philosophical preference but a formal precondition for any law to exist. The second, third, and fourth theorems derive A1, A2, and A3 from this primitive in sequence, each one forced by what precedes it.

The paper makes no use of VERSF-specific machinery in the proofs. The theorems are stated in substrate-neutral language. VERSF enters only at the end, as the natural theoretical framework that takes these structural necessities seriously as its foundation.

What this paper does and does not claim. The paper does not derive detailed microphysics from lawhood alone. It derives only those structural conditions without which lawhood cannot possess an operational domain for finite observers. The conclusion is not that all physics follows from a single primitive, but that A1, A2, and A3 are the minimal structural prerequisites for any physics with accessible facts. A universe that satisfies these three conditions may still exhibit many different dynamical structures; what it cannot do is support law-governed, recoverable physical facts while violating any one of them.

Scope. This paper does not aim to complete the full quantitative VERSF program. Its aim is narrower and more foundational: to show that the axioms on which that program rests are unavoidable in any universe where physical law is possible. The derivation of specific constants — α, Λ, and the numerical value of $\mathcal{T}$_max — is downstream work; it cannot begin until the necessity of the axioms is established. That is what this paper establishes.

Domain of validity. This paper concerns universes capable of supporting: finite observers who register outcomes; accessible physical facts that can be recorded and recovered; and operationally meaningful laws that hold across instances. Universes that satisfy none of these conditions are not the intended subject — they are mathematical structures without physics in the relevant sense. This specification is not a concession; it is the exact class of universes for which the theorems are intended, stated explicitly to prevent irrelevant objections from muddying the argument.

---

2. The Single Primitive: Physics Requires Physical Facts {#2}

Definition (Fact-supporting universe). A universe is *fact-supporting* if it contains finite observers capable of registering outcomes of physical processes and comparing those outcomes across instances.

Premise (Operational Physics Premise). A universe admits physics only if it is fact-supporting. A universe that is not fact-supporting may admit mathematical dynamics but cannot admit operational physical laws.

This is a one-directional conditional: it says that physics requires the fact-supporting condition, not that the fact-supporting condition is sufficient for physics. The distinction between mathematical dynamics and operational physical laws is essential: a formal system of equations can evolve states without any agent capable of registering, comparing, or recovering outcomes; such a system is mathematics, not physics in the operational sense. The premise makes the domain of the theorems explicit. A universe that contains no finite observers, supports no registrable outcomes, and permits no cross-instance comparison may admit mathematical dynamics, but it does not admit physics in the operationally relevant sense. By making this premise explicit, the paper preempts the objection that "your conclusions depend on assuming observers" — the response is: yes, because physics itself requires them. The domain of Theorems 1–4 is precisely the class of universes satisfying the Operational Physics Premise. Theorems 1–4 are therefore statements about any universe that admits physics in this sense.

The theorems below can be summarized as: *any fact-supporting universe must satisfy A1, A2, and A3.*

2.1 Definition of a Physical Fact

Definition. A *physical fact* is a definite outcome of a physical process — an event that has occurred and can in principle be registered.

This is deliberately minimal. It does not yet require that the outcome be stable, reproducible, or distinguishable at any particular resolution. Those properties are not part of the definition of a physical fact; they are what Theorem 1 proves facts must have in order to serve as the domain of physical laws.

2.2 Theorem 1: Lawhood Requires Fact-Stable State Classes

Theorem 1. *A physical law presupposes a state space whose elements are fact-stable distinguishable classes. For the outcomes of physical processes to serve as the domain of a law, they must be stable, recoverable, and reproducible. Without these properties, no state space exists on which any law can act, and the concept of a physical law is formally undefined.*

Proof.

A physical law is formally a mapping L: $\mathcal{S} \times \mathcal{T} \to \mathcal{S}$, where $\mathcal{S}$ is a state space and $\mathcal{T}$ is a time parameter. The outcomes of physical processes (physical facts in the minimal sense of §2.1) must serve as the elements of $\mathcal{S}$. For this to be possible, lawhood imposes three conditions that the definition of §2.1 does not itself contain.

Step 1: The domain $\mathcal{S}$ must be a quotient space of robust equivalence classes — facts must be stable.

Physical states are not mathematical points but equivalence classes stable under perturbation. Given a metric d on the space of configurations, define:

$[s] = \{ s' : d(s, s') < \delta \}$

for some $\delta > 0$. A physical state is the equivalence class [s], not the point s. The physical state space $\mathcal{S}$ is the *quotient space* of configurations under this equivalence relation — the space $\mathcal{C}/\sim$ where $s \sim s'$ iff $d(s, s') < \delta$. Physical law does not act on microscopic configuration points; it acts on equivalence classes under admissible perturbation. Without a quotient space of robust classes — that is, if $\delta \to 0$ — the domain $\mathcal{S}$ collapses: equivalence classes shrink to singletons, the quotient is trivial, and any perturbation maps [s] to a new singleton [s′]. The domain is then not stable under the natural perturbations that any physical implementation of L must tolerate, and L fails to be a well-defined dynamical system on $\mathcal{S}$. A law that is not well-defined on its own domain is not a law. **Lawhood therefore requires facts to be stable.**

Step 2: The mapping L must be reproducible across instances — facts must be recoverable.

A law is a regularity across instances, not a single occurrence. The claim "state s_1 at t_1 maps to state s_2 at t_2" is meaningful only if $[s_1]$ and $[s_2]$ can be identified as the *same* equivalence classes across different instantiations. This requires that the distinctions defining $[s_1]$ and $[s_2]$ persist long enough to be registered by independent agents at different times. **Lawhood therefore requires facts to be recoverable.**

Step 3: The image of L must be confirmable — facts must be reproducible.

A law makes predictions. A prediction is confirmable only if the predicted outcome can be identified as a physical fact and compared with subsequent outcomes of the same preparation. If a fact cannot be reproduced under the same conditions, no comparison is possible, and the law has no empirical content. It is not physics — it is an unverifiable formal expression. **Lawhood therefore requires facts to be reproducible. □**

Corollary. The requirement that any physical law exist entails that facts — in the minimal sense of §2.1 — must also be stable, recoverable, and reproducible. These are not part of the definition of a fact; they are what a fact must be for physics to be possible. Theorem 2 will show that this in turn requires finite distinguishability.

Remark (Quotient structure as mathematical precondition). For physical law $L: \mathcal{S} \times \mathcal{T} \to \mathcal{S}$ to possess a stable operational domain, the equivalence classes forming $\mathcal{S}$ must be: (i) *nontrivial* — each class contains more than one point, so that the class persists under perturbation; (ii) *perturbatively robust* — membership in [s] is stable under the perturbations that any physical implementation of L must tolerate; (iii) *recoverably identifiable across instances* — an agent can determine whether two separately prepared states belong to the same class. The quotient structure is not interpretive decoration; it is the minimum mathematical form required for a physical law to possess a stable operational domain.

2.3 The Derivation Chain

From Theorem 1, we derive A1, A2, and A3. Theorems 2 and 3 are independent first-order consequences of Theorem 1 — neither presupposes the other. Theorem 4 is the closure condition their conjunction forces:

Physics possible
⟹ Stable, reproducible physical facts *(Theorem 1)*
⟹ { **A1: Finite distinguishability** *(Theorem 2)* } and { **A2: Irreversible commitment** *(Theorem 3)* } [independently, both from Theorem 1]
⟹ **A3: Finite localization capacity** *(Theorem 4 — forced by A1 ∧ A2)*

This matches the branching structure of the diagram in §6. A1 and A2 are derived in parallel from the same primitive; A3 is the closure condition that makes them simultaneously satisfiable.

Since lawhood requires a stable operational state space, and stability requires perturbation-robust equivalence classes, the remaining question is whether such classes can exist without a finite lower bound on fact-supporting distinguishability. Theorem 2 shows they cannot.

---

3. Theorem 2: Finite Distinguishability Is Unavoidable {#3}

3.1 Statement

Theorem 2 (VERSF Axiom A1 Is Unavoidable). *Any universe supporting fact-stable state classes must possess finite fact-supporting distinguishability: there exists $\delta_{min} > 0$ such that two states separated by less than δ_{min} are operationally identical as physical facts.*

3.2 Three Kinds of Distinguishability

Before the proof, we eliminate the most common misreading of this theorem by distinguishing three notions of distinguishability:

- **Logical distinguishability**: two states are distinct as formal objects in a mathematical description. Classical field theory has infinite logical distinguishability — field values range over $\mathbb{R}$. This is irrelevant to the theorem.
- **Physical distinguishability**: two states can in principle be separated by some physical operation. This is a dynamical property, not a fact-formation property.
- **Fact-supporting distinguishability**: two states can be stably separated and the separation recorded as a recoverable physical fact. This is *record-supporting separability* — the operationally relevant notion.

The theorem concerns only fact-supporting distinguishability. The existence of a mathematical continuum does not imply that arbitrarily fine distinctions can be stabilized as facts. A classical field assigns continuous labels to states, but those labels are extracted by measurements of finite resolution. The theorem is immune to the objection that "continua exist in mathematics" because that objection confuses logical with fact-supporting distinguishability.

Remark (Logical continuum vs operational continuum). A theory may possess a mathematically continuous configuration space while still admitting only finitely many distinguishable states operationally within any bounded region. The theorems of this paper concern operational distinguishability, not mathematical cardinality. A formally continuous state space is consistent with A1 provided the continuous labels cannot all be operationally resolved — which is exactly the condition $\delta_{min} > 0$ expresses.

3.3 The Information-Cost Argument

The following sequence of lemmas and propositions builds the information-cost argument in increasing generality: from the physical constraints on any fact-forming operation, to the requirements lawhood places on the state space, to the general encoding burden on any refinement family, and finally to the finite distinguishability floor that all of these jointly entail.

Lemma (Physical Realizability). *Any fact-forming operation occupies a finite spacetime region and therefore has finite causal access and finite available action.*

Proof. A fact-forming operation is a physical process — it has a spacetime location and a duration. By relativistic causality, signals propagate at or below c; the causal past of any event at spacetime point P with duration τ has volume $\sim (c\tau)^3$. This is finite for any finite τ. The action available to a process in a bounded spacetime region is bounded by $E \cdot \tau$ where E is the finite energy available within causal reach. Neither bound can be infinite for any physically realizable process. □

This lemma establishes, independently of quantum mechanics or any specific dynamics, that any physically realizable record-formation process operates under finite causal and operational constraints. It is not a physical intuition but a consequence of relativistic causal structure.

Proposition (Lawhood requires finite realizability). *If a fact cannot be produced, recorded, or recovered by any finite physically realizable process, it cannot belong to the state space of a physical law for finite observers.*

Proof. The domain of a physical law must be operationally instantiable: the state space $\mathcal{S}$ consists of equivalence classes [s] whose membership can be determined by physical agents. An element of $\mathcal{S}$ that cannot be instantiated — produced by any finite process, recorded in any finite time, or recovered by any finite agent — is not an element of the operational state space; it is a formal label with no physical referent. Non-instantiable distinctions are therefore not part of the physical state space. Any refinement that requires unbounded encoding burden to produce, record, or recover is excluded from lawhood itself — not by resource limitation, but by the structural requirement that the law's domain be operationally accessible. □

Lemma (Operational closure of state space). *If a distinction cannot be produced, recorded, and recovered by any finite physically realizable process, it cannot belong to the operational state space of a physical law. The operational state space of physics therefore contains only realizable distinctions.*

Proof. Immediate from the Lawhood-Realizability Proposition: non-instantiable distinctions are excluded from the physical state space. The operational state space is closed under realizability

— it contains exactly those distinctions that finite observers can produce, record, and recover. A formally infinite mathematical state space may contain additional elements, but those elements have no physical referents and play no role in any law whose domain must be operationally accessible. □

Lemma (Operational completeness). *A physical state space must be closed under the operations required to instantiate, record, and recover its elements. Distinctions that cannot participate in these operations cannot belong to the operational state space.*

Proof. Suppose a distinction d is in the operational state space $\mathcal{S}$ but cannot be instantiated, recorded, or recovered by any physically realizable process. Then d has no operational role: no physical law can refer to d as a distinct element of its domain, because no physical agent can access d to determine whether a given state is d or not. An element of $\mathcal{S}$ with no operational role is not an element of the operational state space — it is a formal label. The operational state space is therefore closed under instantiation, recording, and recovery in the sense that any element that cannot participate in these operations is not operationally distinguishable from its neighbors and drops out of the physical state space. The step from a formal (mathematical) state space to the operational state space is therefore not a restriction but a structural necessity: the operational state space is exactly the closure of the physically accessible distinctions under these three operations. □

The encoding-burden argument.

Given the Physical Realizability Lemma and the Lawhood-Realizability Proposition, forming a fact that distinguishes a state requires encoding the distinction in a record producible within finite causal reach and finite action, and this record must belong to the operational state space of any law that includes the distinction. The following lemma states the general form of the resulting constraint.

Lemma (General Encoding Burden). *For any fact-forming operation with N distinguishable outcomes, the minimum encoding burden is $\log_2 N$ bits. If $N \to \infty$ under refinement of the operation, fact formation requires unbounded encoding resources, which cannot be met within finite causal reach and finite action.*

Proof. A record of which of N outcomes occurred must distinguish N possibilities; by Shannon's source-coding theorem, at least $\log_2 N$ bits are required. If $N \to \infty$, the encoding burden $\log_2 N \to \infty$, and no process bounded by finite causal reach and finite action can sustain it. □

Lemma (Unbounded refinement implies unbounded encoding burden). *For any family of fact-forming operations whose number of distinguishable outcomes $N(\lambda)$ grows without bound under refinement $\lambda \to 0$, the minimum encoding burden $\log_2 N(\lambda)$ diverges. No physically realizable process can meet this diverging burden.*

Proof. Immediate from the General Encoding Burden Lemma: if $N(\lambda) \to \infty$ as $\lambda \to 0$, then $\log_2 N(\lambda) \to \infty$. By the Physical Realizability Lemma, no process with finite causal reach and finite action can produce or record an unbounded encoding burden. Any refinement family with $N(\lambda) \to \infty$ therefore exceeds the capacity of all physically realizable processes in the limit. □

Proposition (Finite operational state space). *Any state space that serves as the domain of a physical law for finite observers must contain only finitely many distinguishable states within any bounded region of configuration space.*

Proof. By the Physical Realizability Lemma, any fact-forming operation within a bounded region has finite causal reach and finite available action. By the Lawhood-Realizability Proposition, only operationally instantiable distinctions belong to the physical state space. By the General Encoding Burden Lemma, distinguishing among N outcomes requires $\log_2 N$ bits; if N is unbounded within a bounded region, the encoding burden diverges and exceeds the capacity of any physically realizable process within that region. Therefore the number of distinguishable states in any bounded region is finite — there exists a maximum $N_{max} < \infty$ such that no fact-forming operation within the region can exceed N_{max} distinguishable outcomes. The minimum resolvable granularity $\delta_{min} > 0$ is the direct consequence: it is the spacing required to keep N_{max} finite. □

This proposition is the synthesis the preceding lemmas and propositions jointly justify: starting from physical realizability and lawhood requirements, and using the encoding-burden machinery, the existence of $\delta_{min} > 0$ follows as an immediate corollary. It is what Theorem 2 formally establishes for the full state space.

For spatial localization as a concrete case, a fact distinguishing a state's location within a region of volume V to resolution ℓ has $N \sim V/\ell^3$ possible outcomes, giving encoding burden:

$I(\ell) \sim \log_2(V/\ell^3)$ bits [spatial localization example]

As $\ell \to 0$, $I(\ell) \to \infty$. More generally, the relevant quantity is the logarithm of the number of distinguishable outcomes of the fact-forming operation; spatial localization is a representative case, not the only one.

The closure step does not require "a finite physical system has finite Ω" as a premise. It requires only that any physically realizable record-formation process operates within finite causal duration and finite operational resources. An encoding burden that diverges as $\ell \to 0$ cannot be met by any process constrained to finite causal reach and finite operational resources. Record formation at arbitrarily fine resolution is therefore not physically realizable — not because we assume finite systems, but because physical realizability itself implies finite causal and operational bounds.

The full implication chain is:

Infinite distinguishability
→ unbounded encoding burden per fact
→ requires unbounded causal reach or unbounded operational resources
→ violates physical realizability of fact formation

→ no fact can be formed at arbitrarily fine resolution
→ no fact-stable state classes
→ no laws (Theorem 1)

3.4 Proof

Proof via information cost and three contradiction branches.

Suppose fact-supporting distinguishability is infinite: for every $\varepsilon > 0$, two states differing by ε can have their separation stabilized as a physical fact.

Branch 1: Divergent encoding burden exceeds physical realizability.

By the General Encoding Burden Lemma (§3.3), forming a fact at resolution ℓ requires encoding burden $\log_2(V/\ell^3) \to \infty$ as $\ell \to 0$. The spatial localization formula is one representative realization of the more general statement proved by that lemma: whenever refinement drives the number of distinguishable outcomes N without bound, the encoding burden $\log_2 N$ diverges — regardless of what physical quantity is being resolved. This burden cannot be met within finite causal duration and finite operational resources (Physical Realizability Lemma). No physically realizable fact-formation process can sustain it. Therefore no record can stabilize arbitrarily fine distinctions.

Branch 2: No robust records destroys reproducibility.

A reproducible fact requires that a physical record encoding an outcome can be reliably recovered by an independent agent. A record susceptible to erasure by any perturbation s' with $|s - s'| < \varepsilon$ for arbitrarily small ε cannot be reliably recovered — any perturbation below the non-existent lower bound transforms the record into an undecidable one. Therefore no fact is reproducible.

Branch 3: No reproducibility destroys lawhood.

By Theorem 1, physical laws require fact-stable state classes. Without reproducibility, no fact-stable state classes exist. No physical laws hold. This contradicts the assumption that physics exists in this universe. ■

3.5 Minimality

Finite distinguishability is the *weakest* condition that prevents record instability. We do not require integer-valued states, a spatial lattice, or any particular discrete geometry. We require only a fixed lower bound $\delta\_min > 0$ on resolvable differences.

Lemma (δ_min does not imply spatial discreteness). *The existence of a minimum fact-supporting distinguishability scale $\delta\_min > 0$ does not require spacetime to be discrete. It requires only that distinctions below δ_min cannot be stabilized as recoverable facts. A theory may possess a mathematically continuous configuration space while still admitting only finitely* *many operationally distinguishable states per bounded region. δ_min bounds the distinguishability of facts, not the existence of intermediate configurations.*

Proof. The theorem concerns fact-supporting distinguishability — the ability to stabilize and recover a distinction as a physical fact — not the mathematical cardinality of a configuration space. A continuous configuration space assigns uncountably many labels to states, but those labels are extracted by fact-forming processes of finite resolution. The operationally accessible distinctions are those that can be recorded and recovered by finite observers. The existence of δ_min > 0 is equivalent to a bound on the resolution of such operations; it says nothing about whether intermediate states exist mathematically. The difference between "mathematically continuous" and "operationally continuous" is exactly the difference between logical and fact-supporting distinguishability (§3.2). □

This lemma preempts the most common misreading of A1: that VERSF requires a lattice or discrete spacetime. It does not. It requires a finite resolution floor on recoverable distinctions, which is compatible with any continuous mathematical structure.

Representative weaker alternatives all fail. The following candidates illustrate why no weakening of this condition suffices:

- "Resolution finite only on average": fails because worst-case perturbations — which exist whenever δ_min has no fixed lower bound — still destroy individual records. A guarantee that holds only on average provides no protection against the single adversarial perturbation that erases a specific fact.
- "Resolution finite only for macroscopic systems": fails because Theorem 1 requires fact-stability for all fact-forming events, not just macroscopic ones. A single sub-microscopic record with no stability floor is sufficient to undermine reproducibility.
- "Resolution finite only in equilibrium": fails for the same reason — facts must be stable outside thermodynamic equilibrium for any dynamical law to apply.

A1 is therefore not merely sufficient; it is the minimum sufficient condition: the weakest bound that closes all three failure modes.

Lemma (Context-dependent floors do not secure lawhood). *A resolution floor δ_min that varies by context cannot provide a universal operational state space for physical law, because worst-case perturbations remain admissible somewhere in the law's intended domain.*

Proof. Suppose δ_min(c) depends on context c — physical regime, thermodynamic state, location, or any other parameter. Then for any proposed floor value ε > 0, there exists some context c* where δ_min(c*) < ε. In that context, a perturbation of magnitude ε remains below the floor, erasure of records is possible, and the reproducibility required by Theorem 1 fails at scale ε. A physical law that admits even one context where record stability fails does not define a universal state space — the equivalence classes of §2.2 become context-dependent and cannot be identified across instances. Context-dependent distinguishability floors are therefore insufficient to ground lawhood, regardless of how they vary. □

Any constraint weaker than a uniform lower bound on resolvable differences still permits adversarial perturbations below the context-dependent floor.

3.6 Independence Tests

What survives if A1 is removed?

If finite distinguishability is denied, then:

- Physical records have no stability guarantee against sub-granular perturbation.
- No reproducible measurement is possible: every recorded value is susceptible to arbitrary re-specification.
- The concept of a physical constant — a reproducibly measurable quantity — collapses: any measured value could be displaced by an arbitrarily small perturbation and yield a different result.
- Lawhood fails by Theorem 1.

What is preserved? Mathematical formalisms may still be statable, but they describe nothing that a physical agent could ever confirm or use.

3.7 What Is Not Assumed

Theorem 2 does not assume:

- A spatial lattice or any discrete geometry.
- Quantum mechanics or the Heisenberg uncertainty principle.
- Any specific dynamics or field equations.
- Any VERSF-specific substrate structure.
- A minimum length as a primitive.

It assumes only that fact-forming processes are physically realizable — that they operate within finite causal reach and finite operational resources — and that Theorem 1 holds. Both of these are presuppositions of any physical theory, not features of VERSF.

3.8 Independent Appearance Across Physics

Finite distinguishability is not unique to VERSF. Each of the three VERSF axioms corresponds to a known fundamental limit already identified by independent branches of physics:

VERSF Axiom	Known manifestation	Field
A1: Finite distinguishability	Heisenberg uncertainty; Shannon channel capacity	QM; information theory
A2: Irreversible commitment	Decoherence; Landauer's principle	QM; thermodynamics

VERSF Axiom	Known manifestation	Field
A3: Finite localization capacity	Holographic bound; Planck-scale minimum length	Gravity; quantum gravity

The VERSF axioms are not inventions. They are a unifying reinterpretation of limits that physics has discovered independently, showing that each limit is a manifestation of the same structural necessity.

For A1 specifically:

- **Quantum mechanics**: the Heisenberg uncertainty relations bound the simultaneous resolvability of conjugate variables. In the VERSF reading, this is the dynamical *expression* of finite fact-supporting distinguishability, not its origin.
- **Shannon information theory**: the channel capacity theorem establishes a finite upper bound on distinguishable signals in any physical channel of finite bandwidth and power. Infinite distinguishability would require infinite bandwidth — no physical channel has this.
- **Holographic bound**: the Bekenstein bound limits information content per bounded region to $A/4\ell_P^2$ bits — a finite upper bound on fact-supporting distinguishability per region.

---

4. Theorem 3: Irreversible Commitment Is Unavoidable {#4}

4.1 Statement

Theorem 3 (VERSF Axiom A2 Is Unavoidable). *Within the class of local probabilistic physical theories — theories in which local operations have inaccessible dilation variables and the total dynamics is reversible — every fact-forming operation deposits an irreversible correlation residue. No locally-bounded agent has joint control over the full dilation environment, by relativistic causality; every local fact formation is therefore unconditionally irreversible within this class. Deterministic classical mechanics, treated as a complete fundamental ontology without measurement or finite agents, lies outside the scope of this theorem. A universe whose ontology is purely deterministic and contains no probabilistic or measurement layer may still admit mathematical dynamics, but it cannot support the agent-accessible fact formation required by Theorem 1. The scope of Theorem 3 therefore coincides with the class of physical theories capable of producing law-governed facts accessible to finite observers — exactly the class whose existence this paper establishes as the minimum precondition for physics.*

4.2 Deterministic Ontologies with Finite Observers

Proposition (Finite-observer reduction). *Any deterministic ontology that supports law-governed facts accessible to finite observers induces an observer-level local probabilistic theory with inaccessible environmental variables.*

Proof. Let U be a deterministic universe with complete microstate M(t). A finite observer O has a causal horizon — it can only access the information in its past light cone. The set of microstates consistent with O's accessible information at time t forms an equivalence class [M] under O's epistemic access relation. Because M(t) is complete and deterministic but O's accessible state $m(t) \subsetneq M(t)$ is not, the variables $V(t) = M(t) \setminus m(t)$ are inaccessible to O. The observer's local description of any process is therefore a probability distribution over microstates consistent with m(t). The inaccessible variables V(t) function as dilation variables: from O's perspective, local fact-forming operations have outcomes that are not determined by m(t) alone — they depend on the hidden V(t). This is formally equivalent to the Tier I structure of §4.5: local operation on accessible state m(t), environment V(t) traced out, total dynamics deterministic. □

A deterministic ontology does not evade this conclusion once finite observers are introduced. A finite observer cannot access the complete microstate of a deterministic universe. Their effective description of local processes is necessarily probabilistic, with inaccessible environmental variables playing the role of dilation variables. Deterministic mechanics with finite observers is therefore formally equivalent to a local probabilistic theory with inaccessible dilation variables, and falls within the scope of Theorem 3.

This is not an ad hoc maneuver. Any universe capable of supporting law-governed facts accessible to finite observers — the condition established by Theorem 1 as the precondition for physics — is one in which finite observers exist and form facts. The introduction of finite observers is not an additional assumption; it is already implied by the requirement that facts be accessible. Theorem 3 therefore applies to any universe that satisfies Theorem 1's precondition, including deterministic ones.

4.3 The Correlation-Export Formulation

The content of Theorem 3 is not that "facts cost entropy." That statement is too weak — it sounds like the second law, which is a statistical claim about ensembles. Theorem 3 is structural and holds for single fact-formation events regardless of the thermodynamic state of the universe.

The correct formulation is: **fact formation requires exporting correlation structure beyond the agent's causal reach.** Entropy increase is the measurable *footprint* of this export — not its cause, and not the same thing.

This distinction is critical. The claim "facts cost entropy" admits a counter: perhaps in a low-entropy universe, fact formation could be arranged reversibly. The claim "fact formation exports correlations beyond causal reach" does not admit this counter, because the irreversibility is built into causal structure, not thermodynamic state.

4.4 Lemma: Bounded Access from Causal Locality

Lemma (Causal Boundedness). *No finite agent can control degrees of freedom outside its causal past. Joint disentangling of a local system from its environment is physically inadmissible unless the entire environment lies within the agent's causal past — which it never does for any locally-bounded agent in an infinite universe.*

Proof.

The lemma has two independent supports, each sufficient on its own.

Support 1: Relativistic causal structure.

In any relativistic spacetime, signals propagate at or below c. An agent at spacetime point P can only have received information from, and therefore can only have causal control over, events in its causal past $J^{-}(P)$. The dilation environment E extends, in principle, across the universe. Regions spacelike-separated from P lie in E but outside $J^{-}(P)$. No operation at P can read out or manipulate those degrees of freedom without superluminal signaling. Therefore global disentangling requires accessing degrees of freedom outside $J^{-}(P)$, which is physically inadmissible.

Support 2: Finite control capacity.

Even setting aside causal constraints, reversing a correlation export requires that the agent: (i) acquire complete information about the correlation structure exported into E; (ii) store this information; (iii) coordinate a reversal operation on E. The information and storage resources required scale with the size of E. Since the environment grows unboundedly with time — as successive fact-formation events export further correlations — the control resources required for global disentangling grow without bound. No finite agent possesses infinite control resources. Global disentangling is therefore impossible not only in principle (by causality) but in practice (by resource exhaustion).

The two supports are independent: relativistic causality blocks the reversal in any spacetime with finite signal speed; finite control capacity blocks it even in a universe without a speed limit. Together they provide independently grounded support for the irreversibility of commitment. □

This lemma grounds A2 not in thermodynamics but in the causal structure of relativistic spacetime. The irreversibility of fact formation is not a practical limitation arising from complexity or cost — it is a consequence of causal locality.

4.5 Proof via CPTP Maps and the Entanglement Ledger

Proof.

The proof has two tiers. Tier I sets up the abstract structural conditions from which the irreversibility conclusion follows. Tier II shows how quantum mechanics — via CPTP maps and the Stinespring dilation — instantiates that structure exactly. Tier I identifies the structural conditions under which irreversible correlation export occurs; Tier II shows that quantum mechanics realizes these conditions exactly through CPTP maps and analogously in other local probabilistic theories.

Tier I — Abstract structural conditions.

The irreversibility argument applies in any theory satisfying three conditions:

1. **Local operations**: the agent can only directly manipulate a bounded subsystem S.
2. **Inaccessible dilation variables**: S is coupled to a larger environment E whose degrees of freedom the agent cannot fully access — established by the Causal Boundedness Lemma.
3. **Reversibility of the total dynamics**: the global evolution of S ⊗ E is reversible, so information is not destroyed but relocated.

Given these three conditions, any local operation on S that produces a definite outcome must relocate information from S into E. Since E is inaccessible to the agent, that relocation is irreversible from the agent's perspective. This structural argument can be stated as a proposition independently of any specific formalism:

Proposition (Observer-level irreversibility without quantum formalism). *Whenever a finite observer interacts locally with a subsystem whose full determining variables are not accessible, any registration of a definite outcome requires that information distinguishing the outcome be transferred into variables outside the observer's direct control. Since reversal would require joint control over those variables — which is excluded by causal boundedness — the registration is observer-irreversible independently of the specific formalism used to represent the process.*

Proof. Let the observer have direct access to subsystem S and no direct access to environment E. A definite outcome distinguishes one state from alternatives; any physical record of this distinction requires that the distinguishing information be encoded somewhere. If the encoding is within S, it is accessible to the observer and could in principle be reversed — but reversing it would require erasing the distinction, eliminating the fact. If the encoding is within E, it is outside the observer's direct control by assumption. Causal boundedness (§4.4) establishes that the observer cannot jointly control S ⊗ E; therefore the relocation into E is not reversible by any locally-admissible operation. The registration is observer-irreversible. □

This proposition carries the theorem-level weight of Tier I. Tier II shows that quantum mechanics realizes this structure exactly through CPTP maps and the Entanglement Ledger — not as an approximation but precisely. Note that the irreversibility established here arises from causal boundedness — the structural fact that an observer cannot jointly control S ⊗ E — not from thermodynamic irreversibility or entropy production. A2 is therefore a structural necessity, not a statistical tendency.

Tier II — Quantum realization via CPTP maps.

Step 1: All locally admissible quantum operations are CPTP maps.

In quantum mechanics, any physically admissible local operation is a completely positive trace-preserving (CPTP) map. CPTP maps arise from exactly the three Tier I conditions: local operation on S, environment E discarded by tracing out, global dynamics unitary. The Stinespring dilation theorem gives the exact form:

$$\mathcal{N}(\rho_S) = \mathrm{Tr}_E[U (\rho_S \otimes |0\rangle\langle 0|_E) U^\dagger]$$

This is not an approximation — it is the exact quantum representation of the Tier I structure. A critic who accepts Tier I but rejects quantum mechanics still accepts the conclusion; a critic who accepts quantum mechanics has the argument made precise by Tier II.

Step 2: Fact formation reduces local accessible correlations and exports them.

If the CPTP map transitions S to a definite state, accessible correlations within S decrease. Under the Stinespring dilation, global von Neumann entropy $S(\rho_{SE})$ is constant. Therefore:

$S(\rho_E)$ increases by at least $\Delta S_{min} \sim k_B \ln 2$ per bit of fact formed.

Correlations are not destroyed — they are exported from S into E.

Step 3: The Entanglement Ledger is monotone.

Define:

$$\mathcal{L}(S, E) = S(\rho_S) + S(\rho_E) - S(\rho_{SE}) \geq 0$$

Under global unitary evolution, $S(\rho_{SE})$ is constant. Local operations on S alone can only increase $S(\rho_E)$. Therefore $d\mathcal{L}/dt \geq 0$ monotonically for any locally-bounded agent.

Step 4: The export is irreversible by the Causal Boundedness Lemma.

The exported correlations reside in E, which extends outside $J^-(P)$. By the Causal Boundedness Lemma, no locally-bounded agent at P can access, read out, or reverse these correlations. The Ledger increment is therefore not merely costly to reverse — it is causally inadmissible to reverse. ■

4.6 Minimality

Irreversible commitment is the *weakest* condition that prevents reversible erasure of physical history. We do not require thermodynamic equilibration, macroscopic decoherence, a specific collapse mechanism, or any particular quantum interpretation. We require only that local fact-forming operations export correlation beyond causal reach.

Representative weaker alternatives all fail. The following candidates illustrate why no weakening of this condition suffices:

- "Irreversibility only for macroscopic systems": fails because a single microscopic fact-formation event that remains reversible would allow the physical past to be retroactively altered at that event. Theorem 1's requirement for a definite domain of prior states then fails.
- "Irreversibility only in thermodynamic equilibrium": fails for the same reason — facts formed outside equilibrium must still be definite and recoverable, or initial-value formulations of physics break down.
- "Irreversibility only statistically": fails because statistical irreversibility permits individual reversals, which is sufficient to undermine the definiteness of any specific past fact.

A2 is the minimum condition for a universe to have a definite past: the weakest bound that prevents retroactive erasure at every level. A context-dependent weakening — irreversibility that holds only in certain physical regimes — still permits erasure in the remaining regimes, which is sufficient to undermine the definiteness of past facts formed in those regimes.

4.7 Independence Tests

What survives if A2 is removed?

If irreversible commitment is denied, then:

- No boundary between past and future is structurally enforced. Records could be erased without accessing the environment, eliminating the distinction between past and future states.
- The Entanglement Ledger becomes non-monotone, permitting histories in which later states have strictly less correlation structure than earlier ones — allowing physical laws to retroactively fail.
- Initial-value formulations of physical laws lose their meaning: any "initial" state could be undone, making the notion of a definite past incoherent.
- A2 establishes a monotone structural asymmetry between past and future — the Entanglement Ledger increases in one direction only. Whether this asymmetry fully accounts for the phenomenological arrow of time is a further question; what A2 establishes is the minimal structural prerequisite for any such account.

What is preserved? Time-symmetric laws remain statable as mathematics, but they cannot be applied to a universe with a definite past — because without A2, there is no definite past.

4.8 What Is Not Assumed

Theorem 3 does not assume:

- The second law of thermodynamics. (The second law follows from Theorem 3 as a statistical consequence; the reverse does not hold.)
- Thermodynamic equilibration or macroscopic decoherence.
- Any specific collapse mechanism or quantum interpretation.
- Any VERSF-specific structure.

It assumes only: (i) locally admissible operations admit CPTP realization in quantum theory and analogous forms in other local probabilistic theories — this is the scope established in §4.1 and the Tier I argument of §4.5; (ii) relativistic causal structure limits agents to their causal past — standard special and general relativity. Neither assumption is specific to VERSF.

The VERSF Entanglement Ledger is the formal device for tracking $\mathscr{L}$ across cosmic history. It is the *representation* of the structural necessity proved in Theorem 3, not its source.

4.9 Independent Appearance Across Physics

The correspondence table in §3.8 maps each axiom to its independent manifestation in physics. For A2 specifically:

- **Decoherence theory**: the irreversible leakage of quantum coherence into environmental degrees of freedom is exactly the correlation export of Theorem 3. Decoherence is the empirical signature of A2.
- **Landauer's principle**: erasing one bit of information necessarily dissipates at least k_B T ln 2 of energy to the environment. This is a thermodynamic footprint of irreversible correlation export.
- **Black hole information**: the causal inaccessibility of information behind a horizon is a direct instance of the Causal Boundedness Lemma — the horizon is the boundary of $J^-(P)$ for exterior agents.

Each is an independent empirical discovery of the same structural necessity proved in Theorem 3.

---

5. Theorem 4: Finite Localization Capacity Is Unavoidable {#5}

5.1 The Closure Problem

Theorems 2 and 3 establish:

- **A1**: facts require a minimum spatial granularity δ_min and a finite information commitment per fact.
- **A2**: every fact exports irreversible correlation beyond the agent's causal reach.

Taken together, these two necessities generate a potential self-inconsistency. If the substrate permitted arbitrarily dense fact localization — packing facts into ever-smaller regions without limit — then either:

(a) Fact density would require distinctions below δ_{min} — violating A1; or
(b) Correlation export per unit volume would diverge — violating A2's requirement that the environment have finite absorption capacity.

Either outcome is a contradiction with an already-established necessity. Theorem 4 proves that A3 — finite localization capacity — is the closure condition that resolves this inconsistency. It is not a new assumption introduced to patch the framework. It is what A1 and A2 jointly require in order to be simultaneously satisfiable.

5.2 Statement

Theorem 4 (VERSF Axiom A3 Is Unavoidable). *Theorems 2 and 3 jointly entail that any fact-supporting universe must have a finite upper bound on stabilizable distinction density — the maximum number of distinguishable, stably-recorded facts per unit volume. From this bound, finite localization capacity follows immediately: there exists* $\ell > 0$ below which fact formation is impossible. In VERSF, finite localization capacity is represented constitutively as a maximum substrate tension $\mathcal{T}_{max}$.*

The three-stage derivation is deliberate. The primary target is *stabilizable distinction density* — a direct combination of the two already-proven necessities. Finite localization capacity is what a bound on distinction density means geometrically. $\mathcal{T}_{max}$ is what finite localization capacity looks like in the dimensional language of the VERSF void substrate. A critic who objects to $\mathcal{T}_{max}$ as a primitive has already conceded the primary target; the tension representation follows from dimensional analysis given c (established by relativistic causal structure) and $\hbar$ (taken as an empirical input in the present paper).

The exhaustiveness argument that closes this proof depends on the Scale-Independence Lemma (§5.4); the full boxed-in structure is presented there once that lemma is established.

5.3 Lemma A: Impossibility of Unbounded Fact Density

Lemma A. *If localization capacity is unbounded and fact density is preserved as $\ell \to 0$, then Theorem 2 (A1) is violated.*

Proof.

Define the fact density in a region of size ℓ as:

$$\rho_f(\ell) = N(\ell) / \ell^3$$

where $N(\ell)$ is the number of distinguishable facts supportable in a region of size ℓ.

By Theorem 2, the minimum linear distinguishable scale is $\delta_{min} > 0$. A region of linear size ℓ can therefore contain at most (ℓ/δ_{min}) distinguishable sites along each axis, giving:

$$N_{max}(\ell) \lesssim (\ell / \delta_{min})^3$$

and therefore a maximum stabilizable fact density:

$$\rho_f(\ell) = N(\ell) / \ell^3 \lesssim \delta_{min}^{-3}$$

This bound is determined entirely by δ_{min}, which is a fixed constant by Theorem 2 — not a function of ℓ. The maximum fact density is therefore ℓ-independent and finite: $\rho_f \lesssim \delta_{min}^{-3} < \infty$ for all ℓ.

Suppose now that unbounded localization is claimed — that ρ_f can be maintained as $\ell \to 0$ by shrinking δ_{min} with ℓ. Any such claim is equivalent to claiming that facts can be formed at sub-δ_{min} resolution: if the region size ℓ falls below δ_{min}, then a fact localized to that region would require distinguishing a state's location at a scale below the minimum resolvable difference — which is precisely what Theorem 2 prohibits. A δ_{min} that is allowed to shrink with ℓ is therefore not a fixed lower bound at all; it is a bound that vanishes as $\ell \to 0$, which is functionally infinite distinguishability. This contradicts Theorem 2 directly. ■

5.4 Lemma (Scale-Independence of Minimum Fact Commitment)

Proposition (Minimum fact commitment is structural, not scale-conventional). *If a fact is to count as a new physically registrable distinction, the environmental imprint that individuates it must exceed a nonzero structural threshold independent of localization scale. This threshold is not a contingent property of the physical medium — it is the minimum unit of physically registrable novelty required for a new fact to exist as distinct from prior facts.*

This proposition motivates the following lemma. If ΔS_{min} could shrink with scale, the notion of a "new distinct fact" would lose its footing: at sufficiently small scales, new facts would leave no distinguishable imprint, making them indistinguishable from non-events. The lemma proves this possibility is structurally foreclosed.

Lemma (ΔS_{min} is ℓ-independent). *The minimum correlation commitment ΔS_{min} required to register a fact cannot decrease with localization scale ℓ without collapsing either distinguishability (A1) or irreversible commitment (A2).*

Proof.

Structural premise. In a local probabilistic theory with inaccessible dilation variables (the Tier I conditions of §4.5), distinct facts cannot be specified solely by their local state: the physical record of their occurrence must be encoded in correlations with the environment. Two facts are physically distinct only if they produce distinguishable correlation footprints in the environment; without such a footprint, no agent bounded by causal structure can register the distinction. This premise relies only on Tier I structure and causal boundedness, both already established.

Suppose $\Delta S\_min = \Delta S\_min(\ell)$ decreases with ℓ, tending to zero as $\ell \to 0$. Then as facts are localized to smaller regions, the irreversible commitment required to distinguish them approaches zero. Consider what this means:

Branch 1: Indistinguishability collapse. If the commitment cost of registering the distinction between two fact-formation outcomes approaches zero, those outcomes become increasingly indistinguishable in terms of their physical signature. In the limit $\Delta S\_min \to 0$, two distinct facts leave no distinguishable trace in the environment — they produce the same correlation footprint. But facts that leave no distinguishable trace are not distinct facts; they are the same fact. This collapses the minimum distinguishable scale $\delta\_min \to 0$, violating A1.

Branch 2: Reversibility of commitment. Alternatively, suppose the two facts remain distinguishable despite $\Delta S\_min \to 0$. Then arbitrarily fine distinctions would require arbitrarily little irreversible commitment. This severs the link between fact formation and irreversible registration established by Theorem 3: if the commitment cost can be made arbitrarily small, it can be made negligible, and the distinction between reversible and irreversible operations disappears at that scale. This violates A2.

In either branch, $\Delta S\_min \to 0$ violates an already-established necessity. Therefore ΔS_min has a nonzero lower bound independent of ℓ: there exists $\Delta S\_0 > 0$ such that $\Delta S\_min \geq \Delta S\_0$ for all ℓ. □

5.5 Lemma (Finite Correlation-Absorption Capacity of Finite Regions)

Lemma (Finite absorption capacity). *Any finite spacetime region accessible to a fact-forming operation has finite correlation-absorption capacity — it can absorb only a finite number of distinguishable correlation increments.*

Proof. The proof follows from the paper's own earlier machinery in four steps:

(1) Finite causal reach. By the Physical Realizability Lemma (§3.3), any fact-forming operation has finite causal reach — the receiving region is a finite spacetime region.

(2) Finite available action. By the same lemma, any physically realizable process has finite available action $E \cdot \tau$ within its causal region.

(3) Finite independently accessible channels. A finite spacetime region with finite energy content does not imply a finite number of field modes in a quantum field theory; however, only a finite subset of those modes can be excited and accessed by any physical operation within that region. Any attempt to independently address N channels requires at least $\log_2 N$ bits of control information to specify which channel is being accessed; a process of finite action cannot supply unbounded control information. Because the encoding burden required to independently address N channels grows as $\log_2 N$, any attempt to access infinitely many channels within a region of finite action would require unbounded encoding resources, contradicting the Physical Realizability Lemma (§3.3). By the Unbounded Refinement Lemma (§3.3), a process with finite causal reach and finite available action therefore cannot interact with infinitely many

independent channels. The number of independently accessible channels is finite — a direct consequence of the encoding lemmas of §3.3, not a new assumption. The logic is: finite action → finite control information → finite addressable channels → finite correlation capacity.

(4) Finite distinguishable correlation states. The number of independently accessible channels is finite (step 3). Each channel admits a finite number of distinguishable states (by step 2 and the General Encoding Burden Lemma). The number of independently accessible distinguishable correlation channels is therefore finite, and the total number of distinguishable correlation states the receiving region can absorb is correspondingly finite.

The finiteness follows entirely from the earlier lemmas without invoking Bekenstein bounds, ℓ_P, or any VERSF-specific structure. Because the accessible channels are finite and each channel admits only finitely many distinguishable states under finite action, the number of distinct correlation increments that can be written into the region is finite. This bound holds independently of the detailed microphysics of the substrate. □

5.6 Lemma B: Impossibility of Unbounded Correlation Export

Lemma B. *If localization capacity is unbounded, the correlation export per unit volume diverges as $\ell \to 0$, exceeding the finite correlation-absorption capacity of any finite receiving region, in violation of Theorem 3 (A2).*

Proof.

By the Scale-Independence Lemma (§5.4), $\Delta S_{min} \geq \Delta S_0 > 0$ independently of ℓ. Each fact-formation event therefore exports at least ΔS_0 of correlation to the environment, regardless of how small the localization region. For facts localized to regions of size ℓ, the minimum correlation export per unit volume is:

$$\Delta S_{vol} \sim \Delta S_0 / \ell^3 \to \infty \text{ as } \ell \to 0$$

The two cases correspond to the two possible behaviors of $\rho_f(\ell)$:

Case 1: ρ_f finite as $\ell \to 0$.
Fact density is preserved. By Lemma A, this requires $\delta_{min} \to 0$, violating Theorem 2. Already a contradiction.

Case 2: $\rho_f \to 0$ as $\ell \to 0$.
Fact density falls, but the correlation export required per fact-formation event, absorbed by a region of volume $\sim\ell^3$, still produces a local volumetric demand of $\Delta S_0/\ell^3 \to \infty$. That is: a single fact exports ΔS_0; when that fact is localized to a region of size ℓ, the per-volume absorption demand is $\Delta S_0/\ell^3$, which diverges as $\ell \to 0$ regardless of how sparse the facts become. By the Finite Absorption Capacity Lemma (§5.5), any finite spacetime region accessible to the fact-forming process can absorb only finitely many distinguishable correlation increments. An export rate that diverges as $\ell \to 0$ cannot be absorbed by any such region regardless of its size.

Divergent export into a finite-capacity environment is therefore impossible under Theorem 3. Contradiction.

In both cases, unbounded localization produces a violation of an established necessity. ■

5.7 The Boxed-In Structure

With the Scale-Independence Lemma (§5.4) established, the possible behaviors of stabilizable distinction density as $\ell \to 0$ can now be exhaustively classified.

The boxed-in structure. Suppose stabilizable distinction density is unbounded — that facts can be packed into regions of arbitrarily small size without violating any constraint. Then exactly one of two behaviors must hold as $\ell \to 0$:

- *Case 1*: distinction density is preserved as $\ell \to 0$. Then δ_{min} must shrink with ℓ, which is infinite distinguishability — violating Theorem 2 (A1). ✗
- *Case 2*: distinction density falls as $\ell \to 0$. Then the correlation export per unit volume $\Delta S_{vol} \sim \Delta S_0/\ell^3$ still diverges (since $\Delta S_{min} \geq \Delta S_0 > 0$ by §5.4) — violating Theorem 3 (A2). ✗

The Scale-Independence Lemma is precisely what closes off a third behavior: $\Delta S_{min}(\ell) \to 0$ as $\ell \to 0$ would permit Case 2 to escape, but §5.4 proves this leads to contradiction. Within the class of physically admissible behaviors compatible with Theorems 2 and 3, Cases 1 and 2 are exhaustive. The bound on distinction density is structurally entailed — not merely the most natural closure condition, but the one no fact-supporting universe can avoid.

Trilemma. Any attempt to localize facts to arbitrarily small regions must do exactly one of three things, each of which leads to contradiction:

1. *Shrink the distinguishability floor*: allow $\delta_{min} \to 0$ to preserve fact density. This is infinite distinguishability — violating A1. ✗
2. *Shrink the commitment floor*: allow $\Delta S_{min} \to 0$ to avoid divergent correlation export. The Scale-Independence Lemma shows this collapses either distinguishability or the link between fact formation and irreversible registration — violating A1 or A2. ✗
3. *Accept divergent correlation export*: keep $\Delta S_{min} \geq \Delta S_0 > 0$ and ρ_f decreasing. The export per unit volume $\Delta S_0/\ell^3 \to \infty$ exceeds the finite correlation-absorption capacity of any finite receiving region — violating physical realizability under Theorem 3. ✗

There is no fourth option. Arbitrarily fine fact localization is structurally impossible in any universe that satisfies Theorems 2 and 3. The trilemma shows that arbitrarily fine fact localization requires either infinite distinguishability, vanishing commitment, or infinite environmental capacity — each incompatible with the structural conditions already established. The trilemma therefore shows that arbitrarily fine fact localization is incompatible with any universe satisfying lawhood, realizability, and causal boundedness. Equivalently: arbitrarily fine fact localization requires violating either distinguishability (A1), commitment (A2), or realizability — the three pillars the preceding theorems have shown to be unavoidable.

5.8 Theorem 4 from Lemmas A and B

Proof of Theorem 4.

Lemma A: unbounded stabilizable distinction density with ρ_f finite → structurally requires $\delta_{min} \to 0$ → contradicts Theorem 2 (A1).
Lemma B: unbounded stabilizable distinction density with any ρ_f behavior → entails $\Delta S_{vol} = \Delta S_0/\ell^3 \to \infty$ (where $\Delta S_0 > 0$ by the Scale-Independence Lemma) → contradicts Theorem 3 (A2).

The Scale-Independence Lemma (§5.4) is what makes the exhaustiveness claim rigorous: it rules out the third behavior — $\Delta S_{min}(\ell) \to 0$ as $\ell \to 0$ — that would otherwise permit ΔS_{vol} to remain bounded even as $\ell \to 0$. With ΔS_{min} bounded below by ΔS_0, the two cases in Lemmas A and B are exhaustive within the class of physically admissible behaviors compatible with Theorems 2 and 3. No third case exists within this class. Unbounded stabilizable distinction density admits no physically viable alternative to contradiction.

Step 1 — Finite distinction density bound: There exists a maximum stabilizable distinction density σ_{max} such that $\rho_f(\ell) \leq \sigma_{max}$ for all ℓ.

Lemma (Bounded density implies exclusion scale). *If the maximum stabilizable distinction density σ_{max} is finite, then there exists a minimum region size $\ell > 0$ below which no region can host a stabilizable fact.*

Proof. A region of volume V can host at most $\sigma_{max} \cdot V$ stabilizable facts. For a region to host at least one fact, $\sigma_{max} \cdot V \geq 1$, which requires $V \geq 1/\sigma_{max}$. The minimum admissible volume is therefore $V^* = 1/\sigma_{max} > 0$, implying a minimum linear scale $\ell^* \sim \sigma_{max}^{(-1/3)} > 0$. Any region smaller than ℓ^* has volume below V^* and cannot contain even a single stabilizable fact without violating the density bound. □

Step 2 — Finite localization capacity: By the Bounded Density Lemma, $\ell^* > 0$ exists below which fact formation is impossible. Packing a fact into a region smaller than ℓ^* would require either a distinction below δ_{min} (violating A1) or a correlation export that exceeds the environment's finite capacity (violating A2) — both of which are the routes closed by Lemmas A and B.

Step 3 — Constitutive representation in VERSF: The finite localization capacity is a physical property of whatever substrate the universe employs. In VERSF, the substrate's resistance to fact localization is characterized by its tension — the energy cost per unit area of sustaining a fact-supporting boundary. The minimum localization scale ℓ^* corresponds to a maximum sustainable tension:

$$\mathcal{T}(\ell) \sim \hbar c / \ell^2 \leq \mathcal{T}_{max}$$

Here c is the signal-speed bound established by relativistic causal structure, and $\hbar$ is taken as an empirical input — the quantum of action — whose structural derivation within VERSF is developed in companion work. In this paper $\hbar$ enters only as an empirical constant representing the quantum of action. The structural necessity established here is the existence of a finite localization capacity; identifying its numerical value with Planck-scale quantities is a downstream constitutive step developed in companion work, not a premise of the present argument. $\mathcal{T}$_max is not postulated. It is the dimensional form that finite localization capacity necessarily takes given these two constants; its identification with c^4/G is a downstream constitutive step. ■

5.9 Minimality

Finite localization capacity is the *weakest* condition that prevents Theorems 2 and 3 from being jointly inconsistent. We do not postulate a specific minimum length, a lattice, or any Planck-scale structure as a primitive. We require only that there exist *some* finite upper bound on stabilizable distinction density.

Representative weaker alternatives all fail. The following candidates illustrate why no weakening of this condition suffices:

- "Capacity bound only on average": fails because locally divergent distinction density — even in a single small region — produces a locally divergent correlation export that violates A2 at that location, regardless of the global average.
- "Capacity bound only statistically": fails for the same reason — a single statistical fluctuation to arbitrarily high distinction density produces the A1/A2 violation at that event.
- "Capacity bound only at thermodynamic equilibrium": fails because A1 and A2 are structural constraints on every fact-formation event, not ensemble averages. A capacity violation at one non-equilibrium event is sufficient.

A3 is not a third independent assumption. It is the minimum closure condition: the weakest bound that makes A1 and A2 simultaneously satisfiable under all physically admissible conditions. A context-dependent capacity bound — one that holds only on average or only in equilibrium — still permits local violations that produce the A1/A2 contradictions identified in Lemmas A and B.

5.10 Independence Tests

What survives if A3 is removed?

If finite localization capacity is denied, then:

- Theorems 2 and 3 cannot both be satisfied: Lemma A shows A1 is violated, Lemma B shows A2 is violated. The first two necessities become mutually self-undermining.

- No stable fact-forming scale exists: the identity collapse scale ℓ_P — the minimum region that can host a fact — has no value, and the dimensional structure of physics loses its ground.
- The physical constants G and ℓ_P, which are determined by $\mathcal{T}_{max}$, become undefined.

What is preserved? Without A3, A1 and A2 are individually statable but jointly inconsistent. The framework collapses.

5.11 What Is Not Assumed

Theorem 4 does not assume:

- The Planck scale or any specific minimum length as a primitive.
- A spatial lattice or discrete geometry.
- Gravity or any gravitational coupling.
- Any VERSF-specific substrate structure.

A critic who argues A3 is ad hoc must identify which step in Lemma A or Lemma B fails. The theorem holds for any substrate satisfying Theorems 2 and 3.

5.12 Independent Appearance Across Physics

Finite localization capacity appears, independently, across several domains:

- **Quantum gravity**: combining the Heisenberg uncertainty principle with Schwarzschild radius arguments gives a minimum resolvable length of order ℓ_P. This is the standard heuristic argument for the Planck scale — Theorem 4 derives it from structural necessity rather than heuristic combination.
- **Holographic principle**: the Bekenstein–Hawking entropy bound $A/4\ell_P^2$ is precisely a finite localization capacity bound per surface element — no more than one fact per Planck area.
- **Causal dynamical triangulations / loop quantum gravity**: both frameworks predict a minimum length at the Planck scale. Theorem 4 shows this is not a feature of any particular quantum gravity approach but a necessity of fact-bearing physics.

---

6. The Axiomatic Hierarchy {#6}

The logical chain in one sentence: **lawhood requires stable facts; stable facts require finite distinguishability and irreversible commitment; these two necessities jointly require finite localization capacity.**

In its most compact formal expression:

lawhood ⟹ fact-stable states ⟹ { A1 (finite distinguishability), A2 (irreversible commitment) } ⟹ A3 (finite localization capacity)

Everything else in VERSF begins here. The void substrate, the Entanglement Ledger, $\mathcal{T}$_max, the Planck scale, and the structure of gravity are all downstream consequences of this four-step chain. None of them can be questioned without first identifying which link in the chain fails.

The three VERSF axioms are not parallel postulates of equal standing. They form a strict hierarchy of logical dependence:

Level 0 — Existence of Physics (Primitive)

Physics is possible. Some universe supports physical laws.

Level 1 — Theorem 1

Physics requires stable, reproducible physical facts. Without fact-stable state classes, no law can act on any domain. This is the formal precondition for any physical theory.

Level 2 — Theorems 2 and 3 (First-order necessities)

Fact-stable state classes require:

- *Finite distinguishability* (Theorem 2 / A1): the encoding burden of sub-granular fact formation diverges; no physically realizable process with finite causal reach can meet it.
- *Irreversible commitment* (Theorem 3 / A2): causal locality prevents joint control over the dilation environment; every fact export is causally irreversible.

These are independent necessary conditions — each follows directly from Theorem 1 through a different argument.

Level 3 — Theorem 4 (Second-order closure)

Finite distinguishability ∧ Irreversible commitment require:

- *Finite localization capacity* (Theorem 4 / A3): without this, A1 and A2 are jointly unsatisfiable. A3 is the closure condition that makes the first two consistent.

This hierarchy means:

If you accept that physics is possible, you must accept A1.
If you accept that physics is possible, you must accept A2.
If you accept A1 and A2, you must accept A3.

There is no step that involves a choice. Each step is forced by the one before.

```
Physics possible
        │
        ▼
Lawhood requires stable state classes        [Theorem 1 — proved here]
        │
        ├──────────────────┐
        ▼                  ▼
Finite                 Irreversible
distinguishability     commitment
(A1)                   (A2)
[Theorem 2             [Theorem 3
 — proved here]         — proved here,
                          probabilistic
                          theories]
        │                  │
        └────────┬─────────┘
                 ▼
        Finite localization
        capacity (A3)          [Theorem 4 — proved here]
                 │
                 ▼ (- - - companion papers - - -)
        ℓ_P ,  G  [downstream: companion papers]
                 │
                 ▼ (- - - companion papers - - -)
        α ,  Λ  [downstream: companion papers]
```

Legend: Solid arrows (│, ▼) = entailments proved in this paper. Dashed arrows (- - -) = derivations in companion papers.

This diagram is the entailment structure of the paper. Solid arrows are established by Theorems 1–4. Arrows to ℓ_P, G, α, and Λ represent downstream constitutive steps developed in companion work.

---

7. Representation of the Necessities within the VERSF Framework {#7}

The three necessities proved above are substrate-neutral: they hold in any universe where physics is possible, regardless of what the underlying substrate is. Any concrete theoretical framework must *represent* these necessities in specific mathematical language. VERSF is the framework specifically constructed as the natural realization of these structural necessities — where each axiom is not postulated but recognized as the unavoidable condition it has been shown to be.

A1 is naturally realized by the void-supported distinguishability bound.

Any framework taking A1 seriously must have a lowest-level description in which distinctions below δ_min simply do not exist as physical facts — not as a practical limitation but as a structural feature of the substrate. In VERSF, the void substrate cannot mediate fact formation below ℓ_P. This is not imposed; it is the natural representation of A1 in a framework where the substrate is the medium of all fact formation.

A2 is naturally realized by the Entanglement Ledger.

Any framework taking A2 seriously must track the irreversible export of correlation across the history of a universe. In VERSF, the Entanglement Ledger is a global, monotonically non-decreasing record of all correlation exports since the first fact was formed. It is not an add-on mechanism; it is the natural bookkeeping device for a framework built on A2 as a structural necessity rather than a statistical tendency.

A3 is naturally realized by the maximum substrate tension $\mathcal{T}$_max.

Any framework taking A3 seriously must represent the substrate's finite localization capacity as a measurable physical parameter. In VERSF, $\mathcal{T}$_max is the maximum tension the void can sustain while remaining fact-supporting. It is not a free parameter; it is the constitutive representation of the localization bound whose existence was proved in Theorem 4, and its value determines G and ℓ_P.

Gravity as the infrared expression of A3 — motivation and outlook.

Once A3 is in place as a structural necessity, there is a natural path toward general relativity as the unique long-range geometric theory. The Weinberg–Deser theorem establishes that the only consistent theory of a massless spin-2 field with four properties — Lorentz invariance, universal coupling, stress-energy conservation, and infrared locality — is general relativity. The structural necessities A1–A3 motivate an infrared regime in which all four conditions are naturally supported; the full derivation of each condition from the VERSF framework is developed in the companion emergent-gravity paper. The present paper establishes the axiomatic foundation on which that derivation rests. The bullets below indicate the intended correspondence between the axiomatic structure and the Weinberg–Deser premises; they are not by themselves full derivations.

- *Lorentz invariance* is structurally motivated by A1: if all facts are formed at the same minimum scale, no matter type receives a privileged distinguishability bound.
- *Universal coupling* is structurally motivated by A1 for the same reason: the equivalence principle follows when no matter type receives a structurally privileged distinguishability bound.
- *Stress-energy conservation* is structurally motivated by A2: the Entanglement Ledger conserves total correlation structure, motivating local conservation laws in the continuum limit.
- *Infrared locality* is structurally motivated by A3: no sub-Planckian structure propagates to macroscopic scales, suppressing additional long-range fields.

The claim here is not that GR is derived in this paper. The claim is that A1–A3 provide the structural ground from which the Weinberg–Deser premises emerge naturally, making GR the natural infrared completion of the necessity chain. That completion is treated elsewhere.

VERSF is therefore the framework that follows from taking the unavoidability of A1, A2, and A3 seriously at the level of physical ontology. The framework's structure reflects the structural necessities established in this paper.

A note on α and Λ. The present paper establishes the necessity of the structural conditions. The dimensionless constants — the fine structure constant α and the cosmological constant Λ — are structural constraints that emerge downstream of the axioms, not additional postulates. α is constrained by relational closure: the requirement that every distinction-creating interaction at the Planck scale be paired with a resolving one, which bounds the U(1) gauge degree-of-freedom count. Λ is constrained by global geometric closure: the requirement that the total correlation structure of the universe satisfy a topological self-consistency condition. In both cases the axioms define the *structural form* of the constraint. The numerical completion — deriving $\alpha \approx 1/137$ and $\Lambda \sim 10^{-52}\ \mathrm{m}^{-2}$ from first principles — is the subject of companion papers in the VERSF program. The present paper does not claim to provide those derivations; it provides the foundation that makes them well-posed.

---

8. Discussion {#8}

8.1 Why These Are Necessities, Not Assumptions

The distinction between a necessity and an assumption matters. An assumption is a claim that could be false — it is adopted because it is convenient or plausible, and a different framework might adopt a different assumption. A necessity is a claim whose denial produces a contradiction — it cannot be false in any universe where physics holds.

Theorems 1–4 show that A1, A2, and A3 are necessities in this strict sense. Each theorem proves its target condition by showing that its denial contradicts Theorem 1 or a preceding necessity. A critic who wishes to deny A1, A2, or A3 must identify which step in the relevant proof fails — and in doing so, must accept a universe without reproducible facts, without a causal past, or without any consistent localization scale. None of these is a viable alternative physics. All of them are the absence of physics.

8.2 Relation to Other Approaches

Standard quantum mechanics: does not take A1, A2, or A3 as foundational, but is consistent with all three. The uncertainty principle, decoherence, and the Planck scale each encode one of the necessities in specific dynamical language.

Information-theoretic approaches (Wheeler, Zeilinger; Verlinde's entropic gravity proposal): take information as primitive but do not prove that information-theoretic structure is unavoidable. The present work does.

Emergent gravity (Jacobson, Verlinde): derives gravitational equations from thermodynamic assumptions. The present work shows why those thermodynamic assumptions — specifically the irreversibility of entropy production — are structurally necessary, not optional.

Loop quantum gravity / causal dynamical triangulations: predict discrete spacetime at the Planck scale. Theorem 4 shows why the existence of such a minimum length is a structural necessity, independent of the specific framework predicting it. Whether the minimum scale equals the Planck scale as computed in those frameworks requires the numerical identification deferred to companion papers.

8.3 What This Paper Does Not Claim

This paper proves the necessity of the structural conditions. It does not provide, and does not claim to provide, the following:

- The specific numerical value of $\mathcal{T}$_max (and therefore G and ℓ_P). The existence of a finite localization capacity bound is proved; its numerical value requires additional physical input developed in companion papers.
- A first-principles derivation of $\alpha \approx 1/137$. The relational closure constraint on α is identified and structurally motivated in §7; its numerical completion is the subject of a separate paper.
- A first-principles derivation of $\Lambda \sim 10^{-52}\ m^{-2}$. The global geometric closure condition on Λ is identified; its numerical completion requires the full topological accounting of the universe's correlation structure.
- A proof that VERSF is the *unique* framework realizing A1, A2, and A3. The claim is that VERSF is the *natural* realization — the framework whose structure is forced by the necessities. Whether alternative realizations exist is an open question.

A paper that claimed all of the above would be overreaching. This paper's scope is disciplined: it establishes that the axioms cannot be otherwise. What follows from the axioms — the specific quantitative predictions, the detailed derivation of constants, the full structure of the void substrate — is the work of the broader VERSF program.

---

9. Anticipated Objections and Their Resolution {#9}

9.1 Could a universe with infinite distinguishability still support physics?

Objection. A universe might possess infinitely fine distinguishability while still allowing physical laws to operate. Facts could in principle be encoded with arbitrarily high resolution.

Response. Theorem 2 does not forbid infinite *logical* distinguishability in mathematical descriptions. It forbids infinite *fact-supporting* distinguishability. The distinction is crucial. Logical distinguishability refers to mathematical labels in a formal model; fact-supporting distinguishability refers to distinctions that can be physically recorded and recovered. Any fact-forming operation must occur within a finite spacetime region and finite causal duration (Physical Realizability Lemma). Under these conditions, distinguishing among infinitely many outcomes requires unbounded encoding resources — which no physically realizable process can supply. Without fact-supporting distinctions, no reproducible state space exists and lawhood collapses (Theorem 1). Infinite distinguishability is compatible with mathematics but incompatible with physics.

9.2 Could deterministic classical physics avoid irreversible commitment?

Objection. In a purely deterministic classical universe, measurements need not export correlation irreversibly. If the full microstate were known, no irreversibility would arise.

Response. This objection assumes observers with unlimited access to the global microstate. Theorem 3 instead concerns finite observers, which are already required by Theorem 1's condition that facts be accessible. A finite observer cannot access the complete microstate of a deterministic universe. Their effective description of local processes is therefore probabilistic: inaccessible environmental variables function as hidden dilation variables. Deterministic mechanics with finite observers is formally equivalent to a local probabilistic theory with inaccessible dilation variables (§4.2), placing it within the scope of Theorem 3. Irreversible commitment follows from causal boundedness and incomplete access — not from quantum indeterminacy alone.

9.3 Could the minimum correlation commitment shrink with scale?

Objection. One might attempt to evade Theorem 4 by allowing ΔS_{min} to decrease as localization scale shrinks, keeping the correlation export per unit volume bounded.

Response. The Scale-Independence Lemma (§5.4) shows that $\Delta S_{min} \to 0$ collapses either distinguishability or irreversible commitment. If $\Delta S_{min} \to 0$ while distinctions remain meaningful, irreversible commitment disappears, violating A2. If $\Delta S_{min} \to 0$ because distinctions leave no environmental trace, distinguishability collapses, violating A1. In either branch the attempt to avoid the conclusion violates a necessity already established. ΔS_{min} therefore has a nonzero lower bound $\Delta S_0 > 0$ independent of localization scale.

9.4 Could correlation export remain finite if fact density decreases sufficiently quickly?

Objection. If fact density falls rapidly enough as localization scale shrinks, the divergence in correlation export per unit volume might be avoided.

Response. Even if fact density decreases, the correlation export from each remaining fact must be absorbed by a local environment of volume $\sim\ell^3$, producing an absorption demand of $\Delta S_0/\ell^3$ per event. Since $\Delta S_0 > 0$ by the Scale-Independence Lemma, this demand diverges as $\ell \to 0$ regardless of how rapidly fact density falls. The Finite Absorption Capacity Lemma (§5.5) then guarantees that no finite receiving region can meet a divergent absorption demand, producing a contradiction with Theorem 3. There is no rate of density decrease that rescues the conclusion.

9.5 Are the VERSF axioms unique?

Objection. Even if A1–A3 are necessary conditions for physics, many theoretical frameworks might satisfy them. The necessity of the axioms does not imply the necessity of VERSF.

Response. Correct — and the paper does not claim otherwise. The claim is narrower: A1–A3 are structural necessities for any universe capable of supporting law-governed physical facts accessible to finite observers. VERSF is presented as a framework explicitly constructed to represent these necessities in a unified and coherent way. Whether other frameworks can satisfy the same structural conditions equally naturally is an open question this paper deliberately leaves open. The burden the paper places on competitors is not uniqueness but adequacy: any alternative framework must show how it accommodates the necessities established by Theorems 1–4.

---

10. Conclusion {#10}

The Void Energy-Regulated Space Framework rests on three axioms: finite distinguishability (A1), irreversible commitment (A2), and finite localization capacity (A3). This paper has proved that these axioms are not postulates. Each one is a structural necessity — unavoidable in any universe where physical laws can exist.

The argument proceeds through four theorems:

- **Theorem 1** establishes the primitive: physical laws require fact-stable state classes. This is not a philosophical claim but a formal precondition for any law to act on any domain.
- **Theorem 2** proves A1: the information cost of sub-granular fact formation diverges, preventing any finite physical record from sustaining arbitrarily fine distinctions.
- **Theorem 3** proves A2: relativistic causal locality prevents any locally-bounded agent from accessing the dilation environment of a local operation; correlation export is therefore irreversible by causal structure, not thermodynamic accident.
- **Theorem 4** proves A3 from A1 $\wedge$ A2: without a finite localization capacity, the first two necessities are jointly unsatisfiable. A3 is the closure condition that makes them consistent.

Each theorem is stated in substrate-neutral language and holds independently of VERSF-specific machinery. VERSF enters as the natural theoretical framework for representing these necessities: the void substrate's distinguishability bound realizes A1; the Entanglement Ledger realizes A2; the maximum tension $\mathcal{T}$_max realizes A3.

The implications are significant. The VERSF axioms are not assumptions about the particular universe we inhabit. They are conditions on any universe that contains physics. A framework built on them is not describing one possibility among many. It is describing the structural ground that any physical theory must occupy.

The burden on any competing foundational framework is therefore not merely to postulate alternatives, but to show how physical law remains possible while denying one or more of Theorems 1–4. A critic must do more than reject one of the axioms; they must show how finite observers can possess law-governed, recoverable facts without finite distinguishability, without irreversible registration, or without a finite localization bound. Until that is shown, the VERSF axioms stand not as the assumptions of one framework but as the unavoidable conditions of any framework.

Appendix A: Formal Theorem Summary

Primitive (P)
Physics is possible in a universe U: there exist physical laws L: $\mathcal{S} \times \mathcal{T} \to \mathcal{S}$ with non-trivial domains.

Theorem 1
P ⟹ U supports fact-stable state classes: stable, recoverable, reproducible distinctions forming the domain of L.

Theorem 2 (A1)
Theorem 1 ⟹ $\exists\ \delta_{min} > 0$ such that distinctions below δ_{min} cannot be stabilized as physical facts.
Proof route: $I(\ell) \sim \log(V/\ell^3) \to \infty$ as $\ell \to 0$; no physically realizable process with finite causal reach and finite operational resources can meet this encoding burden (Physical Realizability Lemma, §3.3).

Lemma (Causal Boundedness)
Relativistic causality ⟹ no agent at P controls degrees of freedom outside $J^-(P)$.

Theorem 3 (A2)
Theorem 1 + Causal Boundedness ⟹ every local fact-forming operation in a local probabilistic theory deposits a causally irreversible correlation residue in inaccessible environment variables; in quantum theory this is realized by CPTP maps and a monotone Entanglement Ledger.

Lemma A
Unbounded localization + preserved fact density ⟹ violation of Theorem 2.

Lemma B
Unbounded localization ⟹ divergent $\Delta S_{vol} \to \infty$, exceeding the finite correlation-absorption capacity of any finite receiving region; violation of Theorem 3.

Theorem 4 (A3)
Theorem 2 ∧ Theorem 3 ⟹ $\exists\ \ell^* > 0$ (finite localization capacity, equivalently a finite upper bound on stabilizable distinction density). In the broader VERSF program, the finite localization-

capacity scale is later identified with a gravitational tension scale $\mathcal{T}_{max}$; this identification is a downstream constitutive step, not a premise of the present axioms paper.

Appendix B: The Entanglement Ledger — Formal Statement

For a system S and environment E under global unitary evolution U, define:

$\mathscr{L}(S, E) = S(\rho_S) + S(\rho_E) - S(\rho_{SE}) \geq 0$

where $S(\rho)$ is von Neumann entropy. By subadditivity of von Neumann entropy, $\mathscr{L} \geq 0$. Under global unitary evolution, $S(\rho_{SE})$ is constant. For any locally-bounded agent performing a CPTP operation on S alone, $S(\rho_E)$ can only increase:

$d\mathscr{L}/dt \geq 0$

Each fact-formation event increments $\mathscr{L}$ by at least $k_B \ln 2$ per bit of fact formed. The Ledger accumulates monotonically over the history of the universe.

The Ledger is the VERSF representation of A2. Its monotone growth is not a consequence of the second law — it is the mechanism underlying the second law. The second law is the ensemble-level statistical manifestation of the per-event structural necessity proved in Theorem 3.

Appendix C: Mathematical Notation

The following logical objects are used throughout the paper. Definitions are given here for reference; each is introduced at first use in the main text.

Symbol	Definition
$\mathcal{C}$	Configuration space — the space of all microscopic configurations of a physical system
~	Perturbative equivalence relation on $\mathcal{C}$: $s \sim s'$ iff $d(s, s') < \delta$ for some $\delta > 0$
$\mathcal{S} = \mathcal{C}/\sim$	Physical state space — the quotient of configuration space under ~
[s]	Equivalence class of configuration s: $[s] = \{ s' : d(s, s') < \delta \}$
δ_{min}	Minimum resolvable difference — the lower bound on fact-supporting distinguishability (A1)
$N(\lambda)$	Number of distinguishable outcomes of a fact-forming operation at refinement level λ

Symbol	Definition
ΔS_min	Minimum correlation commitment required to register a fact as a new distinct physical event
ΔS_0	Scale-independent lower bound on ΔS_min: $\Delta S_{min} \geq \Delta S_0 > 0$ for all ℓ (Scale-Independence Lemma)
ρ_f(ℓ)	Stabilizable distinction density — the number of distinguishable, stably-recorded facts per unit volume at scale ℓ
N_max(ℓ)	Maximum number of distinguishable fact-sites in a region of scale ℓ: $N_{max}(\ell) \lesssim (\ell/\delta_{min})^3$
$\mathscr{L}$(S,E)	Entanglement Ledger — total correlation exported by fact-formation events: $\mathscr{L}(S,E) = S(\rho_S) + S(\rho_E) - S(\rho_{SE})$
$\mathcal{T}(\ell)$	Localization tension — energy cost per unit area of sustaining a fact-supporting boundary at scale ℓ
$\mathcal{T}$_max	Maximum substrate tension — the finite upper bound on localization tension (A3, VERSF representation)
ℓ*	Minimum localization scale — the smallest region in which a fact can be formed: $\ell^* > 0$
$J^-(P)$	Causal past of spacetime point P

References

Bekenstein, J.D. (1973). Black holes and entropy. *Physical Review D*, 7(8), 2333–2346.

Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183–191.

Shannon, C.E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423.

Stinespring, W.F. (1955). Positive functions on C*-algebras. *Proceedings of the American Mathematical Society*, 6(2), 211–216.

Weinberg, S. (1964). Photons and gravitons in S-matrix theory. *Physical Review*, 135(4B), B1049–B1056.

Wheeler, J.A. (1989). Information, physics, quantum: The search for links. In *Proceedings of the 3rd International Symposium on Foundations of Quantum Mechanics*, Tokyo.

Zeilinger, A. (1999). A foundational principle for quantum mechanics. *Foundations of Physics*, 29(4), 631–643.