Testing the Mathematics: The Speed of Light as a Computational Throughput Limit

VERSF Theoretical Physics Program

Abstract for the General Reader

Why can't anything travel faster than light? Physics textbooks say it's a law of nature, but they rarely explain why *this particular speed—299,792,458 meters per second—is the universe's speed limit.*

This paper explores a radical idea: the speed of light isn't an arbitrary cosmic rule. It's the maximum rate at which the universe can process information and create stable physical facts. Think of it like a computer's processing speed—there's only so fast the hardware can go before errors accumulate and everything breaks down.

We show that this interpretation is mathematically consistent: if the speed of light really is a computational limit, then gravity's strength, the size of the smallest possible things (the Planck scale), and Einstein's relativity all follow as natural consequences.

The key insight: a stable "information packet" (what we call a fold) must be large enough to avoid quantum collapse at tiny scales, yet small enough to maintain coherence across cosmic distances. The optimal size turns out to be the geometric mean of the smallest and largest scales in physics—a proposed mesoscopic "Two-Planck" coherence window at order 10^{-4} m (tens of micrometers).

Given a dimensionless closure ratio fixing electromagnetic fold structure (the fine-structure constant α) and the requirement that physical facts be irreversible, the maximal information-propagation channel is fixed. The speed of light is the embodied expression of that channel.

Executive Summary: What This Paper Proves

1. A speed limit isn't optional.

If reality is made of stable facts (irreversible "yes/no" distinctions), then there must be a maximum rate at which those facts can influence other facts. If influence could spread arbitrarily fast, you'd get contradictions: different parts of reality would disagree about what's true "first," and facts couldn't stay stable.

Result: A hard upper limit on causal propagation must exist. That's the deep reason there is something like a speed of light at all.

2. Light hits the limit because it's the cheapest possible "information package."

A single isolated "bit" can't really travel; it smears out and loses identity. What travels is a *fold*: a minimal, closed, self-consistent packet of information. The electromagnetic packet (light) is the most efficient possible fold—it requires the fewest irreversible "commitments" to exist and remain stable while propagating.

Result: Light defines the maximum speed not because photons are special by fiat, but because EM is the minimal stable carrier of correlation.

3. Gravity and the Planck scale aren't independent of the speed limit.

Once you accept that facts cost action ($\hbar$), there is a maximum substrate tension, and the fastest channel exists, then gravity's strength (G) and the smallest meaningful scale (ℓ^*, numerically matching the Planck length) can't be chosen independently. They are tied together by the same closure logic.

Result: If the universe can only update so fast, gravity must take a corresponding strength so facts don't tear apart.

4. The "Einstein rules" follow once the speed limit is real.

Special relativity is Einstein's 1905 discovery that space and time are not absolute—they stretch and compress depending on how fast you're moving. Specifically:

- **Time dilation:** Moving clocks tick slower. A clock on a speeding spaceship runs slow compared to one on Earth.
- **Length contraction:** Moving rulers shrink. A spaceship traveling near light speed is physically shorter (in its direction of motion) as measured by a stationary observer.
- **Relativity of simultaneity:** Two events that happen "at the same time" for one observer may happen at different times for another.

These effects seem bizarre, but they are *required* if every observer must agree on the same ultimate causal limit. The only consistent way to relate different observers' measurements is the Lorentz transformation. Time dilation and length contraction are the bookkeeping rules required so all observers agree on what can and cannot become a stable fact.

VERSF reinterpretation (see Section 20): Space and time don't actually stretch—what changes is the "frame rate" at which irreversible facts are produced. Moving clocks tick slower because motion diverts update capacity away from internal processes. Length contraction occurs because fewer correlation layers can be maintained along the direction of motion. The Lorentz factor γ is a throughput reallocation factor. Same math, different ontology.

5. A specific "middle scale" (~80 micrometers) emerges from balancing UV and IR constraints.

- Too small → unstable due to UV (identity-collapse) limits
- Too large → loses coherence due to IR (cosmological) limits

Balancing these gives: $\xi_{meso} \sim \sqrt{(\ell^* \cdot L_{IR})} \approx 80$ μm.

Status: The numerical check uses standard inferred scales, so it's best described as strong compatibility plus a concrete target scale to measure independently.

6. The paper proves a closure rule, not "c from nothing."

A closure relationship is proven: if someone measures the mesoscopic coherence scale independently, then c is no longer free—it's fixed by that measurement plus ħ, G, and L_{IR}. If the mesoscopic scale is computed using relations that already include c, then the equation becomes a self-consistency loop (which we explicitly acknowledge).

Proved: c is structurally constrained by closure relations. *Not proved (yet):* A totally independent numerical calculation of c from first principles alone.

7. The "circularity" concern is addressed: meters and seconds are emergent.

Meters and seconds are not fundamental—they are labels attached to stable patterns of irreversible events. The "speed of light" is the ratio of those emergent calibrations. That's why talking about "meters per hop" can look circular if you forget that meters are reconstructed from the bit-stack.

One-sentence summary:

If reality is built from irreversible facts, then a universal speed limit must exist; electromagnetism is the most efficient fact-carrying channel and therefore sets that limit; and once that limit exists, gravity's strength, the smallest meaningful scale, and special relativity follow as consistency requirements—with a predicted mesoscopic coherence scale around ~80 micrometers that can (in principle) be measured independently.

Table of Contents

Executive Summary

- What This Paper Proves

Part I: Foundations

1. Scope and Status of Results
2. Structural Premises
3. Notation Conventions

Part II: Structural Derivation (Track A) 4. Maximum Tension from Information-Theoretic Constraints 5. Theorem 1: Throughput–Geometry Coupling 6. Planck-Length Consistency Check 7. Core Mathematical Test: Does Faster c_T Imply Stronger Gravity? 8. Interpretation: c as a Processing Throughput Limit 9. What Is Proven and What Remains Open

Part III: Bit-Counting Formalism 10. c as Bits per Depth Displacement 11. Speed as Displacement per Stabilized Bit 12. Frames, Not Seconds: Motion as Bit-Stacking

Part IV: Fold Formalism 13. Folds as the Missing Layer 14. Lemma: Five-Commitment EM Closure

Part V: Fixed-Point Derivation (Track B) 15. Theorem 2: Fixed-Point Constraint on Substrate Coherence 16. Theorem 3: UV/IR Stability Determines Mesoscopic Coherence 17. Closure Relation for the Speed of Light 18. Emergence of Metric Units from Causal Ordering

Part VI: Consequences and Consistency 19. Corollary: Propagation Cone Equality 20. Theorem 4: EM Folds Are Maximally Efficient 21. Lorentz Structure and the VERSF Reinterpretation of Special Relativity 22. Testable Prediction: Generation–Propagation Asymmetry

Part VII: Discussion 23. Objections and Resolutions 24. Relation to Existing Work 25. Falsifiability and Empirical Handles 26. Conclusion 27. References

Appendices A. Empirical Anchors for ξ_meso B. Robustness of Theorem 3 C. Clarifying the Reduced-Constants Claim D. Summary of Empirical Handles

---

Scope and Status of Results

This paper establishes three levels of results:

Track A (Structural consistency): We show that interpreting c as a maximal irreversible information-throughput bound is internally consistent with gravitational coupling, Planck-scale relations, and relativistic kinematics.

Track B (Fixed-point constraints): We derive two fixed-point constraints:

- **Theorem 2:** $\xi_{UV} = N_b \cdot \ell^*$ (substrate fold coherence)
- **Theorem 3:** $\xi_{meso} = \sqrt{\ell^* \cdot L_{IR}}$ (mesoscopic coherence as geometric mean of UV and IR scales)

Track C (Closure relation for c): We derive a closure relation:

$c = (L_{IR}^2 \cdot \hbar G / \xi_{meso}^4)^{(1/3)}$ ... (16.5)

Epistemic status of this result:

If ξ_{meso} is...	Then equation (16.5)...
Measured independently (~88 μm)	Determines c from (ħ, G, L_{IR}) — a **prediction**
Derived from $\ell^*(c)$	Yields c = c — a **self-consistency check**

The reduced-constants claim: In standard physics, (ħ, G, c, Λ) are four independent constants. This framework proposes that they are not freely specifiable simultaneously once the coherence closure is imposed. Whether this reduces the count of independent constants depends on whether ξ_{meso} is independently fixed.

Operationally, the closure becomes predictive only to the extent that ξ_{meso} can be fixed independently of (ħ, G, c).

---

1. Structural Premises

For the general reader: This section lists our starting assumptions. We assume the universe has a smallest meaningful size (UV cutoff), a largest coherent scale (IR cutoff), and information-processing limits from quantum mechanics.

We assume:

1. **Finite distinguishability (FD):** Minimum separable state difference enforces UV cutoff ℓ^*.

2. **Irreversible commitment (IC):** Stabilizing any fact requires nonzero action $\sim \hbar$.
3. **Causal locality:** Commitments act only over bounded neighborhoods per update.
4. **Substrate saturation (SAT):** Maximum restoring tension $\mathcal{T}_{max}$ exists.
5. **IR closure:** Maximum coherent scale L_{IR} exists (cosmological horizon / Λ-boundary).
6. **Far from saturation**, geometric response reproduces inverse-square law.

The main results:

- (Theorem 1) c_T, G, and ℓ^* are not independent.
- (Theorem 2) $\xi_{UV} = N_b \cdot \ell^*$ (substrate coherence).
- (Theorem 3) $\xi_{meso} = \sqrt{\ell^* \cdot L_{IR}}$ (mesoscopic coherence).
- (Theorem 4) EM folds saturate the throughput bound.

---

2. Notation Conventions

Symbol	Meaning
≡	Definition
=	Equality
≈	Numerical approximation
~	Scaling relation
∝	Proportionality

Key symbols:

Symbol	Definition
ℓ^*	Identity-collapse length (UV cutoff); $\sim \ell_P$
ℓ_P	Standard Planck length; $\ell_P \equiv \sqrt{\hbar G/c^3}$
L_{IR}	Cosmological coherence scale (IR cutoff)
ξ_{UV}	Substrate fold coherence; $\xi_{UV} = N_b \cdot \ell^*$ (Planck-scale)
ξ_{meso}	Mesoscopic coherence; $\xi_{meso} \sim \sqrt{\ell^* \cdot L_{IR}}$ ($\sim 10^{-4}$ m)
ν^*	Substrate-limited commitment rate; $\nu^* \equiv \mathcal{T}_{max}\, \ell^* / \hbar$
N_b	Commitments per fold ($N_b = 5$ for EM)
$\mathcal{T}_{max}$	Maximum restoring tension
c_T	Causal cone speed
c	Speed of light; $c = c_T$ for maximal carriers

*Note on ℓ vs ℓ_P:** The identity-collapse scale ℓ^* is inferred from finite distinguishability requirements, not defined using c. The standard Planck length ℓ_P is defined from measured (ħ, G, c). After Theorem 4 establishes $c = c_T$, these become numerically coincident.

Convention for L_IR: Unless stated otherwise, we take $L\_IR \propto \Lambda^{-1/2}$ as the non-circular IR coherence scale; $L\_IR \sim c/H_0$ is used only as a cosmological order-of-magnitude proxy.

Definition (Two-Planck Window): The "Two-Planck window" refers to the mesoscopic coherence scale $\xi\_meso \sim \sqrt{(\ell\_P \cdot L\_IR)}$, which involves both the Planck length ℓ_P (UV scale) and the cosmological scale L_IR (IR scale). The name reflects its dependence on *both* Planck-scale and cosmological-scale physics. Numerically, this gives ξ_meso ~ 30–100 μm. This term is introduced in this paper; the ~88 μm value used in calculations corresponds to $\xi\_meso = \sqrt{(1.6 \times 10^{-35}\ \text{m} \times 4.4 \times 10^{26}\ \text{m})}$.

---

3. Maximum Tension from Information-Theoretic Constraints

From causal closure and the Margolus–Levitin bound [2], stabilizing one bit within a cell of size ℓ* over one causal tick $\tau \sim \ell^*/c\_T$ requires minimum energy $E \sim \hbar c\_T/\ell^*$.

Dividing by ℓ* yields maximum restoring tension:

$$\mathcal{T}_{max} \sim \hbar c\_T / \ell^{*2}$$

---

4. Theorem 1: Throughput–Geometry Coupling

Lemma 1.0 (Uniqueness of the IR Coupling Under Saturation)

Assume:

1. **IR linearity:** Far from saturation, the acceleration response to a localized mass M at distance r is inverse-square: $a \propto M/r^2$.
2. **Universality:** The coupling constant does not depend on source composition or test body.
3. **Saturation dominance:** The dominant UV/IR bridging parameter controlling the maximum linear-response strength is the substrate saturation tension $\mathcal{T}_{max}$ (units of force). Other potential scalars (e.g., matter-specific coupling constants) are either absorbed into $\mathcal{T}_{max}$ or contribute only subdominant corrections.
4. **Relativistic throughput ceiling:** The causal cone speed c_T is an invariant conversion constant between temporal update ordering and spatial correlation propagation.

Then: The unique form of the IR coupling constant is:

$$G \equiv c\_T^4 / \mathcal{T}_{max}$$

up to an order-unity dimensionless factor that can be absorbed into the definition of $\mathcal{T}_{max}$.

Proof (outline). In a linear inverse-square law $a = \gamma M/r^2$, the coupling γ must have dimensions $[L^3 M^{-1} T^{-2}] = [G]$. The dominant dimensional quantities allowed by assumptions (3–4) are c_T and $\mathcal{T}_{max}$. The unique combination of these with dimensions of G is $c_T^4/\mathcal{T}_{max}$. Any alternative would require introducing an additional independent dimensional scale at leading order, violating assumption (3). □

Theorem 1 (Throughput–Geometry Coupling).

Assume (i) FD, (ii) IC, (iii) SAT, (iv) linear geometric response in IR.

Then:

$$\mathbf{G \equiv c_T^4 / \mathcal{T}_{max}}$$

$$\boldsymbol{\ell^{*2} = \hbar G / c_T^3}$$

Proof. By Lemma 1.0, $G = c_T^4/\mathcal{T}_{max}$ is the unique IR coupling under saturation control. Substituting $\mathcal{T}_{max} \sim \hbar c_T/\ell^{*2}$ yields:

$$G = c_T^4 \cdot \ell^2 / (\hbar c_T) = c_T^3 \ell^2 / \hbar$$

Solving for ℓ^2: **$\ell^2 = \hbar G/c_T^3$** □

Remark (effective-action view). In the IR linear regime, the most general local scalar functional coupling a matter source ρ to a response potential Φ has the schematic form $S_{eff} \sim \int d^3x\, [(\nabla\Phi)^2/(8\pi\gamma) + \rho\Phi]$. Variation gives $\nabla^2\Phi = 4\pi\gamma\rho$. The identification $\gamma = c_T^4/\mathcal{T}_{max}$ fixes the coefficient of the gradient term in terms of the substrate saturation parameter, making G an emergent elasticity constant of the void response rather than a free coupling.

---

5. Planck-Length Consistency Check

The standard Planck length is defined as:

$$\boldsymbol{\ell_P \equiv \sqrt{\hbar G/c^3} \approx 1.6 \times 10^{-35}\ \mathrm{m}}$$

In this framework, the identity-collapse scale obeys $\ell^2 \sim \hbar G/c_T^3$. *Identifying $c_T = c$ for the maximally efficient carrier (Theorem 4) makes ℓ* numerically coincident with ℓ_P.

---

6. Core Mathematical Test: Does Faster c_T Imply Stronger Gravity?

$G \propto c_T^3$ (at fixed ℓ^*)

Clarification: This scaling holds when the identity-collapse scale ℓ^* is treated as fixed by substrate physics independent of c_T. From Theorem 1, $G = c_T^4/\mathcal{T}_{max}$ with $\mathcal{T}_{max} \sim \hbar c_T/\ell^2$, *giving* $G \sim c_T^3\ell^2/\hbar$. If ℓ^* is held fixed while c_T varies, then $G \propto c_T^3$.

Note that $\ell^2 = \hbar G/c_T^3$ *(the Planck-length relation) is a consistency condition, not an independent equation. It expresses how the quantities must relate once the substrate is specified. One cannot freely vary c_T while holding both ℓ* and G fixed—they are coupled through closure.

The physical interpretation: if the universe had a faster causal limit c_T (with the same distinguishability scale ℓ^*), gravity would be correspondingly stronger. Doubling c_T requires gravity to become 8× stronger.

---

7. Interpretation: c as a Processing Throughput Limit

The speed of light is:

c = (adjacency hops) / (ticks)

Massless excitations saturate this bound; massive systems cannot.

---

8. What Is Proven and What Remains Open

Proven:

- c_T, G, ℓ^* not independent (Theorem 1)
- $\xi_{UV} = N_b \cdot \ell^*$ (Theorem 2)
- $\xi_{meso} = \sqrt{(\ell^* \cdot L_{IR})}$ (Theorem 3)
- $c = (L_{IR}^2 \cdot \hbar G / \xi_{meso}^4)^{(1/3)}$ (structural relationship)

Open:

- Independent measurement of ξ_{meso} to verify ~30–100 μm prediction
- Precision cosmological determination of L_{IR}

---

9. c as Bits per Depth Displacement (Tick–Bit Formalization)

Primitive counts:

Symbol	Meaning
N_tick	Commitment events
N_bit	Stabilized bits
N_hop	Adjacency hops

Dimensionless speed:

$\boldsymbol{\kappa \equiv \Delta D / \Delta B}$ (hops per bit)

Theorem (Bounded Throughput). $\kappa \le \kappa_{max} \sim 1/\beta^*$. □

---

10. Speed as Displacement per Stabilized Bit (Film-Frame Analogy)

$\boldsymbol{v \equiv \Delta x / N_{bit}}$

The maximal channel satisfies $\kappa_{bit} = 1$.

---

11. Frames, Not Seconds: Motion as Bit-Stacking

$\boldsymbol{c = \kappa \cdot N_{sec}}$

---

12. Folds as the Missing Layer: Why Only Folded Bits Propagate

Definition (Fold, Formal)

A **fold** is a finite subgraph F of the commitment adjacency graph such that:

1. **(Closure)** F has no open boundary under the adjacency operator (topological closure).

2. **(Transportability)** There exists a morphism T mapping F to an adjacent subgraph F′ such that distinguishability is preserved: D(F) = D(F′) where D is the stable-distinction functional.

Definition (Stable-Distinction Functional): D(F) is the number of independent binary distinctions stably encoded in configuration F. Formally, $D(F) = \log_2|S\_F|$ where S_F is the set of distinguishable internal states of F that persist under small perturbations. Two configurations F and F′ have equal D if and only if they encode the same information content.

3. **(Stabilizability)** The transport map T requires a finite number N_b of irreversible commitments and no additional external commitments scale with path length (universality).
4. **(Gauge-consistency)** The quotient space of fold states by gauge equivalence contains only physical degrees of freedom under transport (no unphysical modes propagate).

A **propagating fold species** is a family of folds closed under repeated application of T.

Lemma. Isolated bits decohere within O(1) hops. □

Corollary. Only folds (satisfying conditions 1–4) propagate causally. □

---

13. Theorem 2.5: Classification of Minimally Propagating Gauge Folds

Definitions

A **propagating gauge fold** is a closed configuration that:

- **(C1)** Transports a distinguishable state across adjacency without loss
- **(C2)** Is gauge-consistent (no unphysical degrees of freedom propagate)
- **(C3)** Supports universal, medium-independent propagation
- **(C4)** Carries at least one binary physical degree of freedom (to encode polarization class)
- **(C5)** Admits a causal transport rule that is invariant under inertial changes (cone preservation)

An **irreversible commitment** is a constraint satisfaction event that reduces admissible fold states and cannot be undone without entropy export.

Theorem 2.5 (N_b = 5 Classification)

Claim: Any propagating U(1) gauge fold requires at least five independent irreversible commitments, and there exists a U(1) fold achieving exactly five.

Proof (Lower bound ≥ 5):

Requirement	Commitment	Constrained DOF
(C1) Closure of connectivity	Topological closure	Connectivity
(C2) U(1) phase consistency	Gauge loop closure	Phase
(C2) Removing unphysical modes	Physical-mode projection	Mode content
(C4) Binary physical class selection	Polarization commitment	Class
(C3,C5) Universal inertial invariance	Causal transport closure	Transport law

Each requirement constrains a distinct degree of freedom, so the commitments are independent. □

Proof (Achievability = 5):

The electromagnetic fold constructed with these five constraints satisfies C1–C5 and propagates at the universal bound.

Why no sixth constraint is needed: Any candidate sixth constraint must either:

- (a) Be derivable from C1–C5 (hence redundant), or
- (b) Impose an additional restriction that violates one of the propagation requirements.

For example:

- *Chirality constraints* beyond polarization (C4) would over-specify the mode content, violating the binary physical DOF requirement.
- *Additional gauge constraints* beyond U(1) closure would introduce unphysical degrees of freedom or destroy universality (C3).
- *Metric-dependent transport rules* would violate medium-independence (C3) and inertial invariance (C5).

Once a fold is (1) topologically closed, (2) gauge-consistent, (3) physical-mode projected, (4) polarization-committed, and (5) transport-universal, all propagation requirements C1–C5 are satisfied. The five constraints are therefore both necessary and sufficient. □

Conclusion:

N_b^(γ) = 5

#	Commitment	Constrains	Failure if Removed
1	Topological Closure	Connectivity	Fragmentation
2	Gauge Closure (U(1))	Phase	Leakage
3	Physical Mode Projection	DOF	Artifacts
4	Polarization Commitment	Class	Ambiguity

#	Commitment	Constrains	Failure if Removed
5	Causal Transport Closure	Universality	Medium-dependence

---

14. Theorem 2: Fixed-Point Constraint on Substrate Coherence

Theorem 2 (Substrate Coherence Fixed Point).

Assume (i)–(vi) from Section 1.

Then:

$\xi_{UV} = N_b \cdot \ell$

Proof.

Substrate-limited commitment rate: $\nu^* = \mathcal{T}_{max}\, \ell^* / \hbar$

Fold-step length: $\ell_{hop} = \xi_{UV} / N_b$

Throughput: $c_T = \ell_{hop} \cdot \nu^* = (\xi_{UV}/N_b) \cdot (\mathcal{T}_{max}\, \ell^*/\hbar)$

Substituting $\mathcal{T}_{max} \sim \hbar c_T/\ell^{*2}$:

$c_T = c_T \cdot \xi_{UV} / (N_b \cdot \ell^*)$

Dividing by c_T: $\xi_{UV} = N_b \cdot \ell^*$ □

Corollary 2.1. No explicit c_T in the constraint (non-circular). □

Corollary 2.2. The fixed point is an attractor. □

---

15. Theorem 3: UV/IR Stability Determines Mesoscopic Coherence

For the general reader: A stable information packet must be large enough to avoid quantum collapse (UV failure) yet small enough to maintain cosmic coherence (IR failure). The optimal size is the geometric mean of the smallest and largest scales in the universe.

Theorem 3 (UV/IR Geometric Mean).

Assume:

- (i) FD/IC enforce UV cutoff ℓ^*
- (ii) There exists an IR coherence scale L_{IR} (cosmological horizon or Λ-boundary)
- (iii) Stable propagation requires robustness against both UV and IR failure modes

Then: The mesoscopic coherence window scales as:

$$\xi_{meso} \sim \sqrt{\ell \cdot L_{IR}}$$

Lemma 3.0 (First-Order Fragility Scalings)

Let ℓ be the coherence size of a candidate fold.

UV fragility:

Finite distinguishability implies there exists a minimum stable boundary thickness ℓ^*. The dominant UV failure mode is *boundary leakage*: an open or insufficiently thick boundary leaks distinguishability into neighbors at a rate proportional to the fraction of the structure occupied by unstable boundary. To first order, this fraction scales as:

$$\varepsilon_{UV}(\ell) \propto \text{(boundary thickness)/(structure size)} \sim \ell/\ell^*$$

Higher powers $(\ell^*/\ell)^p$ correspond to multi-step leakage models; the first-order term is the minimal monotone scaling.

IR fragility:

Global coherence failure arises when the fold size becomes comparable to the IR closure scale, producing *closure mismatch* with the background coherence frame. For $\ell \ll L_{IR}$, the leading-order mismatch is linear in ℓ/L_{IR} by smoothness:

$$\boldsymbol{\varepsilon_{IR}(\ell) \propto \ell/L_{IR} + O((\ell/L_{IR})^2)}$$

Thus the minimal IR fragility scaling is linear, with higher powers representing subleading corrections.

Conclusion: The choices ℓ^*/ℓ and ℓ/L_{IR} represent the leading-order terms in the small-parameter expansions controlling UV boundary leakage and IR closure mismatch. □

Proof of Theorem 3

UV failure mode. If a fold's coherence length ℓ is too small, it is vulnerable to identity collapse. By Lemma 3.0, the dimensionless UV fragility scales as:

$\varepsilon_{UV}(\ell) \sim \ell/\ell^*$

IR failure mode. If ℓ is too large, the fold cannot maintain coherence across the cosmological substrate. By Lemma 3.0, the dimensionless IR fragility scales as:

$\varepsilon_{IR}(\ell) \sim \ell/L_{IR}$

Total instability functional:

$\varepsilon_{tot}(\ell) = \ell/\ell + \ell/L_{IR}*$

Minimization. Setting $d\varepsilon_{tot}/d\ell = 0$:

$-\ell/\ell^2 + 1/L_{IR} = 0*$

$\ell^2 = \ell \cdot L_{IR}*$

$\ell = \sqrt{(\ell \cdot L_{IR})}*$

We therefore obtain the scaling law $\xi_{meso} \sim \sqrt{(\ell \cdot L_{IR})}*$ (up to order-unity factors). □

Numerical Compatibility Check

Empirical scales:

Using the empirically inferred identity-collapse scale ℓ^* (numerically coincident with ℓ_P when evaluated using measured ħ, G, c) and the cosmological horizon:

- $\ell^* \sim \ell_P \sim 1.6 \times 10^{-35}$ m
- $L_{IR} \sim c/H_0 \sim 4.4 \times 10^{26}$ m

Predicted mesoscopic coherence:

$\xi_{meso} \sim \sqrt{(1.6 \times 10^{-35} \times 4.4 \times 10^{26})} = \sqrt{(7 \times 10^{-9})} \approx 8 \times 10^{-5}$ m

$\xi_{meso} \approx 80$ μm

This is compatible with the proposed Two-Planck window scale (~88 μm).

Epistemic note: This numerical check uses standard values that depend on measured c. The framework is therefore *compatible* with the correct order of magnitude, but the prediction would become fully independent only if ℓ^* were derived from distinguishability closure without reference to c.

Cosmological Connection

$\xi_{meso} \propto \sqrt{L_{IR}}$

If the universe had a different cosmological constant (different L_{IR}), the coherence window would differ. This is testable across cosmologies.

---

16. Closure Relation for the Speed of Light

For the general reader: This section explains precisely what is—and is not—being derived. We do not "pull the speed of light out of thin air." Instead, we show that once certain independently motivated structural scales are fixed, the value of c is no longer free. The result is a closure relation: given three of the four quantities (ħ, G, c, Λ), the fourth is determined.

16.1 The Throughput Closure Equation

In the VERSF framework, the speed of light is interpreted as the maximal rate at which stabilized distinctions (facts) can propagate across the void substrate. For the maximally efficient carrier (the electromagnetic fold), the embodied speed takes the form

$c \equiv \ell_{hop} \cdot \nu^*$

where ℓ_{hop} is the emergent metric representation of one causal hop (a projection of correlation depth into spatial distance), and ν^* is the substrate-limited commitment rate.

For the maximally efficient fold, we take the transport step to realize the minimal hop embodiment:

$\ell_{hop} \sim \ell^*$

This is the statement that the EM fold saturates the minimal distinguishability scale in its transport step.

The substrate-limited commitment rate is:

$\nu \equiv \mathcal{T}_{max}\, \ell$ / ħ**

Thus the throughput relation becomes:

$$\mathbf{c = \mathcal{T}_{max} \ell^{*2} / \hbar} \; ... \; (16.1)$$

Equation (16.1) is not an independent postulate; it simply expresses the statement that the maximal carrier advances one minimal hop per substrate commitment tick.

16.2 Elimination of the Substrate Tension

From Theorem 1 (Throughput–Geometry Coupling), the maximum restoring tension satisfies:

$$\mathcal{T}_{max} \sim c\_T^4/G$$

Substituting into (16.1) yields:

$$c\_T = c\_T^4 \ell^2 / (G\hbar) \Rightarrow \ell^2 \sim \hbar G/c\_T^3 \; ... \; (16.2)$$

Equation (16.2) is the familiar Planck-length relation. In the present framework, however, it is not taken as a definition. Instead, it emerges as a consistency condition linking:

- the identity-collapse scale ℓ^* (from finite distinguishability),
- gravitational coupling G,
- quantum action ħ, and
- the maximal propagation speed c_T.

At this stage, no numerical value of c_T has been derived; we have only established that these quantities cannot be independent. (The identification c = c_T for the maximally efficient carrier is established in Theorem 4.)

16.3 Incorporating the Mesoscopic Coherence Scale

Theorem 3 established that the mesoscopic coherence window scales as:

$$\xi\_meso \sim \sqrt{(\ell \cdot L\_IR)}^* \; ... \; (16.3)$$

where L_IR is the cosmological coherence scale (e.g., horizon or Λ-boundary).

Solving (16.3) for ℓ^*:

$$\ell = \xi\_meso^2 / L\_IR^* \; ... \; (16.4)$$

Substituting (16.4) into the Planck-scale consistency relation (16.2) yields a closure relation for the speed of light:

$$\mathbf{c = (L\_IR^2 \cdot \hbar G / \xi\_meso^4)^{(1/3)}} \; \| \; ... \; (16.5)$$

Equation (16.5) is the central result of this section.

16.4 Interpretation and Epistemic Status

It is crucial to distinguish two logically distinct uses of (16.5):

(i) Closure given an independently measured ξ_{meso}

If the mesoscopic coherence scale ξ_{meso} is measured independently (for example, via experiments probing the Two-Planck window), then equation (16.5) determines the numerical value of c given ħ, G, and L_{IR}.

Note on L_{IR}: For a strictly non-circular IR input, one may take $L_{IR} \propto \Lambda^{-1/2}$, which does not explicitly involve c. Alternatively, using $L_{IR} \sim c/H_0$ creates a slightly different closure structure.

Using:

- $\xi_{meso} \approx 80–100\ \mu m$
- $L_{IR} \sim 10^{26}$ m
- measured ħ and G

L_{IR} definition used: For this calculation, we use $L_{IR} = c/H_0 \approx 4.4 \times 10^{26}$ m (the Hubble radius). Alternative definitions include:

- $\sqrt{(3/\Lambda)} \approx 1.6 \times 10^{26}$ m (de Sitter radius)
- Particle horizon $\approx 4.7 \times 10^{26}$ m

These differ by factors of 2–3. Since c enters the closure relation as $L_{IR}^{(2/3)}$, this propagates to ~15–25% uncertainty in the prediction. The 6% agreement quoted uses $L_{IR} = 4.4 \times 10^{26}$ m; with $L_{IR} = 1.6 \times 10^{26}$ m, the predicted c would be ~40% lower. The choice of Hubble radius is motivated by its role as the causal coherence scale—the maximum distance over which causal correlations can be maintained.

Quantity	Value
L_{IR}^2	1.94×10^{53} m²
ħG	7.04×10^{-45} m⁵/s³
ξ_{meso}^4 (at 88 μm)	6.0×10^{-17} m⁴
$L_{IR}^2 \cdot \hbar G / \xi_{meso}^4$	2.28×10^{25} m³/s³

$c = (2.28 \times 10^{25})^{(1/3)} = 2.83 \times 10^8$ m/s

The observed value is $c = 2.998 \times 10^8$ m/s. **Agreement: within 6%.**

This constitutes a non-trivial numerical closure, conditional on an independent determination of ξ_{meso}.

*(ii) Self-consistency when ξ_{meso} is derived from ℓ(c)**

If, instead, ξ_{meso} is computed by combining (16.3) with (16.2)—that is, if one substitutes $\ell^* = \sqrt{\hbar G/c^3}$ back into the definition of ξ_{meso}—then equation (16.5) reduces identically to:

$c = c$

In this case, the equation expresses self-consistency, not an independent prediction. This is expected: once the loop is closed, no new numerical information can be extracted.

16.4.3 Error Propagation and Sensitivity

From (16.5):

$c = (L_{IR}^2 \cdot \hbar G / \xi_{meso}^4)^{(1/3)}$

Taking differentials:

$\Delta c/c \approx (1/3)(2\Delta L_{IR}/L_{IR} + \Delta\hbar/\hbar + \Delta G/G - 4\Delta\xi_{meso}/\xi_{meso})$

Sensitivity to ξ_{meso}:

Holding other inputs fixed:

$\Delta c/c \approx -(4/3)\, \Delta\xi_{meso}/\xi_{meso}$

A 4.5% shift in ξ_{meso} produces a 6% shift in c.

Value needed for exact match:

If $c_{pred} = 2.83 \times 10^8$ m/s and $c_{obs} = 2.998 \times 10^8$ m/s, then:

$\xi_{req} = \xi_{used} \times (c_{pred}/c_{obs})^{(3/4)}$

Numerically, $c_{pred}/c_{obs} \approx 0.944$, so $\xi_{req} \approx 0.957 \times \xi_{used}$.

With $\xi_{used} = 88$ μm: **$\xi_{req} \approx 84$ μm**

The 6% discrepancy corresponds to a few-micron shift in ξ_{meso}.

Sensitivity to L_{IR}:

$\Delta c/c \approx (2/3)\, \Delta L_{IR}/L_{IR}$

Order-unity differences in the definition of L_IR (e.g., $\Lambda^{-1/2}$ vs $\sqrt{3/\Lambda}$ or horizon variants) can contribute at the few–10% level, comparable to the observed discrepancy.

Constants ħ and G:

ħ is known extremely precisely, while G is comparatively uncertain (ppm to tens of ppm). Because G enters only as G^(1/3), its contribution is suppressed. The dominant uncertainty comes from ξ_meso and the definition of L_IR, not from ħ or G.

Conclusion: The 6% agreement is plausibly within the "order unity" ambiguity of L_IR and a few μm uncertainty in ξ_meso.

16.5 What Has Been Achieved

This section does not claim to derive the speed of light from nothing. Rather, it establishes:

1. **A structural closure relation** linking (ħ, G, c, Λ)
2. **A derived mesoscopic coherence scale** ξ_meso ~ √(ℓ* · L_IR)
3. **A clear criterion** for what would constitute an independent determination of c

In standard physics, (ħ, G, c, Λ) are treated as independent constants. In the present framework, they are related by closure conditions, reducing the number of independent parameters by one.

16.6 The Role of the Fine-Structure Constant and Non-Circularity

A natural objection asks: "Doesn't c appear in $\alpha = e^2/(4\pi\varepsilon_0\hbar c)$? So isn't any derivation involving α circular?"

What we are NOT saying:

- "The fine-structure constant α contains c, so of course we can calculate c." (This would be circular.)
- "Given α alone we can compute c." (This is false in standard physics and remains false here.)

What we ARE saying:

α is a *dimensionless* constant. It does not encode meters or seconds—it encodes *closure ratios*. Specifically:

1. **α fixes EM fold geometry:** The fine-structure constant determines the closure geometry of electromagnetic folds—their polarization structure, gauge phase cost, and coupling efficiency.
2. **Irreversible fact formation fixes N_b:** The requirement that physical facts be irreversible determines how many commitments are needed per fold (N_b = 5 for EM).

3. **Together, these fix the dimensionless throughput bound κ_max** for electromagnetic propagation.
4. **Only when this throughput is embodied** (choosing meters per hop and bits per second) does a numerical c emerge.

The correct statement is:

"The fine-structure constant does not numerically determine the speed of light. Rather, it fixes the closure geometry of electromagnetic folds. Combined with the requirement of irreversible fact formation, this geometry selects a unique maximal propagation throughput. The numerical value of c arises only when this throughput is embodied in metric units."

Why this is not Planck-length circularity:

Quantity	Standard definition	VERSF derivation
Planck length ℓ_P	Defined using c	—
Identity-collapse scale ℓ*	—	Inferred from fact stability (no c input)
Planck relation	Definition	Discovered as consistency condition

The identity-collapse scale ℓ* is inferred from finite distinguishability, not defined using c. Only *afterwards* do we discover that ℓ*, G, ħ, and c satisfy the Planck relation. This is analogous to *discovering* $E = mc^2$, not *defining* mass as E/c^2.

The core derivation chain:

1. Irreversible facts exist → maximum hops-per-bit throughput κ_max
2. Depth is an ordering, not a dimension → propagation is void-surface update
3. EM is the minimally closed fold → EM saturates κ_max
4. α fixes EM fold geometry → determines closure efficiency
5. Therefore the maximal physically realizable propagation channel has a fixed throughput
6. Once we embody that throughput (choose meters per hop and bits per second), we obtain numerical c

Nothing in steps 1–5 requires c as an input.

Future work on α: In this paper, α is treated as an empirical dimensionless closure ratio. Deriving α from deeper closure requirements (e.g., from the structure of U(1) gauge folds or topological constraints on phase closure) is delegated to companion work.

---

17. Emergence of Metric Units from Causal Ordering

For the general reader: In everyday life, we measure speed as distance divided by time. But in this framework, neither distance nor time exists at the most fundamental level. Both emerge from a deeper process: the irreversible creation of facts. This section explains how familiar units such as meters, seconds, and the speed of light arise as projections of causal ordering, rather than as primitive quantities.

17.1 No Primitive Meters, No Primitive Seconds

In the VERSF framework, neither spatial distance nor temporal duration is fundamental.

- There is no underlying spatial metric.
- There is no background time parameter.
- There are only irreversible commitments (bits) and their causal ordering.

What exists fundamentally is:

- **Fact creation** (irreversible distinctions)
- **Ordering of those facts** (depth)
- **Correlated propagation** of facts via folded causal structures

Everything else—geometry, duration, motion—is inferred.

17.2 Depth as Causal Ordering, Not a Dimension

Let $\{B_i\}$ denote irreversible commitments (bits). If B_i must exist before B_j, then:

$$\mathbf{B_i \prec B_j}$$

This partial order defines **depth**.

Depth is not a length, a coordinate, or a dimension. It is an ordering relation among facts.

What we experience as "space" is a holographic reconstruction of this ordering, inferred from stable correlation patterns among bits.

17.3 Propagation Without Motion

In this framework, nothing moves through space.

What is usually called "propagation" is the creation of correlated facts across the void substrate:

- A signal does not travel through meters.
- It generates a new fact that is correlated with a previous fact.

- Repeating this process creates a chain of correlations.

Light is special because it is the minimal folded structure capable of sustaining such correlations across depth without internal bookkeeping overhead.

17.4 How Time Emerges

A clock is any physical system that produces irreversible transitions in a regular, reproducible way.

Let:

- N_{bit} = number of irreversible commitments produced by a process
- N_{sec} = commitments defined as "one second" by an observer

Then experienced time is:

$$t_{exp} \equiv N_{bit} / N_{sec}$$

Time is therefore:

- not fundamental,
- not universal,
- but a label applied to a count of irreversible events.

Different clocks count different bits; this is why time dilates.

17.5 How Distance Emerges

Similarly, spatial distance is not primitive.

Two events are considered "far apart" if many irreversible distinctions must exist between them to maintain stable correlation.

Distance is defined operationally as: the number of causal updates separating two correlated facts, compressed into a geometric representation.

Meters are not fundamental objects; they are units assigned to stable correlation depth.

17.6 The Meaning of Speed in This Framework

Because neither meters nor seconds are fundamental, speed is not fundamentally distance divided by time.

The invariant quantity is:

$\kappa \equiv$ (correlated fact updates) / (irreversible commitments)

This is a dimensionless throughput: how efficiently correlations can be propagated per irreversible fact.

Finite distinguishability implies a strict upper bound:

$\kappa \leq \kappa_{max}$

This bound is the true invariant.

17.7 The Speed of Light as an Emergent Ratio

The speed of light appears only when this invariant throughput is expressed in emergent units:

$c \equiv$ (correlation depth labeled as one meter) / (bit count labeled as one second)

Both numerator and denominator are counts of irreversible commitments.

Thus:

- c is not a fundamental speed
- It is a ratio of two emergent conventions
- Both derived from the same underlying causal process

This is why:

- clocks run slow,
- rulers contract,
- but the speed of light remains invariant.

They all rescale together because they are built from the same bit substrate.

17.8 Why "Meters per Hop" Language Is Shorthand

Earlier expressions such as "meters per hop" or "bits per second" are shorthand. They should be read as:

- "meters" = a chosen projection of correlation depth
- "seconds" = a chosen projection of irreversible bit count
- "hops" = steps in causal ordering

These are not independent primitives, and treating them as such introduces apparent circularity.

Once this is recognized, the circularity disappears: all such quantities are different views of the same underlying process.

17.9 Restating the Core Claim Without Circularity

The fundamental claim of this paper is:

The universe admits a maximal rate at which correlated facts can be created without destroying stability. Electromagnetic folds saturate this rate. The numerical value of the speed of light arises when this invariant throughput is expressed in conventional metric units built from the same irreversible processes.

This is the sense in which the speed of light is **explained**, not merely postulated.

17.10 Consequence for the Closure Relation

With this clarified, the closure relation:

$$c = (L_{IR}^2 \cdot \hbar G / \xi_{meso}^4)^{1/3}$$

should be interpreted as a closure relation between emergent quantities, not as a definition of a primitive constant.

It shows that once:

- irreversible facts exist,
- correlations must propagate stably, and
- large-scale coherence is enforced,

then the emergent metric ratio we call "the speed of light" is no longer free.

Equation (16.5) should therefore be read as a constraint on how emergent spatial and temporal calibrations can be jointly consistent with stable fold propagation under global coherence.

---

18. Corollary: Propagation Cone Equality and Commitment-Cost Inequality

GW170817 bounds $|c_T - c|/c$ at $\sim 10^{-15}$ [3]. Higher N_b manifests as generation difficulty, not slower propagation. □

---

19. Theorem 4: Electromagnetic Folds Are Maximally Efficient Causal Carriers

Theorem 4.

Claim: Among all propagating fold species satisfying (C1)–(C5), electromagnetic folds minimize N_b and saturate the causal cone speed c_T.

Proof (outline):

1. **Minimality of N_b:** By Theorem 2.5, any U(1) gauge fold requires $N_b \geq 5$, and EM achieves exactly $N_b = 5$.
2. **Throughput saturation:** The throughput of a fold species is $c = \ell_{hop} \cdot \nu^*$. For minimal N_b, the hop length $\ell_{hop} = \xi/N_b$ is maximized for fixed coherence ξ.
3. **No lower-cost carrier:** Any fold with $N_b < 5$ would violate one of (C1)–(C5) and fail to propagate stably.
4. **Therefore:** EM saturates the maximum throughput bound, so $c = c_T$.

Corollary: The identification $c = c_T$ holds for the maximally efficient carrier (EM). □

---

20. Consistency Check: Lorentz Structure from Invariant Fold Throughput

Under A1 (relativity), A2 (isotropy), A3 (invariant c_T), kinematics must be Lorentzian.

$$x' = \gamma(x - vt)$$

$$t' = \gamma(t - vx/c_T^2)$$

where $\gamma = 1/\sqrt{1 - v^2/c_T^2}$

□

Note: This section is a consistency reinterpretation, not an independent derivation of Lorentz invariance. The Lorentz transformation follows from standard SR axioms once c_T is identified as the universal causal limit.

20.1 The VERSF Reinterpretation of Special Relativity

For the general reader: Special relativity does not require space and time themselves to stretch or shrink. What changes between observers is the rate at which irreversible facts are produced and registered—the "frame rate" of reality.

Important: This reinterpretation preserves the Lorentz transformation and all experimentally verified predictions of special relativity; it changes only the ontology (what the symbols mean).

What is invariant:

Standard Relativity	**VERSF**
Spacetime interval	Maximum causal throughput (κ_{max})

What is *not* invariant: the number of irreversible commitments a given physical process can perform per observer-defined "second." That quantity must vary between observers, otherwise causality would break.

20.2 Time Dilation = Reduced Local Frame Rate

In VERSF, a clock is just a physical process that produces irreversible commitments at some rate.

When an object is moving relative to an observer:

- Part of its available update capacity is consumed maintaining correlations with the external frame (motion bookkeeping)
- Fewer updates remain available for internal processes (atomic transitions, decay, oscillations)

So the clock:

- Does **not** "run slow" because time itself stretches
- Runs slow because its **internal frame rate drops**

Formally:

proper time ∝ number of irreversible updates

Different observers disagree on time because they count different numbers of updates for the same physical process.

This is why:

- Moving clocks tick fewer times
- Unstable particles live longer when moving fast

20.3 Length Contraction = Fewer Frames per Spatial Correlation

Similarly, length is not a static geometric extent.

In VERSF, "length" is reconstructed from how many stable correlations exist across an object.

For a fast-moving object:

- Fewer irreversible updates are available to maintain internal spatial correlations along the direction of motion
- Transverse correlations are unaffected

So:

- The object does **not** physically compress
- The observer reconstructs fewer correlation layers → shorter measured length

This matches exactly the directional nature of Lorentz contraction.

20.4 Why the Lorentz Factor Still Appears

The familiar Lorentz factor $\gamma = 1/\sqrt{(1 - v^2/c^2)}$ still appears—but its meaning changes.

In VERSF:

- γ is a **throughput reallocation factor**
- It quantifies how much of the universal update budget is diverted into maintaining motion-related correlations instead of internal updates

γ measures how much the local frame rate is reduced relative to the maximal causal update rate.

In this ontology, γ is interpreted as the ratio of available internal irreversible updates in the comoving frame to those available in the lab frame under the invariant throughput bound.

This keeps all the mathematics of special relativity intact while changing the ontology underneath it.

20.5 Why All Observers Agree on c

This is the key consistency check.

Even though observers have different frame rates, they all agree on the maximum causal throughput. That's because:

- c is not "distance per time"
- It is the **upper bound on correlated update propagation**

If an observer tried to measure a signal exceeding c:

- They would need more correlation updates than their frame rate allows
- Which is impossible

So invariance of c is automatic—it's the ceiling everyone shares.

20.6 Relativity of Simultaneity = Disagreement About Update Ordering

In standard SR, simultaneity is relative because spacetime slices differ.

In VERSF, simultaneity is relative because different observers group irreversible updates differently.

Two events may:

- Be registered in the same update batch for one observer
- But in different batches for another

There is no contradiction, because:

- There is no absolute global "now"
- Only local update orderings

20.7 Summary: Same Predictions, Different Ontology

This reinterpretation:

- Reproduces all tested predictions of special relativity
- Explains *why* those predictions exist
- Removes the need for spacetime as a fundamental object

It aligns naturally with quantum irreversibility, entropy production, computational limits, and fold-based propagation.

Nothing experimental changes. Only the explanation changes.

In the VERSF framework, special relativistic effects admit a natural reinterpretation. Time dilation and length contraction do not require spacetime itself to stretch or deform. Instead, they reflect differences in the rate at which irreversible facts are produced and registered by physical systems in relative motion. A clock is a process that generates irreversible commitments; motion diverts part of the available update capacity toward maintaining external correlations, reducing the number of updates available for internal dynamics. Length contraction arises because fewer correlation layers can be maintained along the direction of motion. The Lorentz factor quantifies this redistribution of update capacity. The invariance of the speed of light follows because it is not a distance-per-time ratio but the universal upper bound on causal correlation throughput, shared by all observers regardless of their local frame rate.

21. Testable Prediction: Generation–Propagation Asymmetry

Carrier	N_b	L_{min}	Power scaling
EM	5	1	$P \propto (v/c)^4$
Gravity	>5	2	$P \propto (v/c)^6$

$E_{gen}(X)$ is monotone increasing in $N_b(X)$

---

22. Objections and Resolutions

(O1) Is the derivation of c circular? *Resolution:* See Section 16.4–16.6 for detailed analysis. The short answer: the closure relation $c = (L_{IR}^2 \cdot \hbar G / \xi_{meso}^4)^{(1/3)}$ is self-consistent when ξ_{meso} is derived from Theorem 3 (yielding c = c). If ξ_{meso} is measured independently (~80 μm), c is directly determined—a non-trivial prediction.

(O2) Doesn't α contain c, making any derivation involving α circular? *Resolution:* No. The fine-structure constant α is *dimensionless*—it encodes closure ratios, not meters or seconds. α fixes the geometry of electromagnetic folds; irreversible fact formation fixes $N_b = 5$; together these determine the dimensionless throughput κ_{max}. Only when embodied in metric units does c emerge. See Section 16.6.

(O3) Why does the framework predict $\xi_{meso} \approx 80$ μm? *Resolution:* The proposed Two-Planck window is the geometric mean of UV and IR scales—exactly what Theorem 3 predicts. This constitutes an order-of-magnitude compatibility check (not yet an independent prediction, since ℓ^* currently uses measured c).

(O4) GW170817 and gravity speed. *Resolution:* Higher N_b affects generation difficulty, not cone speed.

(O5) The 30-order-of-magnitude gap (ξ_{UV} vs ξ_{meso}). *Resolution:* Explained by Theorem 3: $\xi_{meso}/\xi_{UV} \sim \sqrt{(L_{IR}/\ell^*)}/N_b \sim 10^{30}$.

(O6) What sets L_{IR}? *Resolution:* L_{IR} is the cosmological horizon or Λ-boundary. Its value comes from observation. The framework shows how c depends on L_{IR}.

---

23. Relation to Existing Work

Framework	How c Appears	VERSF Difference
GR	Structural invariant	Derives relationship $c = f(\hbar, G, L\_IR, \xi\_meso)$
QFT	Built into Lorentz structure	Time emerges from commitments
Emergent Gravity	Origin unspecified	UV/IR balance determines coherence
Holography	IR/UV connection	Similar spirit; different mechanism

The UV/IR connection has parallels to holographic ideas [7], but emerges here from stability requirements.

---

24. Falsifiability and Empirical Handles

For reviewers: This section identifies concrete empirical tests that could confirm or refute the framework's claims.

The framework makes three testable commitments:

H1: Mesoscopic coherence crossover at ξ_meso ~ 10^{-4} m

The framework predicts a characteristic coherence scale $\xi\_meso \sim \sqrt{(\ell^* \cdot L\_IR)} \approx 30\text{–}100$ μm.

Potential experimental signatures:

- Quantum decoherence experiments at mesoscopic scales
- Coherence length measurements in interferometric setups
- Anomalous behavior in optomechanical systems near this scale

Falsification criterion: If precision experiments definitively rule out any coherence feature in the 10–200 μm range, the geometric-mean prediction fails.

H2: ξ_meso should track √L_IR across cosmological conditions

If Λ were different (different cosmological epoch, different universe), the mesoscopic coherence scale should shift:

$$\xi\_meso \propto \sqrt{L\_IR} \propto \Lambda^{-1/4}$$

In practice: This is difficult to test directly, but:

- Cosmological models with different Λ values can be analyzed for consistency
- Early-universe conditions (smaller L_{IR}) would predict smaller ξ_{meso}
- This provides a consistency check across cosmological regimes

Falsification criterion: If a consistent cosmological model with different Λ shows ξ_{meso} scaling differently than $\sqrt{L_{IR}}$, the UV/IR balance mechanism fails.

H3: Frame-rate reinterpretation must reproduce SR exactly

The VERSF reinterpretation of special relativity (Section 20) changes only the ontology, not the predictions:

- Time dilation: $\Delta t' = \gamma \Delta t$ (unchanged)
- Length contraction: $L' = L/\gamma$ (unchanged)
- Lorentz transformation: exact (unchanged)

Commitment: The framework predicts **zero deviation** from standard SR kinematics.

Falsification criterion: Any measured deviation from Lorentz invariance (e.g., in high-energy cosmic rays, precision atomic clocks, or Michelson-Morley-type experiments) would falsify both standard SR and this reinterpretation.

H4: Generation–propagation asymmetry for different carriers

The framework predicts that carriers with higher N_b (more commitments per fold) should be harder to generate but propagate at the same speed:

Carrier	N_b	Generation difficulty	Propagation speed
EM	5	Baseline	c
Gravity	>5	Higher	c

Falsification criterion: If gravitational waves were found to propagate at a speed measurably different from c (beyond the 10^{-15} bound from GW170817), the universal cone-speed claim fails.

Summary: What would kill this framework?

Observation	Framework Status
No coherence feature near 10^{-4} m	ξ_meso prediction fails
ξ_meso doesn't scale with √L_IR	UV/IR mechanism fails
Any Lorentz violation	SR reinterpretation fails
Gravity waves ≠ light speed	Universal cone fails

The framework is falsifiable. It commits to specific predictions that can, in principle, be tested.

---

25. Conclusion

Main Theorems:

1. **Theorem 1:** G, c_T, ℓ* not independent; Planck length emerges.
2. **Theorem 2:** ξ_UV = N_b · ℓ* (substrate coherence fixed point).
3. **Theorem 3:** ξ_meso = √(ℓ* · L_IR) (UV/IR geometric mean).
4. **Theorem 4:** EM folds minimize N_b and saturate c_T.

The closure relation for c:

$c = (L_{IR}^2 \cdot \hbar G / \xi_{meso}^4)^{(1/3)}$ ... (16.5)

What is achieved:

- A structural closure relation linking (ħ, G, c, Λ)
- Order-of-magnitude compatibility: ξ_meso ≈ 80 μm (matching ~88 μm proposed Two-Planck window)
- A clear criterion for independent determination of c
- Potential reduction of independent constants (contingent on independent ξ_meso measurement)

Epistemic status:

- If ξ_meso is measured independently → c is determined (non-trivial prediction)
- If ξ_meso is derived from ℓ*(c) → c = c (self-consistency check)
- Current numerical checks use standard values that depend on measured c; the framework is *compatible* with the correct order of magnitude

What remains open:

- Independent precision measurement of ξ_meso

- First-principles derivation of ℓ* from distinguishability closure (not using c)
- First-principles derivation of L_IR (why Λ takes its value)
- Experimental test: ξ_meso should track √L_IR across cosmologies

The bottom line: This paper does not derive c from nothing. It establishes that (ħ, G, c, Λ) are related by closure conditions, so that given any three plus the mesoscopic coherence scale, the fourth is determined. The ~6% numerical agreement with observed c—conditional on ξ_meso ≈ 88 μm—is a non-trivial consistency check.

Summary of epistemic posture: The framework establishes structural constraints that are mathematically proven, numerical compatibility checks that are consistent with observation within stated uncertainties, and explicit falsification criteria that render the framework empirically testable. While some inputs remain empirical (notably ξ_meso and α), the closure relations significantly reduce arbitrariness in the relationship between fundamental constants and offer a novel explanatory perspective consistent with known physics.

References

[1] R. Landauer, "Irreversibility and Heat Generation in the Computing Process," *IBM J. Res. Dev.* **5**, 183 (1961).

[2] N. Margolus and L. B. Levitin, "The maximum speed of dynamical evolution," *Physica D* **120**, 188 (1998).

[3] B. P. Abbott et al., "GW170817 and GRB 170817A," *Astrophys. J. Lett.* **848**, L13 (2017).

[4] T. Baker et al., "Strong Constraints on Cosmological Gravity from GW170817," *Phys. Rev. Lett.* **119**, 251301 (2017).

[5] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (W. H. Freeman, 1973).

[6] S. Lloyd, "Ultimate physical limits to computation," *Nature* **406**, 1047 (2000).

[7] J. Maldacena, "The Large N Limit of Superconformal Field Theories and Supergravity," *Adv. Theor. Math. Phys.* **2**, 231 (1998).

[8] T. Jacobson, "Thermodynamics of Spacetime: The Einstein Equation of State," *Phys. Rev. Lett.* **75**, 1260 (1995).

[9] J. A. Wheeler, "Information, Physics, Quantum: The Search for Links," in *Complexity, Entropy, and the Physics of Information* (Addison-Wesley, 1990).

[10] G. 't Hooft, "The Cellular Automaton Interpretation of Quantum Mechanics," *arXiv:1405.1548* (2014).

[11] E. Verlinde, "On the Origin of Gravity and the Laws of Newton," *JHEP* **2011**, 29 (2011).

[12] J. D. Bekenstein, "Black Holes and Entropy," *Phys. Rev. D* **7**, 2333 (1973).

Appendix A: Empirical Anchors for the Mesoscopic Coherence Scale

Addresses the concern: "The empirical anchor for ξ_meso is still soft."

Empirical Status

The mesoscopic coherence scale $\xi_{meso} \sim \sqrt{\ell^* \cdot L_{IR}}$ emerges in this framework as the characteristic length separating UV identity-collapse dynamics from IR coherence loss. At present, **ξ_{meso} should be regarded as a target scale, not a confirmed constant.**

However, the framework does not leave ξ_{meso} empirically unconstrained. It provides specific experimental handles by which this scale could be independently probed or bounded.

Potential Empirical Anchors

1. Decoherence crossover experiments

Interferometric or optomechanical systems operating at mesoscopic length scales (10–100 μm) may exhibit a crossover in coherence behavior not attributable to thermal, electromagnetic, or environmental noise alone. A sharp deviation from expected scaling near ξ_{meso} would constitute direct evidence.

2. Precision force or noise anomalies

Experiments measuring Casimir forces, vacuum fluctuations, or unexplained noise spectra at sub-millimeter separations could reveal departures from standard predictions near the proposed scale.

3. Quantum-limited mechanical resonators

Systems engineered to approach the quantum–classical boundary may show a saturation or transition in coherence persistence when characteristic dimensions approach ξ_{meso}.

4. Null results as constraints

Even the absence of anomalies at the 10–100 μm scale constrains the framework, tightening allowable ranges for ℓ^*–L_{IR} coupling and potentially falsifying the proposed UV/IR balance.

Key point: The framework predicts *where to look*. The absence of a detected effect is not a failure of testability, but a meaningful empirical constraint.

Appendix B: Robustness of Theorem 3 Under Alternative Instability Functionals

Addresses the concern: "Theorem 3 depends on a specific modeling choice."

The Modeling Choice

Theorem 3 derives the scaling $\xi_{meso} \sim \sqrt{\ell^* \cdot L_{IR}}$ by minimizing a simple additive instability functional:

$\varepsilon_{tot}(\ell) = \ell/\ell + \ell/L_{IR}^*$

This choice is not arbitrary. It represents the **minimal monotone combination** of two independent failure modes:

- UV instability increasing as ℓ decreases
- IR instability increasing as ℓ increases

Robustness Under Alternative Combinations

Crucially, the geometric-mean scaling is **robust under a broad class of alternative combinations:**

Weighted sums:

$\varepsilon(\ell) = a(\ell^*/\ell) + b(\ell/L_{IR})$

still minimize at $\ell \sim \sqrt{\ell^* \cdot L_{IR}}$ up to order-unity factors.

Multiplicative combinations:

$\varepsilon(\ell) = (\ell^*/\ell)^p \cdot (\ell/L_{IR})^q$

yield extrema at $\ell \propto (\ell^* \cdot L_{IR})^{(1/2)}$ for any positive p, q.

Smooth monotone interpolations:

Any functional satisfying:

- $\varepsilon(\ell) \to \infty$ as $\ell \to 0$

- $\varepsilon(\ell) \to \infty$ as $\ell \to L_{IR}$

admits an interior minimum whose location is controlled by the product $\ell^* \cdot L_{IR}$.

Conclusion

The appearance of the geometric mean is **not a fine-tuned artifact**, but a structural consequence of balancing independent UV and IR instabilities. The exact numerical prefactor may vary, but the scaling itself is stable.

Appendix C: Clarifying the Reduced-Constants Claim

Addresses the concern: "The reduced-constants claim is contingent."

What the Framework Does NOT Assert

The framework does not assert that the speed of light c is uniquely derivable from $(\hbar, G, \Lambda)$ without further empirical input.

What the Framework DOES Assert

It establishes a **closure structure:**

$$F(\hbar, G, c, \Lambda, \xi_{meso}) = 0$$

The logical content of this result is:

1. These quantities **cannot be freely specified independently** once coherence closure is imposed.
2. **Fixing any four determines the fifth.**
3. **Predictive power arises only if ξ_{meso} is fixed independently of $(\hbar, G, c)$.**

The Two Cases

If ξ_{meso} is...	Then the closure relation...
Independently measured	Makes c no longer a free constant
Derived using relations involving c	Reduces to a consistency check

This is why the paper repeatedly distinguishes:

- **structural constraint** from
- **numerical derivation**

The Precise Claim

"The framework proposes that (ħ, G, c, Λ) are mutually constrained by coherence closure. Whether this reduces the number of empirically independent constants depends on the independent fixation of ξ_meso."

This formulation is intentionally conservative and falsifiable.

Interpretive Note: Explanatory Gain

The framework does not eliminate empirical input; it reorganizes it. Standard physics treats (ħ, G, c, Λ) as independent. Here, global coherence supplies L_IR (linked to Λ by cosmology), and closure relations constrain how the remaining constants can co-exist. The gain is not "no observation," but **reduced arbitrariness**: the constants are no longer freely specifiable simultaneously once stability of fact propagation is imposed.

Appendix D: Summary of Empirical Handles and Falsifiability

Addresses the concern: "The framework needs explicit falsifiability criteria."

The VERSF framework makes no claim that new physics must appear beyond established experimental bounds. Instead, it identifies specific empirical handles by which it can be tested or constrained.

H1: Mesoscopic Coherence Crossover

A detectable change in coherence behavior, noise scaling, or stability near 10^{-4} m would support the existence of ξ_meso.

Falsification: Absence of any anomaly or boundable effect across this regime constrains or excludes the proposed UV/IR balance.

H2: Scaling with Cosmology

The framework predicts: **ξ_meso ∝ √L_IR ∝ $\Lambda^{-1/4}$**

Any observational or theoretical context in which L_IR is effectively altered (e.g., alternative cosmological models) should shift the coherence scale accordingly.

Falsification: Demonstrated independence of coherence scales from L_IR.

H3: No Deviation from Special Relativity

The frame-rate reinterpretation of SR commits to exact agreement with all standard relativistic predictions.

Falsification: Any experimentally confirmed deviation from Lorentz invariance attributable to this ontology would refute it.

H4: Radiation Efficiency Hierarchy

The predicted monotonic relationship between commitment cost N_b and radiation generation efficiency is testable in principle across different interaction channels.

Falsification: Discovery of a fundamental carrier that propagates at c with lower generation cost than EM would contradict the framework.

Summary Table

Handle	Prediction	Falsification Criterion
H1	Coherence crossover at $\sim 10^{-4}$ m	No effect in 10–200 μm range
H2	$\xi_{meso} \propto \sqrt{L_{IR}}$	ξ_{meso} independent of L_{IR}
H3	Zero deviation from SR	Any Lorentz violation
H4	EM is minimally costly carrier	Carrier with $N_b < 5$ at speed c

The framework is falsifiable. It commits to specific predictions that can, in principle, be tested.