

Hexagonal Interface Light Propagation: A Comprehensive Framework

A Unified Admissibility-Based Theory of Emergent Electromagnetism

Abstract

We present a unified framework in which light propagates fundamentally on a two-dimensional causal interface rather than as a primitive three-dimensional bulk process. Using Bit Conservation & Balance (BCB) and admissibility as selection principles, we demonstrate that: (1) isotropy is required rather than assumed, (2) hexagonal coordination is uniquely selected among regular tilings, (3) Lorentz structure emerges and Maxwell electrodynamics is recovered in the continuum limit given link-based $U(1)$ degrees of freedom, and (4) primitive 3D propagation is non-admissible under finite-information constraints.

The framework is stress-tested via renormalization-group analysis, information theory, and comparison with related discrete-spacetime programmes. We explain nine structural features of light that standard physics describes but does not derive, and yield a concrete falsifiable experimental signature: a six-fold (6Ω) harmonic in rotating optical-cavity experiments. Detection of four-fold (4Ω) anisotropy instead would falsify the hexagonal selection, providing a sharp experimental discriminant between competing discrete geometries.

For the General Reader

What if the three-dimensional space we inhabit is not fundamental, but emerges from something simpler?

This paper argues that light does not travel through three-dimensional space in the way we usually imagine. Instead, it propagates on a two-dimensional surface—like ripples on a pond rather than sound waves filling a room. The 3D world we experience is reconstructed from patterns on this 2D surface, much like a hologram creates the illusion of depth from a flat image.

An analogy: spreadsheets all the way down. Imagine reality's finest level is a vast 2D spreadsheet—every cell containing detailed data about what's happening at that location. Now imagine creating a summary spreadsheet that averages every 10×10 block of cells into a single value. Then another summary of *that* summary, averaging again. Stack these layers of increasingly coarse summaries on top of each other, and you've created a third dimension: depth

corresponds to "how zoomed out" you are. The original spreadsheet is the fundamental 2D interface; the third dimension emerges from how information at different scales relates to each other. Two points are "close in the third dimension" if their fine-grained details are correlated; they're "far apart" if you have to zoom out considerably before their data starts to look similar.

Why believe this? The argument comes from thinking carefully about *information*. Every moment, the universe must "decide" what happens next at every point in space. But there's a limit to how much deciding can happen at once—a cosmic information budget. When light spreads outward in three dimensions, it must spread this budget across an ever-growing sphere. Eventually, the budget per point becomes too thin to actually determine what happens there. The mathematics shows this problem doesn't arise in two dimensions, where the budget stays balanced at all distances.

If light really propagates on a 2D surface, what shape should that surface have? We show that a honeycomb pattern—hexagonal tiles—is the unique geometry that satisfies all the consistency requirements. This isn't arbitrary; four convergent mathematical arguments all point to hexagons.

The payoff is a testable prediction. If light propagates on a hexagonal grid, there should be a tiny six-fold pattern in how the speed of light varies with direction—like a subtle grain in the fabric of space. Modern experiments can detect variations in light speed smaller than one part in ten thousand trillion. If they find a six-fold pattern, this theory gains support. If they find a four-fold pattern instead (which would indicate a square grid), the theory is wrong. That's what makes this science rather than speculation.

The deeper implication: the speed of light isn't just a number that happens to be what it is. It's the maximum speed at which the universe can commit new facts into existence—the speed limit of reality's record-keeping. Space, time, and the laws of electromagnetism all emerge from this more fundamental constraint on information processing.

Table of Contents

1. Introduction and Motivation
 - 1.1 The Explanatory Gap in Standard Physics
 - 1.2 The Interface Hypothesis
 - 1.3 Relation to Existing Discrete-Spacetime Programmes
 - 1.4 Questions This Framework Addresses
2. Foundational Principles: BCB and Admissibility
 - 2.1 Primitive Concepts
 - 2.2 Bit Conservation & Balance (BCB)
 - 2.2a Emergent Time and Fact-Commitment
 - 2.3 The Distinguishability-Energy Distinction
 - 2.4 The Ontological Status of Δi_{\min}
 - 2.4.1 The Minimum Recordable Distinction
 - 2.5 Admissibility
 - 2.6 Isotropy as a BCB Consistency Condition
3. The Non-Admissibility Theorem
 - 3.1 Purpose and Scope
 - 3.2 Formal Setup
 - 3.3–3.6 Lemmas 1–4
 - 3.7 Main Theorem
 - 3.8 Addressing the Energy Analogy Objection
4. Hexagonal Selection
 - 4.1 The Selection Problem
 - 4.2 Dispersion Symmetry Analysis
 - 4.3 Renormalization Group Flow
 - 4.4 Information-Theoretic Optimality
 - 4.5 BCB Admissibility
 - 4.6 Synthesis
5. Emergent Lorentz Symmetry and Maxwell Structure
 - 5.1 Emergent Lorentz Invariance
 - 5.2 Gauge Structure from Link Variables
 - 5.3 Field Strength from Plaquettes
 - 5.4 Extending Gauge Structure to the Emergent Bulk
 - 5.5 Polarization Degrees of Freedom
6. Bulk Reconstruction: Emergent Three-Dimensional Space
 - 6.1 The Reconstruction Problem
 - 6.2 Coarse-Graining Depth as the Third Dimension
 - 6.3 Operational Definition of Bulk Distance
 - 6.4 Emergent Metric
 - 6.5 What Physical Process Corresponds to "Moving in z"?
 - 6.6 Why Observers Agree on Three Dimensions
 - 6.7 Correlation Structure from Distinguishability Geometry
7. Phenomenology: Structural Explanations and Research Directions
8. Falsifiable Predictions
 - 8.1 The Core Prediction: Six-Fold Anisotropy

- 8.2 Distinguishing Prediction: 6Ω vs 4Ω
 - 8.3 Parameter Estimation
 - 8.4 Experimental Channel: Rotating Optical Cavities
 - 8.5 Current Experimental Status
 - 8.6 Implied Bounds on Interface Scale
 - 8.6a Auxiliary Constraints from Independent Channels
 - 8.7 Proposed Experimental Protocol
 - 8.8 Falsification Logic and the Role of ε
 - 8.9 Determining the Suppression Order p
9. Anticipated Objections and Responses
 10. Open Questions and Future Directions
 11. Conclusion

References

Appendices

- Appendix A: Glossary of Key Terms
- Appendix B: Selection Channel Summary
- Appendix C: Experimental Bounds
- Appendix D: Rejection Costs and Alternative Ontologies
- Appendix E: Assumptions Ledger
- Appendix F: Overlap Uniqueness from BCB Geometry (Sketch)
- Appendix G: Minimal Toy Model of Finite-Capacity Fact-Commitment Propagation
- Appendix H: Lattice-Field Toy Model with Explicit Correlation Functions
- Appendix I: Scope, Ontological Forks, and Limiting Assumptions

How to Read This Paper

For general readers and scientists from other fields: Each major section begins with a boxed summary (marked "For the general reader") that explains the key ideas in plain language. You can read just these boxes plus the Introduction and Conclusion to get the main argument. The technical sections between them provide the rigorous backing.

For physicists: The core technical content is in Sections 2–4 (foundational principles and the non-admissibility theorem), Section 5 (emergence of standard field theory), and Section 8 (experimental predictions). Section 6 (bulk reconstruction) is an outline rather than a complete derivation.

For experimentalists: Skip directly to Section 8 for the testable prediction (6Ω harmonic signature) and proposed protocol. The background motivation is in Section 1.

Quick summary: We argue that finite information-processing capacity forces light propagation onto a 2D interface rather than primitive 3D space. Hexagonal geometry is uniquely selected by

four convergent arguments. Standard physics emerges in the continuum limit. The prediction: a tiny six-fold angular variation in light speed, detectable in precision rotating-cavity experiments.

1. Introduction and Motivation

For the general reader: This paper proposes that light doesn't actually travel through three-dimensional space the way we normally imagine. Instead, light propagates on a two-dimensional "surface" (we call it an interface), and the three-dimensional world we experience is reconstructed from patterns on this surface—somewhat like how a hologram creates a 3D image from a 2D film. This isn't just philosophical speculation: we show that consistency requirements *force* this conclusion, and we predict a specific experimental signature that could confirm or refute the idea.

1.1 The Explanatory Gap in Standard Physics

Standard electromagnetism and special relativity provide extraordinarily accurate descriptions of light. Maxwell's equations and Lorentz transformations have been verified to remarkable precision across an enormous range of phenomena. This work does not dispute their empirical success.

Instead, this work addresses a different question: *why* those structures arise at all.

Standard formulations describe *how* light behaves. They do not explain:

- Why the speed of light is invariant across all inertial frames
- Why Lorentz symmetry exists rather than some other spacetime symmetry
- Why gauge redundancy is necessary for consistent electrodynamics
- Why light and geometry are so intimately connected
- Why vacuum exhibits no birefringence or dispersion

These structural features are postulated or derived from other postulates, but the postulates themselves remain unexplained within the standard framework.

1.2 The Interface Hypothesis

We propose that **causality and finite information precede spacetime geometry**. Light corresponds to fundamental causal propagation on a two-dimensional interface from which three-dimensional space is emergent. The interface is not a mathematical boundary condition but a physical causal substrate—the arena in which distinguishability is processed and propagated.

This hypothesis, combined with consistency requirements (BCB/admissibility), uniquely selects:

1. **Isotropy** as a derived requirement, not an assumption
2. **Hexagonal coordination** as the uniquely admissible discrete geometry

3. **Lorentz symmetry** as emergent continuum invariance
4. **Maxwell gauge structure** from link-based degrees of freedom
5. **Three-dimensional space** as a reconstruction from interface correlations

1.3 Relation to Existing Discrete-Spacetime Programmes

Several research programmes have proposed discrete or emergent spacetime structures. We briefly situate this framework relative to them:

Causal Set Theory (Sorkin, Bombelli, et al.): Proposes spacetime as a discrete partial order of events. Our framework shares the emphasis on causality as primitive but differs in several important respects:

- *Geometry*: Causal sets use random Poisson sprinkling to preserve Lorentz invariance statistically. Our framework proposes a specific geometric substrate (2D hexagonal interface) that breaks Lorentz symmetry at the Planck scale but recovers it in the continuum limit through RG suppression.
- *Dimensionality*: Causal sets assume the target dimension; our framework derives 2D as the unique admissible substrate.
- *Lorentz violation*: Causal sets are designed to be exactly Lorentz-invariant (no preferred frame). Our framework predicts residual Lorentz violation—the 6Ω anisotropy—as a testable signature. This is a key empirical discriminant: causal set theory predicts no such anisotropy; our framework predicts a specific one.

The hexagonal lattice explicitly breaks continuous rotation symmetry, leaving only six-fold discrete symmetry. This is *not* hidden by the continuum limit—it is suppressed but potentially detectable. Whether this is a feature or a bug depends on experiment: detection of 6Ω would support our framework over causal sets; continued null results at increasing sensitivity would favour causal sets or require pushing our interface scale to smaller sizes.

Why doesn't the non-admissibility argument apply to causal sets? A reader might ask: if 3D propagation fails the distinguishability budget, why do causal sets escape this argument? The answer is that causal sets do not have a fixed lattice with a deterministic wavefront. In causal set theory, the "elements" are randomly sprinkled, so there is no geometric surface that grows as r^2 requiring budget allocation. The causal structure is probabilistic, and Lorentz invariance is preserved *statistically* across the ensemble of possible sprinklings. Our argument applies to deterministic substrates with fixed geometry; causal sets avoid it by having no fixed geometry at all. This is not a refutation of our framework—it is an alternative ontological choice with different empirical consequences (no residual anisotropy vs. 6Ω signature).

Scope clarification: The non-admissibility theorem does not assume a rigid lattice; it assumes that, at the level where "propagation" is defined, causal influence reaches an expanding set of addressable registers whose cardinality scales like the measure of an effective wavefront. Any substrate with a well-defined coarse-grained front (even if microstructure is disordered) falls into this class. Causal sets evade the argument because they intentionally avoid a fixed wavefront geometry: the "front" is a random sprinkling ensemble, and Lorentz invariance is preserved

statistically. This is not a loophole but a clean fork: fixed-geometry substrates predict symmetry-protected residue (here 6Ω), whereas Lorentz-statistical substrates predict no such residue.

Note on stochastic generalisation: A stochastic generalisation of BCB for ensemble-based substrates (such as causal sets) is an open question. The present work analyses deterministic/fixed-front admissible runtimes; extending the framework to stochastic ensembles may require different constraint formulations. Whether a stochastic substrate with well-defined ensemble-averaged wavefront scaling (intermediate between our deterministic case and causal sets' random sprinkling) would face a modified version of the non-admissibility argument remains unexplored. Such substrates might predict intermediate signatures—partially suppressed but nonzero 6Ω —distinct from both the present framework and exact-Lorentz approaches.

Loop Quantum Gravity: Proposes discrete quantum geometry via spin networks. Our framework operates at a different level—we address the causal substrate for field propagation rather than quantizing geometry directly. The approaches may be complementary rather than competing.

Wolfram's Physics Project: Proposes hypergraph rewriting as fundamental dynamics. Our framework shares the computational/information-theoretic spirit but makes specific geometric commitments (2D interface, hexagonal coordination) that yield concrete experimental predictions absent from the hypergraph approach.

Holographic Principles (AdS/CFT, tensor networks): Standard holography asserts bulk/boundary duality—both bulk and boundary are presumed to exist, and the framework provides a correspondence between them. Our framework addresses a different question: why a bulk description arises at all from an interface-native causal substrate. Causal propagation is interface-native, not merely encoded; the bulk is emergent, not dual.

The distinctive contribution of this framework is the combination of (a) specific geometric selection via multiple convergent channels, (b) concrete falsifiable predictions, and (c) derivation of standard field theory as the continuum limit.

1.4 Questions This Framework Addresses

The interface framework provides structural answers to nine fundamental questions:

Question	Standard Status	Interface Answer
Why is c invariant?	Postulated (SR axiom)	Emerges from fixed interface update rate
Why Lorentz symmetry?	Postulated or derived from c -invariance	Emerges from isotropic finite-speed dynamics
Why does light define causality?	Geometric feature of metric	Causality is primary; light cones are reachable regions
Why gauge redundancy?	Required for consistent matter coupling	Local phase reference freedom on links †

Question	Standard Status	Interface Answer
Why two polarizations?	Poincaré rep theory + gauge fixing	Transverse interface modes in emergent bulk †
Why no vacuum birefringence?	Empirical constraint	Forbidden by hexagonal symmetry
Why apparent continuity?	Assumed; discreteness bounded by experiment	RG suppression of discreteness
Why emergent space? ‡	Not addressed (space is fundamental)	Correlation-based reconstruction
Why light-geometry link? ‡	Encoded in metric; unexplained	Light defines the causal structure geometry encodes

Status key:

- **Derived** (items 1–3, 6–7) = follows as a structural consequence within the framework
- † **Conditional** (items 4–5) = derived given U(1) link variables as input (gauge group not derived)
- ‡ **Reframed** (items 8–9) = **not claimed as derivations**; conditional interpretations contingent on a future rigorous reconstruction mechanism (Section 6). Included to indicate explanatory direction, not to claim closure.

The BCB/admissibility perspective has been developed more broadly within the VERSEF programme; here we present a self-contained application focused specifically on photon propagation and discrete geometry selection.

2. Foundational Principles: BCB and Admissibility

For the general reader: Before diving into the technical argument, we need to establish the ground rules. The key idea is simple: *information processing has limits*. Just as a computer can only perform so many calculations per second, the universe can only "process" so much information in each moment. We call this **Bit Conservation & Balance (BCB)**.

Think of it like a budget: each tick of the cosmic clock has a fixed "information budget" that can't be exceeded. This seemingly innocent constraint has profound consequences—it determines what kinds of propagation are even *possible*.

2.1 Primitive Concepts

Before stating BCB, we must define the primitive concepts on which it rests. These primitives are pre-geometric—they do not presuppose spatial structure.

Definition 1 (Event): An event is an elementary occurrence that either happens or does not happen. Events are the atoms of the ontology.

Definition 2 (Causal Influence): Event A causally influences event B if and only if the occurrence of A is necessary for the occurrence of B. This defines a partial order on events.

Definition 3 (Tick): A tick is the minimal unit of causal update—the smallest "step" by which causal influence propagates. Each tick has bounded capacity for processing distinctions.

Remark (Tick invariance): A tick is not a coordinate time interval. It is a primitive causal update step of the substrate—an invariant count along causal influence chains. Different observers may assign different coordinate times to a sequence of ticks (time dilation), but the tick count is part of the underlying causal structure. The threshold Δi_{\min} is therefore frame-independent in the same sense that causal order and update count are invariant, even when coordinate descriptions differ.

Definition 4 (Distinguishability): Distinguishability is the capacity to resolve whether one state of affairs obtains rather than another. It is the fundamental resource of the framework.

Crucially, distinguishability is not the same as energy, though they are related:

Property	Energy	Distinguishability
Conservation	Conserved (1st Law)	Conserved (BCB)
Dilution behaviour	Can spread continuously	Cannot dilute below resolution threshold
Threshold effects	None (arbitrarily small energies meaningful)	Finite minimum per resolvable fact
Aggregation	Linear superposition	Subject to combinatorial constraints

Distinguishability is not "energy in bits". The framework's threshold is not a statement about the minimum energy of a propagating excitation. Quantum electrodynamics allows states with arbitrarily small field amplitude and arbitrarily small mean photon number. A coherent state with $\langle n \rangle \ll 1$ remains a perfectly valid quantum state that propagates under local field dynamics; the field amplitudes superpose and interfere, and detection outcomes are obtained only upon measurement.

By contrast, distinguishability in BCB is the resource required to **commit a binary fact**—to resolve "F is true rather than false" as an objective outcome of a causal update. A committed fact is not a small-amplitude excitation that can remain in superposition; it is a discrete ontological transition. Consequently, there is a meaningful threshold Δi_{\min} for fact-commitment: below this capacity, no definite bit can be committed at that location in that update.

This is why the dilution issue is specific to fact-resolution and does not apply to arbitrarily weak electromagnetic fields. Weak fields can propagate in 3D because propagation does not require committing independent binary facts at each point on the wavefront; it requires only local

evolution of amplitudes. The threshold arises when the ontology demands local, path-independent, finite-budget fact commitment—not when it evolves continuous (or quantized but superposable) dynamical degrees of freedom.

When does fact-commitment become relevant? Fact-commitment is required whenever the causal structure must resolve a definite outcome: a measurement result, a particle detection, a scattering event, or any interaction that leaves an irreversible trace. The interface framework asserts that *causal structure itself*—the network of "what influences what"—is built from committed facts, not from amplitude evolution. Electromagnetic field propagation, in this view, is an emergent effective description of underlying fact-commitment dynamics on the interface. The 3D Maxwell theory we observe is the continuum limit of 2D interface processes; it appears to propagate in 3D bulk because we reconstruct the bulk from interface correlations (Section 6). The constraint bites at the fundamental level, not at the level of effective field theory.

2.2 Bit Conservation & Balance (BCB)

The finitude-irreversibility connection: Before stating BCB, we establish why it is not an optional interpretive choice but a consequence of finite information capacity.

Reversibility—the requirement that every prior state be recoverable from the current state—demands:

1. No information ever discarded
2. No truncation at any precision level
3. Perfect fidelity preservation through all transformations

This is equivalent to *infinite information resolution*. Not as a contingent feature but as a necessary condition: if any distinction is ever lost, the prior state containing that distinction cannot be recovered.

But the physical world exhibits manifestly finite capacity:

- Finite entropy bounds (Bekenstein bound, holographic bounds)
- Finite channel capacities
- Finite energy densities
- Finite causal speeds
- Finite measurement resolution

The direct contradiction: Finite capacity and perfect reversibility are incompatible. If capacity is finite, some distinctions must be discarded at some step. Discarding distinctions *is* irreversibility. Irreversibility *is* fact-commitment—the transition from "could be A or B" to "is A."

This reframes the entire dialectic. We are not defending an unusual ontological choice. We are taking seriously what everyone already implicitly assumes when they say "the experiment showed X." The Everettian who insists on fundamental reversibility is actually committed to

infinite information capacity—the universal wavefunction preserving all phase correlations forever with perfect precision. *That* is the extraordinary claim, not BCB.

BCB is what finite capacity means: With this understanding, BCB is not a new postulate but the formalisation of finitude:

BCB Principle: Distinguishability—bit-level resolvable information—is neither freely created nor destroyed per causal update. The cost of resolving a given fact must be:

- **Bounded:** Each tick processes at most ΔI_{\max} distinguishability
- **Path-independent:** The total cost to resolve a fact between endpoints cannot depend on the causal path taken
- **Direction-independent:** No preferred directions exist for information propagation at the fundamental level

Violation of BCB implies:

- *Unbounded processing:* Infinite computation per finite time—physically unrealisable
- *Path-dependence:* The same physical fact has different "costs" depending on how you check it—inconsistent accounting
- *Direction-dependence:* Preferred directions without physical mechanism—unexplained anisotropy

Burden of proof inverted: The question is not "why accept fact-primitive ontology?" but "do facts exist at all, anywhere in your ontology?" If yes, BCB-type constraints apply at whatever level they enter. If no, you have abandoned the ordinary language of experimental physics—the language of "the detector clicked," "the measurement yielded," "the outcome was."

The Everettian says: "There are no facts, only relative states across branches." But then they must explain how the *appearance* of facts (detectors clicking, papers being written about outcomes) emerges from pure amplitude dynamics. Decoherence explains why interference vanishes; it does not explain why one outcome becomes *the* outcome. If "facts" enter even as emergent phenomena, the question becomes: at what level do they enter, and what are the constraints on that level?

A fully reversible, amplitude-primitive ontology implicitly assumes infinite information capacity. BCB is the alternative: finite capacity, forced irreversibility, primitive fact-commitment. The 6Ω prediction is the empirical discriminator between these positions.

2.2a Emergent Time and Fact-Commitment Are the Same Commitment

There is a deeper route to the same conclusion that does not begin with finite capacity but with the nature of time.

The block universe does not escape: One might think that if time is fundamental—a pre-existing dimension like space—then one can maintain "no facts, just states at coordinates t ." This

is the block universe: all times equally real, past and future on the same ontological footing, nothing specially "committed."

But this reinterpretation fails. A "fact" that could in principle be different at a later coordinate is not a fact in the operative sense—it is merely a configuration at a time. The word "fact" already contains irreversibility. To say "it is a fact that X" means X is settled, closed, not revisable by future evolution. A state that might be revised is not a fact; it is a possibility that happens to obtain at coordinate t.

The block universe doesn't eliminate facts; it either (a) denies that anything is ever settled (in which case it cannot account for records, memories, or the manifest difference between past and future), or (b) accepts that the block *contains* settled closure relations (in which case facts exist within it). Option (a) is empirically untenable; option (b) concedes the point.

If time is emergent, what does it emerge from?

Time emerges from sequential closure: Time emerges from "this happened, then that happened"—from the accumulation of settled matters. Each moment is not a location along a pre-existing axis but a *closure event*: something that was open becomes closed; something that could have been otherwise becomes fixed. The structure of time—before/after, earlier/later—is the structure of these closure relations.

If nothing ever became settled—if everything remained revisable superposition indefinitely—there would be no sequence, no before/after, no time. The "flow" of time is facts accumulating. The asymmetry of time (why we remember the past but not the future) is the asymmetry of closure (closed facts constrain; open possibilities do not).

Emergent time IS fact-commitment: These are the same commitment viewed from different angles:

Concept	Time perspective	Fact perspective
A tick	One step of temporal sequence	One closure event
Before/after	Earlier tick / later tick	Closed fact / fact depending on it
Arrow of time	Asymmetry of sequence	Asymmetry of closure
Present	The current closure frontier	Facts being committed now

You cannot have emergent time without fact-commitment. They are two descriptions of the same structure.

Implication for the paper's logic: If you accept emergent time (which most approaches to quantum gravity do—time is notoriously difficult to locate in the Wheeler-DeWitt equation, in causal sets, in loop quantum gravity), you have already accepted facts. The question is not *whether* facts exist but *what constrains their structure*.

Finite capacity then tells you facts are *rate-limited*—there is a maximum closure rate per region, a budget per tick. This gives BCB its quantitative bite. But the *existence* of facts is already settled the moment you say time emerges rather than being a pre-existing coordinate.

The logical chain can therefore be stated two ways:

From finite capacity: Finite capacity → forced discarding → irreversibility → facts → BCB constraints → 2D hexagonal

From emergent time: Emergent time → sequential closure → facts (same thing) → BCB constrains their structure → finite capacity gives rate limits → 2D hexagonal

Both routes arrive at the same destination. The framework is overdetermined by independent philosophical commitments that most physicists already accept.

2.3 The Distinguishability-Energy Distinction

A central claim of this framework is that distinguishability behaves differently from energy under geometric spreading. This distinction requires careful justification.

Energy dilution: In standard electromagnetism, energy spreads over expanding wavefronts. The energy density at radius r from a point source scales as:

$$\rho_E(r) \propto 1/r^2$$

This is unproblematic because energy is a continuous quantity with no minimum meaningful value. Arbitrarily small energy densities remain physically meaningful.

Distinguishability dilution: Distinguishability, by contrast, has a **resolution threshold**. To distinguish whether fact F obtains requires a minimum distinguishability budget Δi_{\min} . Below this threshold, the distinction cannot be resolved—it is not merely "weak" but *undefined*.

This threshold behaviour follows from the discrete nature of facts: either a bit is resolved or it is not. There is no "half-resolved" bit.

The dilution problem: Consider isotropic propagation from a point in D spatial dimensions. After one tick, causal influence reaches a shell with:

- $D = 1$: $O(1)$ elements (two endpoints)
- $D = 2$: $O(r)$ elements (circumference)
- $D = 3$: $O(r^2)$ elements (surface area)

If the per-tick budget is ΔI_{\max} , the distinguishability available per element is:

$$\Delta i(r) = \Delta I_{\max} / N_D(r)$$

For $D = 3$: $\Delta i(r) \propto 1/r^2$

At sufficiently large r , this falls below Δi_{\min} , and the propagation can no longer resolve facts—it becomes causally impotent.

Energy does not face this problem because it need not resolve discrete facts. A photon with energy 10^{-100} eV is still a photon (in principle). But a "bit" with 10^{-100} of the resolution capacity is not a bit at all.

BCB is not a competing field theory: The present framework does not modify the unitary propagation of QFT amplitudes, nor does it claim a new Hamiltonian for the electromagnetic field. Rather, it addresses the missing ontological layer in standard practice: how definite causal records arise at all. BCB is posed as a constraint on any physical mechanism of fact-commitment (collapse, record-primitive substrates, or other completions), and the non-admissibility theorem applies to that commitment layer—not to the unitary S-matrix computation of scattering probabilities. The distinguishing claim is empirical: if record-primitive causation is fundamental, micro-geometry should leave a symmetry-protected residue (6Ω); if amplitude-primitive completions erase all residue, no such harmonic should appear.

2.4 The Ontological Status of Δi_{\min}

The threshold Δi_{\min} is the framework's most basic assumption. We state its status explicitly and argue that such a threshold is *generic* to any ontology treating facts as stable physical records.

Definition (Minimal recordable distinction): Δi_{\min} denotes the minimal distinguishability required for a binary distinction to be *stably recorded* in the causal substrate—i.e., to persist as an irreversible constraint on future updates. Below this level, variations may exist as dynamical influence or amplitude-level correlations, but they do not constitute a record and therefore do not define a committed fact.

Operational characterisation (error tolerance): Fix a target stability level ε_{rec} (maximum probability per tick that a committed record spontaneously flips, erases, or decoheres into ambiguity). Δi_{\min} is the smallest commitment such that the substrate can implement a two-state record with error $\leq \varepsilon_{\text{rec}}$ under its ambient noise and perturbations. This makes Δi_{\min} a parameterised threshold defined by a stability requirement, not a mystical constant.

Why a threshold is generic

Any physical record must be robust to noise. Robustness requires a finite separation between alternative record states (energy barrier, free-energy gap, redundancy, error exponent, or equivalent). In a noisy substrate, record reliability typically exhibits an error-exponent form:

$$P_{\text{err}} \sim \exp(-\alpha R)$$

where R is the invested stabilisation resource (barrier height, redundancy, free-energy gap, or equivalent). Achieving $P_{\text{err}} < \epsilon$ therefore requires $R \gtrsim (1/\alpha) \log(1/\epsilon)$. This is the precise sense in which "zero resources cannot define a reliable record."

Hence any ontology that treats facts as physically instantiated records generically implies a nonzero minimal commitment scale, represented here by Δi_{min} . (This parallels Landauer's principle [16] connecting logical irreversibility to thermodynamic cost, though we do not claim numerical identity.)

Commitment vs dynamics

Δi_{min} is *not* a minimum field amplitude and is *not* a statement about energy quantization. Quantum fields can propagate in arbitrarily weak states (e.g., coherent states with $\langle n \rangle \ll 1$); smallness affects detection probability, not state validity. Δi_{min} instead concerns the transition from amplitude-level evolution to fact-commitment—the point at which "F rather than $\neg F$ " becomes a causal ancestor of subsequent events.

This connects directly to the quantum measurement problem: when do superpositions become definite facts? Our framework parametrises this transition via Δi_{min} without specifying the mechanism. Different interpretations of quantum mechanics (collapse theories, decoherence, many-worlds) could be mapped onto different accounts of how Δi_{min} is realised physically.

Ontological role

Δi_{min} is not an additional dynamical law; it is the formal statement that "fact" means "reliably recordable distinction" in a finite, noisy substrate. Without a nonzero minimal commitment, whether a distinction counts as a fact would depend on the observer's coarse-graining or noise floor, making ontology gauge-like—a mere labelling convention rather than a substrate property.

Epistemic vs ontological thresholds: A skeptic might object: "Any measuring device has a detection threshold—that's epistemic, not ontological. What makes Δi_{min} different?" The distinction is this: epistemic thresholds concern what observers can *detect* given their apparatus; ontological thresholds concern what the substrate can *commit* as a fact independent of observation. A weak signal below a detector's threshold still exists as a physical state; a distinction below Δi_{min} does not exist as a committed fact in the causal structure. The former affects what we know; the latter affects what is true. This is precisely the distinction between amplitude-primitive and fact-primitive ontologies (Appendix D). We do not claim to prove Δi_{min} is ontological rather than epistemic—that is the interpretive commitment of the framework. The 6Ω prediction provides one route to test it empirically.

Tick and invariance

Because a tick is a primitive causal update step (not coordinate time), Δi_{min} is defined per causal update and is frame-independent in the same sense as causal order and update count. This parallels how causal set theory defines temporal order via the causal partial order without assuming pre-existing time.

Minimality without numerology

The results of Sections 3–4 require only $\Delta i_{\min} > 0$ and $\Delta I_{\max} < \infty$. As shown in Section 2.4.1, $\Delta i_{\min} = 1$ causal bit—the minimal binary constraint on future evolution. This is not a new physical constant but the logical unit of fact-commitment determined by the binary nature of records. What remains undetermined is the *physical implementation cost*: the energy, time, or substrate resources required to commit one causal bit. This cost depends on microphysics (barrier heights, redundancy, error correction) and may vary across substrates. The interface scale a and suppression parameter ε encode this implementation-level physics.

Indirect experimental constraints: Although we cannot measure Δi_{\min} directly, the framework's predictions (Section 8) provide indirect constraints. Null results bound the combination $\varepsilon \cdot a^p$, which constrains the interface scale a . Since a sets the scale at which resolution thresholds become relevant, improved experimental sensitivity progressively narrows the viable parameter space.

Summary: Δi_{\min} is not postulated as a new physical constant. It is the minimal consequence of treating causal influence as the propagation of distinguishability in a finite, local system. Any framework that denies Δi_{\min} must either (i) abandon operational causality, (ii) allow infinite information density, or (iii) treat records as purely emergent bookkeeping. The present work explicitly explores the complementary regime.

2.4.1 The Minimum Recordable Distinction

We now sharpen the status of the resolution threshold Δi_{\min} . The argument does not rely on information-theoretic entropy, probability measures, or topological quantization. It rests solely on the operational meaning of records in a causal system.

Records as binary constraints

A record is defined as a physical structure whose presence or absence constrains future evolution in a counterfactual sense. That is, there exist physically admissible interventions such that downstream outcomes differ depending on whether the record is present.

This definition immediately implies that records are *binary* objects. Either the constraint exists or it does not. A putative "fractional" record that does not condition future evolution into at least two distinct equivalence classes is not a record at all, but merely a sub-threshold dynamical influence.

Thus, the minimum recordable distinction corresponds to the smallest nontrivial partition of future evolution into two classes. We denote this minimal unit as **one causal bit**.

Causal bits vs information-theoretic bits

The term "bit" here does not refer to Shannon entropy, von Neumann entropy, or the information content of a probability distribution. A causal bit is not a measure of uncertainty or surprise; it is

a unit of *constraint*. It specifies whether a particular fact has been committed into the causal structure.

A causal bit may require arbitrarily large or small physical resources to implement, depending on substrate, redundancy, and error tolerance. The value $\Delta i_{\min} = 1$ causal bit therefore fixes only the *logical unit* of recordability, not its energetic or thermodynamic cost.

Logical vs physical threshold: "One causal bit" is the logical minimum for a record (a binary constraint that partitions future evolution). The substantive physical claim is that committing even this minimal record requires nonzero resources in any finite, noisy substrate (stability against perturbations), hence $\Delta i_{\min} > 0$ in the ontological sense. This separates the definitional point (records are binary) from the physical point (implementing a binary record is not free).

Necessity of a threshold

Because records must be stable against perturbations, committing a causal bit requires a nonzero resource investment (as established in Section 2.4). Sub-threshold influences may bias probabilities or amplitudes, but they do not constitute records and cannot function as causal antecedents. Consequently:

$\Delta i_{\min} = 1$ causal bit

This conclusion is not a new physical postulate. It is a direct consequence of treating causation as operationally meaningful and records as physical constraints on future evolution.

Why one bit, not two? A skeptic might ask: why couldn't a substrate require, say, two redundant bits minimum for stability? The answer is that redundancy is already incorporated into the definition of a "record-capable register" (Section 3.2). A register is defined as the smallest degree of freedom that can support a stable binary record at the required reliability level ϵ_{rec} . If stability requires internal redundancy, that redundancy is *inside* the register—it's part of what makes the register record-capable. The threshold $\Delta i_{\min} = 1$ causal bit is therefore the *effective* minimum per register, after whatever internal structure is needed for stability. Different substrates may have different physical implementations of a register (with different internal redundancies), but by definition each register can commit exactly one causal bit when it receives threshold resources.

Implication for propagation

With Δi_{\min} fixed as one causal bit, the arguments in Section 3 become strict counting arguments. If the distinguishability available per addressable register falls below one causal bit, no new records can be committed at that location. Propagation then fails *categorically*, not gradually. The special role of two-dimensional interfaces follows directly: 2D is the unique dimensionality in which the budget-per-register can be maintained at one causal bit indefinitely.

Note: The emergence of integer-valued record units invites comparison with topological quantization in gauge theories (where flux is quantized due to compact gauge groups). While

such analogies may prove fruitful in future work, no topological structure is assumed or required for the present results.

Box: Causal Structure Requires Stable Records (Formal Argument)

The following theorem shows that Δi_{\min} is not an invented constant but a *necessary consequence* of physically meaningful causation.

Definition (Physical causation): $A < B$ is physically meaningful if it supports counterfactual dependence: there exists at least one physically admissible intervention on A that would (with nonzero probability) change the probability distribution of B, holding fixed the rest of the relevant past.

Definition (Stable record): A record R of event A is *stable* if its correlation with A persists long enough to constrain later events B—i.e., it is robust to typical perturbations over the timescale required to influence downstream events.

Theorem (Record Necessity for Physical Causation): If $A < B$ is physically meaningful, then there must exist at least one stable record R that carries information about A forward to the causal domain of B.

Proof sketch:

1. **Counterfactual meaning requires a mechanism:** If changing A can change the distribution of B, some physical influence of A must reach B's neighbourhood.
2. **Influence without retained information cannot systematically alter downstream events:** If all consequences of A are instantly thermalised or decoupled before reaching anything that interacts with B, then interventions on A cannot systematically alter B. A may precede B in time, but it does not *cause* B in the counterfactual sense.
3. **Therefore a carrier of information must exist:** There must be intermediate degrees of freedom whose states depend on A and which influence B.
4. **A carrier that is not stable cannot support causal constraint:** If correlation with A decays to zero before reaching B's domain, B becomes insensitive to interventions on A. Persistence across at least one causal step is precisely stability.
5. **Thus a stable record exists,** mediating A's causal influence on B. ■

Corollary (Stable Records Imply Minimal Commitment Threshold): A "record" that can be overwritten by arbitrarily small perturbations at arbitrarily low cost is not stable. Therefore, stable records require nonzero resource separation (barrier, redundancy, error exponent). That nonzero separation is exactly what Δi_{\min} represents.

Conclusion: Δi_{\min} is not "invented"—it is the minimal requirement for causal structure to be physically meaningful. Any framework with primitive causality implies $\Delta i_{\min} > 0$.

Pre-empting the "no facts, only amplitudes" objection: A referee might argue that fundamental dynamics is purely unitary amplitude evolution with no facts until measurement. Our response: if so, then physical causation is emergent too, because the language of "A causes B" already presupposes facts and records. This paper explicitly explores the alternative hypothesis—that causality is primitive (facts/records are in the substrate). The 6Ω prediction then discriminates empirically between "primitive substrate with residual lattice symmetry" and frameworks enforcing exact Lorentz invariance with no lattice residue. We do not claim to defeat amplitude-only interpretations; we show that *if* you want primitive causality, stable records are not optional.

2.5 Admissibility

Definition 5 (Admissibility): A propagation process is *admissible* if and only if it satisfies:

- (a) **Finite budget:** Processes at most ΔI_{\max} distinguishability per tick
- (b) **Locality:** No instantaneous global broadcasts
- (c) **Path-independence:** Resolution cost for any fact is path-independent
- (d) **Isotropy:** No preferred directions at macroscopic scales

2.6 Isotropy as a BCB Consistency Condition

Isotropy is often assumed as a symmetry principle. Here we show it is a *consistency condition* within BCB—a no-arbitrage requirement—rather than claiming to derive it from first principles alone.

Theorem 1 (Isotropy as No-Arbitrage Condition): Under BCB, any admissible substrate that supports route-independent fact-commitment must lie in (or flow toward) an isotropic universality class at macroscopic scales.

Proof sketch: Suppose propagation is anisotropic—faster or "cheaper" in some directions than others. Consider two paths P_1 and P_2 connecting the same endpoints A and B, with P_1 along a "cheap" direction and P_2 along an "expensive" direction.

The total distinguishability cost to resolve a fact about B given A is:

$$C(P_1) \neq C(P_2)$$

This violates path-independence. The same fact—"what is the state at B given A?"—has different resolution costs depending on the path chosen.

This creates **distinguishability arbitrage**: systematic differences in commitment cost by direction would allow "cheaper" resolution of facts along preferred directions, violating the principle that fact-commitment cost is intrinsic to the fact, not the route. Under BCB with path-independence, this is inconsistent.

Therefore, admissible substrates must have equal propagation cost in all directions at macroscopic scales—which is isotropy. ■

What this argument establishes: The theorem shows that anisotropy + path-independent fact-commitment cannot coexist without contradictions. It does *not* show that anisotropic worlds are logically impossible—only that they are outside the BCB-admissible class.

Empirical isotropy as evidence: Observed vacuum isotropy therefore functions as evidence that the physical substrate sits near an isotropic fixed point of the admissible class. BCB explains *why* that fixed point is singled out among admissible substrates, not *why* anisotropic worlds are metaphysically impossible.

Clarification on logical status: This is not circular reasoning but mutual constraint. BCB-consistency requires isotropy; observed isotropy is consistent with (though not uniquely explained by) BCB. Alternative explanations exist (e.g., Lorentz symmetry as a fundamental postulate). What distinguishes BCB is that it also yields the 6Ω residual signature—a prediction not shared by frameworks that simply postulate exact isotropy. The mutual support is therefore testable.

3. The Non-Admissibility Theorem

For the general reader: This is the heart of the paper. We're going to prove that light *cannot* fundamentally propagate through 3D space the way we usually imagine—it would violate the information budget.

Here's the intuition: Imagine you're a radio tower broadcasting a signal. As your signal spreads outward, it forms an expanding sphere. The surface area of that sphere grows as the *square* of the distance (that's why we have the inverse-square law for light intensity). But your broadcasting power is fixed—you only have so much energy to transmit each second.

Now replace "energy" with "information capacity." If the universe has a fixed information budget per moment, and light spreading in 3D has to "inform" an ever-growing spherical surface, eventually the information gets spread too thin to actually *commit* any definite facts about what's happening at each point. The math shows this breaks down for 3D, but *not* for 2D, where the "surface" only grows linearly with distance.

3.1 Purpose and Scope

This section formalises the central claim: **primitive three-dimensional bulk propagation is non-admissible**. The result is conditional on BCB/admissibility constraints.

Critical caveat: The argument's force depends entirely on accepting that fundamental propagation requires local fact-commitment at wavefront elements. If one rejects this—arguing

that propagation involves only amplitude evolution with fact-commitment occurring only at measurement—then the dilution argument does not bite. We frame the result conditionally:

If causal structure is built from committed facts (not just amplitude correlations), *and* fact-commitment requires a minimum distinguishability threshold, *then* primitive 3D propagation is non-admissible.

However, as established in Section 2.2, this conditional is not as optional as it might appear. Accepting finite information capacity immediately forces irreversibility, and irreversibility immediately forces fact-commitment at some level. The only way to avoid BCB constraints is to accept infinite information capacity—the universal wavefunction preserving all phase correlations forever with perfect precision. That is the extraordinary claim, not finite-capacity fact-commitment.

The reversibility argument (restated): Reversibility requires that every prior state be recoverable from the current state. This demands no information ever be discarded, no truncation at any precision level, and perfect fidelity preservation through all transformations. That is equivalent to infinite information resolution—infinite capacity.

But the physical world manifestly has finite capacity (Bekenstein bounds, finite entropy, finite causal speeds). Therefore, some distinctions must be discarded somewhere. Discarding distinctions *is* irreversibility. Irreversibility *is* fact-commitment.

The question is not "should we accept fact-primitive ontology as an interpretive choice?" but "do facts exist at all, anywhere in your ontology?" If yes, BCB-type constraints apply at whatever level they enter. If no, you have abandoned the ordinary language of experimental physics.

The Everettian position—"there are no facts, only relative states across branches"—must then explain how the *appearance* of facts emerges. Decoherence explains why interference vanishes; it does not explain why one outcome becomes *the* outcome. The Everettian who insists on fundamental reversibility is committed to infinite information capacity: the universal wavefunction containing every branch with every phase correlation preserved forever at infinite precision. *That* is the radical position. BCB takes finite capacity seriously and follows the consequences.

Interpretive commitment reframed: The framework does not assume fact-commitment as an "interpretive choice that could go either way." It takes finite capacity seriously. Given finite capacity, irreversibility follows. Given irreversibility, fact-commitment follows. The only remaining question is *where* in the physical hierarchy fact-commitment enters. We argue it is fundamental (substrate-level), leaving a symmetry-protected residue (6Ω). The alternative—that facts are purely emergent from infinite-precision amplitude dynamics—requires infinite capacity to be coherent. The 6Ω prediction tests which position is correct.

We emphasise that this does not claim Maxwell electrodynamics is wrong. Maxwell theory describes emergent, effective behaviour with extraordinary accuracy. The theorem constrains the *ontological substrate*, not the effective equations.

3.2 Formal Setup

Definition 6 (Primitive D-dimensional propagation): Propagation is *primitive* (i.e., ontologically fundamental) in D dimensions if D-dimensional spatial structure is fundamental rather than emergent—if causal updates occur natively in a D-dimensional arena.

Definition 7 (Resolution threshold): There exists a minimum recordable distinguishability Δi_{\min} corresponding to one *causal bit*—the smallest binary constraint capable of conditioning future evolution (see Section 2.4.1). Below this threshold, no fact is resolved; influences may exist but do not constitute committed records.

Note: This causal bit need not correspond to one Shannon bit or to a unit of thermodynamic entropy. It is a unit of causal constraint, not information-theoretic uncertainty. The physical resources required to implement one causal bit depend on substrate microphysics.

Definition 8 (Addressable Causal Register): The causal substrate is discretised into minimal *addressable causal registers*—the smallest units that can receive and commit a binary distinction in one tick. Each tick can update at most $O(1)$ registers locally per source site. The "number of elements" in subsequent lemmas refers to the count of addressable registers intersecting a region.

Operational clarification: A register is defined as the smallest local degree of freedom that can support a stable binary record at the target reliability level ϵ_{rec} . Thus Δi_{\min} is the minimal commitment required per register-update, while $N_D(r)$ counts how many such record-capable registers lie on the wavefront. Registers need not be "Planck cells"—they are defined by the stability requirement, not by a pre-assumed length scale.

Reframing the register count: The count $N_D(r)$ is not derived from pre-assumed microphysics; it represents the maximum number of independent binary constraints that could be committed at the wavefront under optimal partitioning of the available budget. The non-admissibility argument then shows that even this maximum is insufficient in $D \geq 3$: the budget-per-register falls below any fixed threshold at large r . The argument does not require knowing the precise register structure—only that partitioning into more registers than the budget can support yields no benefit, since sub-threshold registers cannot commit facts. Whatever the actual register structure turns out to be, 3D fails.

Definition 9 (Coarse-grained wavefront): A substrate admits a *coarse-grained wavefront* if, at the level where "propagation" is operationally defined, causal influence from a localised source reaches an expanding set of addressable registers whose cardinality scales like the measure of an effective $(D-1)$ -dimensional surface. This does not require a rigid lattice—disordered or fluctuating substrates qualify provided the effective wavefront scaling is well-defined after appropriate averaging. Substrates that do not admit well-defined wavefronts (e.g., causal sets with random Poisson sprinkling) fall outside the scope of the non-admissibility theorem.

Note on "appropriate averaging": The phrase "appropriate averaging" is an auxiliary assumption whose precise content depends on the substrate class. For deterministic substrates with local disorder (e.g., random bond strengths on a fixed lattice), averaging means taking the large- r limit

where local fluctuations smooth out. For stochastic substrates with fixed ensemble statistics, averaging means ensemble expectation. The assumption is that $N(r)$ has a well-defined leading-order scaling $N(r) \sim c \cdot r^{(D-1)}$ with $c > 0$ and D a well-defined integer (or effective dimension). Substrates for which this scaling does not exist, or fluctuates wildly, are outside the theorem's scope. This auxiliary assumption is listed in Appendix E.2.

Operational definition of dimension: We do not assume metric geometry; we assume an operational notion of "distinct reachable locations after r updates." The effective dimension D is defined by the scaling of the number of newly reachable distinguishable locations after r updates: $N(r) \propto r^{(D-1)}$. This is not assumed geometry but the definition of effective propagation dimensionality.

3.3 Lemma 1: Wavefront Scaling

Lemma 1: Under primitive isotropic propagation in D spatial dimensions, the number of addressable causal registers reached after propagation to radius r scales as:

$$N_D(r) \propto r^{(D-1)}$$

Proof: Isotropic propagation from a point reaches a $(D-1)$ -dimensional shell at radius r . The number of addressable registers is proportional to the shell's $(D-1)$ -volume, which scales as $r^{(D-1)}$. ■

3.4 Lemma 2: Budget Dilution

Lemma 2: If the per-tick distinguishability budget is ΔI_{\max} and propagation must maintain resolvability across all reached elements, then primitive 3D propagation eventually fails the resolution threshold.

Proof: In 3D, $N_3(r) \propto r^2$. The distinguishability per element is:

$$\Delta i(r) = \Delta I_{\max} / N_3(r) \propto 1/r^2$$

For any fixed $\Delta i_{\min} > 0$ and $\Delta I_{\max} < \infty$, there exists r^* such that:

$$\Delta i(r) < \Delta i_{\min} \text{ for all } r > r^*$$

Beyond r^* , propagation cannot resolve facts at individual elements. ■

Box 1: Accumulation Does Not Save Primitive 3D Propagation

A natural objection is that propagation proceeds over many ticks, so one might "accumulate" distinguishability over time and thereby avoid dilution. BCB allows accumulation along a causal process, but admissibility requires that commitment capacity per tick remains finite and local.

Pre-positioning is nonlocal: One might suggest "pre-allocating" distinguishability to distant registers in anticipation of a future emission. But this would require those distant registers to already depend on the source's future choice to emit (or on the emission's content) before any causal influence has reached them. That is equivalent to a superluminal broadcast or retrocausal coordination, and therefore violates locality/causality—the very structure the framework is analysing.

After r ticks, the maximum accumulated committed distinguishability that can originate from a localised source scales at most linearly:

$$I_{\text{acc}}(r) \leq r \cdot \Delta I_{\text{max}} \propto r$$

Under primitive isotropic propagation in D spatial dimensions, the number of newly addressable registers on the wavefront scales as:

$$N_D(r) \propto r^{(D-1)}$$

Hence the available committed distinguishability per newly reached register satisfies:

$$\Delta i_{\text{eff}}(r) \lesssim I_{\text{acc}}(r) / N_D(r) \propto r / r^{(D-1)} = r^{-(D-2)}$$

For $D = 3$: $\Delta i_{\text{eff}}(r) \propto 1/r$, which still falls below any fixed threshold Δi_{min} at sufficiently large r . Accumulation therefore does not remove the failure radius in 3D; it only weakens the divergence from $1/r^2$ (single-tick) to $1/r$ (multi-tick), which still vanishes.

For $D = 2$: $\Delta i_{\text{eff}}(r) \propto r^0 = \text{constant}$. This is why the 2D interface is special: linear accumulation matches linear wavefront growth, so admissible fact-commitment can remain viable at all scales.

Note on extended sources: The above assumes localised (point-like) sources. What if budget from a spatially extended source region all converges on the wavefront? Extended sources face the same constraint applied to each localised sub-source independently. Budget cannot be "borrowed" from causally disconnected regions of the source without violating locality. The total budget arriving at the wavefront from an extended source of size L is at most proportional to $L^{(D-1)} \cdot r$ (the source volume times ticks), while the wavefront area grows as $r^{(D-1)}$. For $D = 3$, this still yields per-register budget $\propto L^2/r$, which vanishes at large r for any finite source size.

Note on selective accumulation: A further objection might be: what about concentrating the accumulated budget on a *subset* of wavefront registers rather than distributing it uniformly? This strategy could maintain threshold resolvability for the selected subset while abandoning the rest. However, any such strategy violates isotropy: the subset of registers receiving sufficient budget defines preferred directions in which facts can be committed, while other directions remain below threshold. The trilemma (Lemma 3) below formalises this: maintaining resolvability for all directions in 3D is impossible without violating one of the admissibility constraints.

3.5 Lemma 3: The Trilemma

Lemma 3: To maintain resolvability beyond r^* in primitive 3D propagation, at least one of the following must be violated:

- (a) **Locality:** Propagation requires global coordination to concentrate budget
- (b) **Path-independence:** Different paths to the same endpoint yield different resolution costs
- (c) **Isotropy:** Some directions receive preferential budget allocation

Proof:

The budget shortfall at $r > r^*$ is:

$$\Delta I_{\text{deficit}}(r) = N_3(r) \cdot \Delta i_{\text{min}} - \Delta I_{\text{max}} \propto r^2 - \text{const} > 0$$

This deficit must be covered somehow:

Option (a): Draw budget from distant regions to supplement local propagation. This requires nonlocal coordination—knowing where to draw from before the information could have propagated there. Violates locality.

Option (b): Accept that resolution cost depends on which path information took to reach a point. Violates path-independence.

Option (c): Allocate budget preferentially to some directions, leaving others under-resolved. Violates isotropy.

No other options exist within the framework's constraints. ■

3.6 Lemma 4: Interface Propagation Avoids the Trilemma

Lemma 4: On a 2D interface, isotropic propagation maintains resolvability without violating locality, path-independence, or isotropy.

Proof: In 2D, $N_2(r) \propto r$. The distinguishability per register is:

$$\Delta i(r) = \Delta I_{\text{max}} / N_2(r) \propto 1/r$$

This is a slower dilution than 3D. More importantly, the *total distinguishability* needed to maintain resolution scales as:

$$I_{\text{total}}(r) = N_2(r) \cdot \Delta i_{\text{min}} \propto r$$

For propagation that has processed r ticks, the accumulated budget is:

$$I_{\text{accumulated}}(r) = r \cdot \Delta I_{\text{max}} \propto r$$

The accumulated budget scales identically with the required budget. There is no r^* beyond which propagation fails—the system remains balanced at all scales.

Why 3D cannot use the same accumulation argument: One might ask whether 3D propagation could similarly accumulate budget over multiple ticks. The answer is no. In 3D, even allowing accumulation over r ticks, the available committed distinguishability grows as r (one tick's budget times r ticks), while the number of newly addressable registers on the wavefront grows as r^2 . Therefore, per-register commitment still decays as $r/r^2 = 1/r$, and there remains a failure radius r^* where this falls below threshold.

The 2D case is special precisely because the wavefront growth rate ($\propto r$) matches the accumulation rate ($\propto r$), maintaining constant per-register commitment indefinitely.

Furthermore, hexagonal coordination (derived below) minimises angular variance in budget distribution, ensuring no directional preference in how the budget is allocated.

Note: This lemma shows that 2D propagation *can* maintain resolvability, not that it automatically does. Isotropy in 2D still requires the interface to have an isotropic discrete structure—which is not guaranteed for arbitrary 2D tilings. Section 4 derives that hexagonal coordination is uniquely selected among regular tilings to achieve this. Lemma 4 is therefore contingent on the hexagonal selection result. ■

3.7 Main Theorem

Theorem 2 (Non-Admissibility of Primitive 3D Propagation): Under BCB and admissibility constraints, primitive three-dimensional bulk propagation of light is non-admissible.

Proof: By Lemma 2, primitive 3D propagation eventually fails the resolution threshold. By Lemma 3, maintaining resolvability requires violating locality, path-independence, or isotropy. Each violation contradicts admissibility. Therefore, primitive 3D propagation is non-admissible. ■

Corollary: Fundamental causal propagation must occur on a substrate of dimension $D \leq 2$. The minimal dimension supporting isotropic propagation is $D = 2$ (a 1D substrate cannot support isotropy). Therefore, **the causal substrate is two-dimensional**.

Scope clarification: The theorem rules out $D \geq 3$ under the stated assumptions; isotropy excludes $D \leq 1$; therefore $D = 2$ is the only admissible integer dimension. Substrates with non-integer effective dimension (fractals) would require separate analysis and are not claimed to be excluded here.

Note on collective modes: The theorem does not assume field modes decompose into independent local bits at the amplitude level; it assumes that records are locally instantiated constraints, and counts the minimal record-capable registers required for causal commitment. Collective or delocalised modes may propagate freely; the constraint applies only when definite causal facts must be committed at distinguishable locations.

Collective modes do not escape the budget constraint: A potential objection is that QFT modes are delocalised, so local budget constraints cannot apply. However, collective modes relocate the problem rather than solving it. A record instantiated across N correlated degrees of freedom requires N -way coordination. The budget for achieving such correlation must still propagate from the source through the causal structure. The 3D scaling problem reappears as: available correlation budget grows as r (ticks \times rate), while the number of degrees of freedom requiring correlation grows as r^3 (volume) for bulk records or r^2 (surface) for wavefront records. For fact-commitment involving distant degrees of freedom, the deficit is even worse than the local-register argument suggests.

Box 2: Single-Tick Distinguishability Budget on Expanding Wavefronts

Setup: A point source emits a causal signal. After r ticks, the wavefront reaches radius r .

Parameter	Meaning
ΔI_{\max}	Maximum distinguishability budget per tick
Δi_{\min}	Minimum distinguishability to commit one binary fact
$N_D(r)$	Number of addressable registers on wavefront at radius r

Wavefront scaling by dimension:

- 1D: $N_1(r) = 2$ (two endpoints)
- 2D: $N_2(r) = 2\pi r/a \propto r$ (circumference)
- 3D: $N_3(r) = 4\pi r^2/a^2 \propto r^2$ (surface area)

Per-register distinguishability: $\Delta i(r) = \Delta I_{\max} / N_D(r)$

Accumulated budget after r ticks: $I_{\text{acc}}(r) = r \cdot \Delta I_{\max}$

Required budget to resolve all registers: $I_{\text{req}}(r) = N_D(r) \cdot \Delta i_{\min}$

Dimension	Per-register $\Delta i(r)$	$I_{\text{acc}}(r)$	$I_{\text{req}}(r)$	Ratio $I_{\text{acc}}/I_{\text{req}}$	Status
2D	$\propto 1/r$	$\propto r$	$\propto r$	constant	✓ Balanced
3D	$\propto 1/r^2$	$\propto r$	$\propto r^2$	$\propto 1/r \rightarrow 0$	✗ Fails at large r

The key insight: In 2D, budget accumulation keeps pace with wavefront growth. In 3D, wavefront growth outpaces accumulation—there is always a radius r^* beyond which propagation cannot commit facts to all registers.

Numerical example: Let $\Delta I_{\max} = 100$ units/tick, $\Delta i_{\min} = 1$ unit.

- 2D at $r = 1000$: $N_2 \approx 6280$, $I_{\text{acc}} = 100,000$, $I_{\text{req}} = 6280$. Ratio ≈ 16 . ✓

- 3D at $r = 1000$: $N_3 \approx 12.6 \times 10^6$, $I_{\text{acc}} = 100,000$, $I_{\text{req}} = 12.6 \times 10^6$. Ratio ≈ 0.008 . χ

The 3D system is $125\times$ short of required budget at $r = 1000$, and the deficit grows with r .

3.8 Addressing the Energy Analogy Objection

Objection: "Energy dilutes as $1/r^2$ in 3D electromagnetism without any inconsistency. Why can't distinguishability do the same?"

Response: The analogy fails because energy and distinguishability have fundamentally different threshold structures:

1. **Energy has no resolution threshold:** A wave with energy density 10^{-100} J/m³ is still a wave. It carries real physical energy that can, in principle, be detected and accumulated.
2. **Distinguishability has a fact-commitment threshold:** A sub-threshold update may carry dynamical influence (e.g., amplitude-level correlations), but it cannot by itself commit a binary distinction as an objective fact. The threshold concerns the minimum resource needed to make "F rather than $\neg F$ " true in the ontology at that register.
3. **This is not about energy quantization:** Even though energy exchange occurs in quanta, quantum fields admit states with arbitrarily small amplitudes and arbitrarily small mean photon number (e.g., coherent states with $\langle n \rangle \ll 1$). Such states remain valid and propagate locally; their smallness affects detection probability, not the existence of the state. Fact-commitment is different: there is no analogue of a "half-committed fact in superposition" that already counts as a resolved bit.
4. **Commitment vs detectability:** Long integration and matched filtering increase *detectability* of weak fields, but detectability is an epistemic property of observers. BCB concerns an *ontological* constraint: whether the substrate can locally commit independent binary distinctions under finite per-update capacity without introducing nonlocal coordination or path-dependent cost.

The relevant threshold is not "detectability with unlimited repetition" but **single-tick fact-commitment capacity**. Sub-threshold contributions may exist as analog influence, but they do not constitute a resolved fact until sufficient distinguishability is concentrated within the same admissible causal update. The 3D scaling problem persists even with multi-tick accumulation (see Box 1).

Toy model: what goes wrong in an amplitude-only world? Consider a concrete scenario: a source emits a signal that must determine "which detector clicked" at a distant location. In the amplitude-only picture, the signal spreads as a superposition to all detectors, with amplitudes ψ_i that can be arbitrarily small. When "measurement" occurs, one detector registers a click.

The question is: *when and where* does the fact "detector 3 clicked rather than detector 7" become true? In amplitude-only frameworks, this fact is either (i) never "true" in a fundamental sense

(Everett: all detectors click in different branches), or (ii) determined by hidden variables that were already correlated with the outcome (pilot-wave).

In contrast, the fact-commitment picture requires that at some point, the substrate *locally* commits "detector 3" as the outcome. This commitment requires sufficient distinguishability to reach detector 3's location. If distinguishability dilutes below threshold before reaching distant detectors, the system cannot commit definite outcomes there—not because of observer limitations, but because the substrate lacks the resources to make the fact true.

This is observable in principle: if fact-commitment has a threshold, then sufficiently weak signals to sufficiently distant detectors should show anomalous statistics—not just low detection probability, but *failure to establish definite outcomes*. The 6Ω prediction is one indirect probe of this regime.

4. Hexagonal Selection

For the general reader: If light propagates on a 2D surface, what shape should that surface's underlying grid be? You might think this is arbitrary, but it turns out the universe is very picky.

Think about tiling a bathroom floor. You can use squares, triangles, or hexagons—these are the only regular shapes that tile perfectly. Now imagine signals spreading across each type of tiled floor. On a square grid, signals naturally travel faster along the grid lines (horizontally and vertically) than diagonally. This creates a subtle "preferred direction" problem.

Hexagonal grids are special: they're the *most isotropic* regular tiling. Honeybees discovered this for structural efficiency; we're discovering it for information-theoretic consistency. Four convergent mathematical arguments all point to hexagons as the unique solution. When four different paths lead to the same answer, that answer is probably not a coincidence.

4.1 The Selection Problem

Given that propagation must occur on a 2D interface, which discrete geometry should the interface adopt? Among the three regular planar tilings (triangular, square, hexagonal), we show that hexagonal coordination (equivalently, triangular lattice / hex-Voronoi dual) is uniquely selected.

We present **four convergent selection channels**, each arriving at the same conclusion through different reasoning.

4.2 Selection Channel 1: Dispersion Symmetry Analysis

Consider a translation-invariant, local propagation operator on a 2D interface with nearest-neighbour coupling. The linearised dynamics are governed by a discrete Laplacian with Fourier-space symbol:

$$\Lambda(\mathbf{k}) = \sum_n [1 - \cos(\mathbf{k} \cdot \mathbf{a}_n)]$$

where $\{\mathbf{a}_n\}$ are neighbour vectors of equal length.

Note on nearest-neighbour assumption: We restrict to nearest-neighbour coupling because (i) locality is a BCB requirement—longer-range couplings would allow "shortcuts" that violate path-independence, and (ii) next-nearest-neighbour couplings introduce additional parameters without changing the qualitative selection result. The hexagonal advantage persists for any finite-range local coupling.

Finite-range locality does not change the harmonic order: The nearest-neighbour restriction is a simplifying choice, not a logical requirement. Let the linear propagation operator be any translation-invariant, finite-range coupling on the interface lattice: $L\psi(\mathbf{x}) = \sum_{\mathbf{r} \in \mathbf{R}} w_{\mathbf{r}} [\psi(\mathbf{x} + \mathbf{r}) - \psi(\mathbf{x})]$ with \mathbf{R} a finite set of displacement vectors and weights $w_{\mathbf{r}}$. Its Fourier symbol is $\Lambda(\mathbf{k}) = \sum_{\mathbf{r} \in \mathbf{R}} w_{\mathbf{r}} [1 - \cos(\mathbf{k} \cdot \mathbf{r})]$, which admits a small- \mathbf{k} expansion whose anisotropic terms are constrained solely by the lattice symmetry group. Finite-range couplings can change the coefficients c_2, c_4, c_6, \dots but not the first symmetry-allowed angular harmonic: on a square lattice the leading anisotropy generically appears as $\cos(4\theta)$, whereas on a triangular/hex lattice the first nonvanishing anisotropy is $\cos(6\theta)$. Thus "nearest-neighbour" affects amplitudes but not the qualitative selection (order and symmetry class of anisotropy), provided the coupling remains finite-range, translation-invariant, and isotropic in the large.

For small $|\mathbf{k}|$, expand:

$$\Lambda(\mathbf{k}) = c_2|\mathbf{k}|^2 + c_4|\mathbf{k}|^4 + d_4|\mathbf{k}|^4 \cos(m\theta) + O(|\mathbf{k}|^6)$$

where θ is the propagation direction and m reflects lattice symmetry.

Analysis by lattice type:

Lattice	Coordination	First anisotropic term	Order
Square	4	$\cos(4\theta)$	
Triangular/Hex	6	$\cos(6\theta)$	

Result: Hexagonal coordination pushes anisotropy to sixth order, two orders higher than square coordination. The leading isotropic behaviour extends further into the UV.

4.3 Selection Channel 2: Renormalization Group Flow

Under coarse-graining by factor $b > 1$, wavevectors rescale as $\mathbf{k} \rightarrow \mathbf{k}/b$. The dimensionless anisotropy coupling transforms as:

$$g_{\text{aniso}}(b) \propto b^{-(p-2)}$$

where p is the order of the first anisotropic term.

Comparison:

Lattice	p	Anisotropy decay
Square	4	$g \sim b^{-2}$
Triangular/Hex	6	$g \sim b^{-4}$

Result: Anisotropy decays quadratically faster under RG flow for hexagonal coordination. The hexagonal lattice is an **RG-accelerated isotropy fixed point**—it flows toward perfect isotropy more rapidly than any other regular tiling.

4.4 Selection Channel 3: Information-Theoretic Optimality

In one interface tick, a node can causally influence only its nearest neighbours. This defines discrete angular sampling of outgoing causal directions.

Angular resolution analysis:

Lattice	Neighbours	Max angular gap
Square	4	90°
Triangular/Hex	6	60°

Hexagonal coordination reduces the maximum angular blind spot by 33%.

Angular harmonic suppression:

Define the angular kernel $K(\theta) = \sum_n \delta(\theta - \theta_n)$. Deviation from isotropy is quantified by Fourier components:

$$\epsilon_m = |Z^{-1} \sum_n \exp(im\theta_n)|$$

For six-fold coordination: $\epsilon_m = 0$ for all m not divisible by 6.

This annihilates all low-order angular harmonics ($m = 1, 2, 3, 4, 5$), which would otherwise dominate anisotropic information leakage.

4.5 Selection Channel 4: BCB Admissibility

Any discrete interface samples directions non-uniformly. BCB requires that the resolution cost for a fact cannot depend on direction. This is equivalent to requiring minimal variance in "distinguishability throughput per tick" across directions.

Among regular planar tilings with equal-length edges, hexagonal coordination **minimises angular variance** in neighbour distribution. Square coordination has larger gaps and correspondingly larger variance.

4.6 Synthesis

Lemma (Uniqueness within the regular tiling class): Among translation-invariant regular planar tilings with equal-length local couplings (no preferred locations; finite-range locality), the triangular/hexagonal (C_6) lattice is uniquely selected as the maximally isotropy-accelerating fixed point. For any finite-range, lattice-symmetric operator on such tilings, the first allowed anisotropic harmonic is fixed by the point group: $C_4 \rightarrow \cos(4\theta)$, $C_6 \rightarrow \cos(6\theta)$. Extension to aperiodic or disordered tilings is an open problem.

Four convergent selection pressures point to the same geometry:

Channel	Mechanism	Selected geometry
Dispersion symmetry	Anisotropy order in Fourier expansion	Hexagonal
RG flow	Anisotropy decay rate under coarse-graining	Hexagonal
Information theory	Angular blind spot minimisation	Hexagonal
BCB/Admissibility	Directional variance minimisation	Hexagonal

Hexagonal coordination is not assumed—it is uniquely selected as the sole admissible regular discrete geometry.

Irregular and aperiodic tilings: The selection result above is stated for regular, translation-invariant, equal-edge local couplings—the natural admissibility class if one requires "no preferred locations" in addition to no preferred directions. Quasicrystalline or aperiodic tilings can exhibit high local rotational order (e.g., 5-fold or 10-fold) but lack global translation invariance.

Why BCB requires translation invariance: The path-independence axiom requires that the total distinguishability cost to resolve a fact between endpoints A and B cannot depend on which causal path connects them. In a non-translation-invariant substrate, the *same displacement* traversed from different starting points would encounter different local structures (different edge lengths, coordination numbers, or coupling strengths). This means the "cost per unit displacement" varies with location—not just with direction. Consider two paths from A to B and from A' to B' where $A' - B' = A - B$ (same displacement vector): if the local substrate structure differs at A vs A', the path costs will differ even for identical displacements. This is a form of path-dependence in the extended sense: the cost to establish a fact about "displacement Δx " depends on where you start. For strict path-independence, the substrate must look the same everywhere, which is precisely translation invariance. (This argument distinguishes spatial *inhomogeneity* from directional *anisotropy*—both violate admissibility, but through different mechanisms.)

Exploring whether any nonperiodic tiling can satisfy the full admissibility constraints is an open problem. For quasicrystals with perfect local matching rules, the location-dependence might be bounded or average to zero over large scales; whether this suffices for admissibility remains unexplored. If such a structure exists and predicts (say) 10Ω , it would be a competing hypothesis

distinguishable from hexagonal by harmonic order. Among regular tilings, six-fold coordination is the unique isotropy-accelerating fixed point.

5. Emergent Lorentz Symmetry and Maxwell Structure

For the general reader: Here's where the payoff comes. Starting from our 2D hexagonal interface with its simple update rules, we can *derive* the fundamental equations of electromagnetism and special relativity—equations that Einstein and Maxwell had to postulate.

The speed of light being constant for all observers? That emerges automatically from the fixed update rate of the interface. The equations governing electric and magnetic fields? They emerge from how phase information flows along the links of our hexagonal grid. We're not assuming these laws—we're deriving them from something deeper.

5.1 Emergent Lorentz Invariance

The interface update dynamics define:

- A **universal speed limit**: the rate at which causal influence propagates (one lattice spacing per tick)
- **No preferred frame**: the discrete dynamics are invariant under lattice symmetries, which flow to continuous rotation/boost invariance under RG

In the continuum limit, the wave equation emerges:

$$\partial^2\phi/\partial t^2 = c^2 \nabla^2\phi$$

where c is set by the ratio of interface spacing to tick duration.

Lorentz symmetry is not postulated but **emerges** as the continuum symmetry of isotropic, finite-speed interface dynamics. The residual discrete symmetry (six-fold) is an experimentally accessible signature of the underlying structure.

5.2 Gauge Structure from Link Variables

Scalar node fields are insufficient to reproduce electromagnetism. Following lattice gauge theory [9, 10, 15], we associate phase variables to oriented links.

Important note: The U(1) gauge structure is an *additional input* to the framework, not derived from BCB or hexagonal selection. BCB selects the interface geometry; it does not determine what fields propagate on that geometry. We assume U(1) link variables because we aim to recover electromagnetism, but this is a choice justified by phenomenological success rather than first-principles derivation. The framework therefore explains why gauge redundancy is natural in

a link-variable description and why hexagonal interface dynamics supports Maxwell form in the continuum, while leaving the origin of the specific gauge group (U(1) vs others) as an open selection problem.

Future directions for gauge selection: Possible selection principles include topological constraints on interface excitations (e.g., requiring well-defined parallel transport), BCB-type arguments applied to link variables (finite-capacity constraints on phase accumulation), or consistency requirements from matter coupling. These remain unexplored and constitute an important direction for future work.

For a link connecting nodes $i \rightarrow j$:

$$U_{ij} = \exp(ia A_{ij})$$

where A_{ij} is a discrete gauge potential and a is the interface micro-length.

Under a local phase transformation $\{\alpha_i\}$ at each node:

$$U_{ij} \rightarrow \exp(i\alpha_i) U_{ij} \exp(-i\alpha_j)$$

This is precisely **U(1) lattice gauge symmetry**—the discrete precursor of electromagnetic gauge invariance.

5.3 Field Strength from Plaquettes

Gauge-invariant observables arise from closed loops. For a minimal hexagonal plaquette p :

$$U_p = \prod_{\{ij \in p\}} U_{ij} = \exp(ia^2 F_p)$$

where F_p is the discrete field strength (curvature).

The Wilson action:

$$S = \sum_p (1 - \text{Re}(U_p))$$

expands for small curvature to:

$$S \rightarrow \int d^2x dt (1/4) F_{\mu\nu} F^{\mu\nu}$$

which is the **Maxwell action**.

5.4 Extending Gauge Structure to the Emergent Bulk

A subtlety arises: lattice gauge theory typically assumes a pre-existing 3+1D lattice. Our interface is 2D. How does the third dimension enter?

The resolution involves the bulk reconstruction (Section 6). The key insight:

1. The 2D interface carries U(1) gauge fields
2. The bulk is reconstructed from interface correlations across coarse-graining scales
3. Gauge fields on the interface induce gauge structure in the reconstructed bulk
4. The emergent 3D gauge theory inherits its structure from the interface

Formally, this is analogous to how holographic dualities relate boundary gauge theories to bulk dynamics—but here the boundary (interface) is fundamental and the bulk is derived.

5.5 Polarization Degrees of Freedom

In the interface picture:

- **Interface modes:** Transverse excitations of the gauge field on the 2D interface
- **Bulk reconstruction:** These modes, viewed across coarse-graining depths, appear as propagating degrees of freedom in 3D
- **Polarization count:** The interface gauge field has components along two independent tangent directions, yielding exactly two transverse polarizations

The counting matches standard electrodynamics: two physical photon polarizations, with gauge redundancy removing longitudinal and timelike modes.

6. Bulk Reconstruction: Emergent Three-Dimensional Space

For the general reader: If reality is fundamentally a 2D surface, why do we experience three dimensions? This is like asking how a hologram creates a 3D image from a 2D film.

The answer involves *scale*. Imagine looking at the interface through microscopes of different powers. At high magnification, you see fine details. At low magnification, you see only large-scale patterns. Each magnification level gives a slightly different picture. Stack all these pictures together, and the "direction" from fine to coarse becomes a third dimension.

This isn't just a metaphor—it's how distances in the third dimension are actually defined. Two points are "close in z" if their correlations persist across many scales. The 3D space we inhabit is reconstructed from how patterns on the 2D interface relate to each other across different resolutions.

Scope note: The non-admissibility theorem and hexagonal selection do not depend on a fully rigorous bulk reconstruction. This section provides a **speculative outline**—correlation geometry across renormalisation depth—suggesting *one possible mechanism* by which an emergent 3D description might arise from interface data.

Important limitations: We do *not* derive that correlation-based reconstruction yields Euclidean 3D geometry specifically. Correlation decay could, in principle, generate many different geometries depending on the interface dynamics. We also do not show that observers made of interface excitations would actually perform this reconstruction, only that it is mathematically conceivable. A complete treatment would require (i) deriving the correlation structure from interface microdynamics, (ii) showing this yields approximately flat 3D geometry, and (iii) demonstrating that observer-experiences reconstruct this geometry. These remain open problems.

The core results of this paper—the non-admissibility theorem and hexagonal selection—stand independently of whether this particular reconstruction mechanism is correct.

Alternative reconstruction routes: Tensor-network/entanglement-based emergence of geometry, or other holography-inspired maps (e.g., AdS/CFT-style bulk reconstruction), could replace the specific coarse-graining construction presented here without affecting the core results of Sections 2–4 (non-admissibility and hexagonal selection) or Section 8 (the symmetry-protected 6Ω discriminator). Section 6 presents one candidate mechanism, not a keystone of the framework.

6.1 The Reconstruction Problem

If propagation is fundamentally 2D, how do observers experience three spatial dimensions? This section outlines one possible reconstruction mechanism, acknowledging its speculative status.

6.2 Coarse-Graining Depth as the Third Dimension

Introduce a continuous parameter z representing **coarse-graining depth**. Each value of z corresponds to a renormalised description of the interface at length scale $\ell(z) = \ell_0 e^z$.

We model the emergent bulk as the **family of coarse-grained interface descriptions indexed by z** :

Bulk point $(x, y, z) \leftrightarrow$ Interface observable at position (x, y) smoothed to scale $\ell(z)$

Under translation invariance, isotropy, and scale-local RG flow, the natural geometry of the space of scales is one additional dimension. This provides the structural basis for emergent 3D geometry.

6.3 Operational Definition of Bulk Distance

Let $O(x; z)$ be a local interface observable at position x and scale z . Define the correlation function:

$$C(x, y; z) = \langle O(x; z) O(y; z) \rangle$$

The bulk distance is operationally defined:

$$d_z(x, y) = -\ell(z) \log[C(x, y; z) / C(x, x; z)]$$

This identifies **bulk proximity with correlation strength**. Points that remain correlated under coarse-graining are "close" in the emergent bulk.

6.4 Emergent Metric

For infinitesimal separations:

$$d_z^2(x, x + dx) \approx g_{ij}(z, x) dx^i dx^j$$

If interface correlations are isotropic in the tangent plane (guaranteed by hexagonal coordination at leading order), then:

$$g_{ij} \propto \delta_{ij}$$

at leading order—a **flat, isotropic metric** emerges.

6.5 What Physical Process Corresponds to "Moving in z"?

In this framework, z labels coarse-graining depth: larger z corresponds to descriptions in which fine interface degrees of freedom have been averaged into effective variables. A key concern is whether such coarse-graining is merely a mathematical convenience or has a physical counterpart.

Physical counterpart: finite-resolution coupling

Any real interaction—whether scattering, emission/absorption, or measurement—couples to interface degrees of freedom with finite bandwidth. Concretely, the coupling strength to interface modes is weighted by a response function $W_\Lambda(k)$ that suppresses modes above some effective cutoff Λ (set by energy, temperature, finite detector size, and finite interaction time). A process with smaller Λ "sees" only coarse features; a process with larger Λ resolves finer features.

The relationship between cutoff and z -coordinate is:

- High Λ (high energy, fine resolution) \leftrightarrow small z
- Low Λ (low energy, coarse resolution) \leftrightarrow large z

This is not a metaphor but standard physics: effective field theory is built on exactly this scale-dependent coarse-graining.

Why observers integrate over a range of z -scales

Observers are composite systems built from many interacting processes spanning a hierarchy of energy and length scales (electronic, vibrational, thermal, mechanical). Their effective dynamics therefore involve simultaneous couplings to a range of cutoffs Λ —i.e., a range of coarse-graining depths z .

In this sense, "integrating over z " is not an extra postulate: it is the physical fact that no observer couples to only one scale. Stable macroscopic records are maintained precisely by discarding microscopic detail into coarse degrees of freedom (entropy production). An observer "perceives" 3D space because its constituent processes collectively sample interface correlations across multiple z -values.

Analogy to holographic reconstruction

This is analogous to how holographic models (AdS/CFT, tensor networks) reconstruct bulk geometry from boundary CFT data, but with the interface playing the role of the fundamental boundary. The key difference: in standard holography, the bulk is "dual" to the boundary; here, the bulk is *emergent from* coarse-graining the interface.

What a successful reconstruction must show

This section provides a concrete mechanism by which a family of coarse-grained interface descriptions can be physically relevant. A complete emergence account would require:

- (i) **Explicit interface microdynamics**: A specific Hamiltonian or update rule for interface degrees of freedom
- (ii) **Derived correlation functions**: Analytical or numerical calculation of $\langle O(x; z) O(y; z) \rangle$ from the microdynamics
- (iii) **Euclidean geometry proof**: Demonstration that the induced correlation geometry is approximately Euclidean and 3D over the relevant regime
- (iv) **Observer consistency**: A mechanism by which matter excitations implement finite-bandwidth coupling that samples these correlations consistently across observers

These steps remain future work. The present discussion establishes the conceptual framework and physical interpretation; the core theorems (non-admissibility, hexagonal selection) do not depend on the bulk reconstruction being completed. Section 6 can be read as a research direction indicating how the framework would connect to emergent spacetime.

Symmetry preservation requirement: A complete bulk reconstruction must verify that the correlation-depth map preserves the interface's six-fold symmetry at leading order. If the reconstruction itself introduces additional anisotropy (e.g., from boundary conditions or coarse-graining kernel), this could contribute to the observed signal with effective suppression order lower than the interface baseline. Section 8.9 discusses how to interpret such scenarios empirically.

Connection to tensor networks: The coarse-graining reconstruction has structural parallels with tensor network models, particularly MERA (Multi-scale Entanglement Renormalization Ansatz), where successive RG layers define an emergent "bulk" dimension. The key difference: standard tensor networks encode entanglement structure of a pre-existing quantum state, whereas our framework proposes the interface as a causal substrate from which even the quantum state description emerges. Nonetheless, tensor network methods may provide technical tools for making the bulk reconstruction rigorous.

6.6 Why Observers Agree on Three Dimensions

All observers constructed from interface dynamics experience the same reconstruction:

1. Matter is composed of interface excitations (developed in future work)
2. Interactions between matter components probe interface correlations
3. The correlation structure is objective (observer-independent)
4. Therefore, all observers reconstruct the same bulk geometry

The "third dimension" is not a convention but an objective feature of the correlation structure.

6.7 Correlation Structure from Distinguishability Geometry

A common concern is that bulk reconstruction appears to require an explicit microscopic action or Hamiltonian in order to compute correlations $\langle O(x)O(y) \rangle$. We emphasise that this is not the case. Within the BCB framework, correlations are not arbitrary dynamical objects but geometric consequences of finite distinguishability capacity, gauge redundancy, and causal update structure.

BCB provides a natural notion of correlation independent of amplitude dynamics. Two interface locations x and y are correlated to the extent that facts committed at those locations share common causal ancestry within the finite-capacity update network. Operationally, correlation measures the degree to which distinguishability committed at x constrains future evolution at y , and vice versa. This definition is intrinsic to the fact-commitment ontology and does not presuppose unitary evolution or Hilbert space structure.

At the geometric level, BCB fixes the admissible overlap structure. Finite information capacity together with gauge redundancy forces the distinguishability manifold to carry a Fisher-type metric with a compact $U(1)$ phase fibre. The only overlap functional compatible with (i) monotonicity under coarse-graining, (ii) composition stability, and (iii) gauge invariance is the Fisher–Bhattacharyya overlap with phase modulation, which coincides with the quantum inner product structure in the continuum limit. (This uniqueness result is developed in companion work on BCB geometry; a brief derivation sketch appears in Appendix F.) This result implies that correlation functions are not free inputs: their functional form is fixed up to scale by BCB geometry alone.

Importantly, this does not require specifying a microscopic action. Any admissible correlation function must:

1. Respect hexagonal symmetry on the interface
2. Decay with separation in a manner consistent with finite Fisher length
3. Preserve path-independent distinguishability cost

Different microphysical realisations may correspond to different members of this constrained class (e.g., exponential versus power-law decay), but all yield the same large-scale correlation geometry and hence the same reconstructed bulk dimensionality.

What this section establishes: The present paper does not commit to a specific dynamical mechanism for correlation propagation. The non-admissibility theorem and hexagonal selection results depend only on the existence of such a correlation structure, not on its detailed microphysical origin. We are deriving what correlations *must* look like before deriving how they arise dynamically—the correct order for a foundational framework. Providing an explicit action or update rule is an important direction for future work, but is not required for the core conclusions of this paper.

7. Phenomenology: Structural Explanations and Research Directions

For the general reader: Standard physics tells us *what* the laws are, but often not *why* they're that way. Why is the speed of light the same for everyone? Why do photons have exactly two polarizations? Physics textbooks present these as facts to be accepted.

The interface framework is different: it *explains* these facts. Each property that standard physics takes as given emerges naturally from the interface picture. This section walks through nine such explanations—not to replace the standard equations (which remain correct and useful), but to show they're not arbitrary. They're consequences of something deeper.

Note on scope: Items 1–7 are **structural explanations** derived within the framework. Items 8–9 are **research directions** rather than completed explanations; they indicate how the framework would connect light to emergent geometry under a successful reconstruction mechanism (Section 6), and are included to show explanatory direction rather than to claim closure.

7.1 Why Is the Speed of Light a Universal Constant?

Standard status: Postulated as the second axiom of special relativity.

Interface explanation: The speed of light emerges as the fixed update rate of the causal interface—one lattice spacing per tick. Because all physical processes are interface-mediated, all observers measure time and distance using interface dynamics. The maximum propagation speed is therefore invariant by construction.

Interpretive remark (entropy-limited causality): Within the interface framework, the invariant speed of light acquires a direct causal interpretation: it is the maximum speed at which new, irreversible information (entropy) can be propagated and committed as a stable record. Energy or amplitude may exhibit superluminal features under certain conditions (phase velocity, group reshaping, entanglement correlations), but entropy production and fact-commitment cannot. Because causal order and record-formation are observer-independent, the maximum speed of entropy propagation must be invariant across frames. Lorentz symmetry and the invariant speed c therefore arise as consistency requirements of finite-rate entropy communication rather than as

independent postulates. This aligns with information-theoretic reconstructions of spacetime and with Jacobson's thermodynamic derivation of gravitational dynamics [11].

7.2 Why Does Lorentz Symmetry Exist?

Standard status: Postulated or derived from the invariance of c .

Interface explanation: Lorentz symmetry is the continuum limit of the discrete symmetries of isotropic interface dynamics. It is emergent rather than fundamental—a consequence of how finite-speed, isotropic propagation appears at scales large compared to the interface micro-length.

7.3 Why Does Light Define Causality?

Standard status: Light cones are geometric features of the Minkowski metric.

Interface explanation: Causality is primary; geometry is derived. Light cones correspond to the reachable region of the interface after successive update cycles. Spacetime geometry *encodes* this causal structure rather than generating it.

7.4 Why Does Gauge Redundancy Exist? (*Derived; conditional on $U(1)$ input*)

Standard status: Gauge symmetry is a mathematical redundancy required for consistent formulation of electrodynamics.

Interface explanation: Given link-based $U(1)$ degrees of freedom (an input to the framework, not derived—see Section 5.2), gauge redundancy reflects local freedom in phase reference on interface links. It is not an abstract symmetry imposed for mathematical convenience but an inherent feature of link-based dynamics: the physics depends on phase *differences* around loops, not absolute phases. The framework explains why gauge redundancy is *natural* given link variables, but does not derive why $U(1)$ specifically.

7.5 Why Are There Exactly Two Photon Polarizations? (*Derived; conditional on $U(1)$ input*)

Standard status: Follows from representation theory of the Poincaré group and gauge fixing.

Interface explanation: The 2D interface supports two independent transverse directions for gauge field excitations. Gauge redundancy removes unphysical modes. The embedding into the emergent 3D bulk preserves this count, yielding exactly two physical polarizations.

7.6 Why Is Vacuum Not Birefringent? (*Derived*)

Standard status: An empirical fact, constrained to extreme precision by observation.

Interface explanation: Birefringence requires different propagation speeds for different polarizations. Hexagonal symmetry treats both transverse directions equivalently at leading order. Any birefringence would require breaking this symmetry, which can only occur at very high order (sixth order in lattice harmonics).

7.7 Why Does Light Appear Continuous Despite Possible Discreteness? (*Derived*)

Standard status: Continuity is assumed; discrete models are constrained to be "small enough" to avoid detection.

Interface explanation: Hexagonal coordination pushes discreteness effects to sixth order in the dispersion relation. Under RG flow, these effects decay as b^{-4} , rendering the interface effectively continuous at observable scales. Discreteness is not merely "too small to see" but is systematically suppressed by the geometry.

7.8 Why Can Space Be Emergent Rather Than Fundamental? (*Reframed; depends on Section 6*)

Standard status: Spacetime is typically assumed as the fundamental arena.

Interface explanation: Space emerges operationally from correlations across interface states at different coarse-graining scales. Geometry is constructed from dynamics rather than assumed as a background. The interface provides the causal substrate; space is derived.

7.9 Why Are Light and Geometry So Deeply Linked? (*Reframed; depends on Section 6*)

Standard status: The link between light and geometry is encoded in metric structure but remains conceptually unexplained.

Interface explanation: Light propagation *defines* the causal structure from which geometry is constructed. The intimate connection between light and spacetime is therefore structural and necessary, not coincidental or merely conventional.

8. Falsifiable Predictions

For the general reader: A theory that can't be tested isn't science—it's philosophy. This section describes how to test our framework experimentally.

The key prediction: if light really propagates on a hexagonal grid, there should be a *tiny* six-fold pattern in how the speed of light varies with direction. Imagine the grid has a slight "grain," like wood. Light traveling along the grain versus across it would have imperceptibly different speeds.

Modern experiments can detect variations in light speed smaller than one part in 10^{17} —that's like measuring the width of a human hair from the distance of the nearest star. If the hexagonal pattern exists at all, these experiments should eventually find it. If they find a *four*-fold pattern instead (or no pattern at all), our framework is wrong. That's what makes this real science.

8.1 The Core Prediction: Six-Fold Anisotropy

If residual anisotropy leaks through from the interface micro-geometry, it must respect hexagonal symmetry. The leading allowed angular dependence is:

$$\delta c/c(\theta) = \epsilon (ka)^p \cos(6\theta) + \text{higher harmonics}$$

where:

- θ = propagation direction relative to a fixed axis
- ϵ = dimensionless amplitude (microphysics-dependent)
- $k = 2\pi/\lambda$ = optical wavenumber
- a = interface micro-length
- p = suppression order ($p \geq 4$ from hex symmetry + RG arguments)

8.2 Distinguishing Prediction: 6Ω vs 4Ω

Critical test: The *lowest* allowed angular harmonic discriminates between underlying geometries:

Geometry	Predicted lowest harmonic	Signature
Hexagonal	$\cos(6\theta)$	6Ω in rotating cavity
Square	$\cos(4\theta)$	4Ω in rotating cavity
Isotropic	None	No angular dependence

Detection of a **4Ω signal without 6Ω** would falsify the hexagonal selection arguments. Detection of **6Ω without 4Ω** would support the framework. Detection of neither constrains the parameter space.

8.3 Parameter Estimation

The framework does not uniquely determine ϵ from first principles—this depends on detailed interface microphysics. However, we can bound the "natural" range:

Upper bound on ϵ : Symmetry arguments suggest ϵ should not exceed $O(1)$, as larger values would indicate fine-tuning against the symmetry suppression.

Heuristic ϵ -floor (auxiliary assumption): Any quantitative lower bound on ϵ requires microphysical modelling of the interface (its stiffness, defect spectrum, and fluctuation environment). We therefore treat $\epsilon \gtrsim 10^{-6}$ **not as a derived bound but as a physically**

motivated heuristic corresponding to a "generic substrate" regime in which symmetry-breaking from imperfections or fluctuations is not exponentially suppressed.

Physical motivation for the heuristic: If the interface has any finite-temperature dynamics, thermal fluctuations induce lattice distortions with amplitude $\delta a/a \sim \sqrt{(k_B T / \kappa a^2)}$, where κ is the effective stiffness. Even at $T \rightarrow 0$, quantum zero-point fluctuations contribute $\delta a/a \sim \sqrt{(\hbar\omega / \kappa a^2)}$. These fluctuations break perfect six-fold symmetry locally and should generate $\varepsilon \sim (\delta a/a)^2 \sim 10^{-2}$ to 10^{-6} depending on stiffness. Topological defects (dislocations, disclinations) would contribute additional symmetry-breaking. The upshot: $\varepsilon = 0$ requires either (i) zero temperature with infinite stiffness (unphysical), or (ii) a topological protection mechanism that exactly preserves six-fold symmetry against all perturbations. The latter is possible in principle but constitutes a strong theoretical commitment.

Estimated range (heuristic): $10^{-6} \lesssim \varepsilon \lesssim 1$

This range is not derived; it reflects order-of-magnitude expectations for generic substrates without exotic protection mechanisms.

Physical interpretation of ε : In the absence of exact micro-dynamics, ε can be interpreted as an effective symmetry-breaking parameter induced by interface imperfections (defects, strain, topological disorder) and/or fluctuation-driven departures from perfect six-fold order. In a "nearly perfect" interface, ε should be controlled by the density of such defects or fluctuations rather than being arbitrarily tuneable.

Qualitative vs quantitative prediction: The framework makes two distinct types of predictions:

1. **Qualitative (robust):** The *lowest-order* angular harmonic is $\cos(6\theta)$, not $\cos(4\theta)$. This is independent of ε . A clean detection of 4Ω without 6Ω would falsify the hexagonal selection regardless of amplitude.
2. **Quantitative (parameter-dependent):** The *amplitude* of the 6Ω signal depends on ε and p , which are not determined from first principles. Null results constrain the parameter space but cannot definitively falsify the framework unless sensitivity reaches levels where even minimal fluctuation-induced ε would be detectable.

The qualitative prediction is the sharper test. The quantitative prediction becomes sharper as experimental sensitivity improves and as theoretical work constrains ε from microphysical considerations.

Possibility of strong suppression: In principle, a substrate that implements powerful error-correction or symmetry-protection at the interface level could suppress ε exponentially, potentially pushing the 6Ω signal beyond foreseeable detectability. This does not negate the qualitative prediction (the lowest allowed harmonic is six-fold), but it weakens near-term falsifiability by amplitude. Conversely, a clean detection of 6Ω would be doubly informative: it would support the interface picture *and* simultaneously disfavor strong "UV error-correction that perfectly erases lattice residue" at the relevant scales.

ϵ as an assumption-class rather than a constant: The framework does not require a universal lower bound on ϵ ; rather, ϵ encodes how strongly micro-geometry leaks into macroscopic propagation. We distinguish three regimes (see **ϵ taxonomy** in Section 8.8 for falsification implications):

1. *Generic substrate*: finite stiffness, nonzero defect/strain density \rightarrow nonzero floor set by fluctuation and disorder ($\epsilon \sim 10^{-6}$ to 1)
2. *Symmetry-protected substrate*: strong UV error-correction \rightarrow ϵ exponentially suppressed, amplitude tests impractical
3. *Random-geometry substrate* (e.g., causal sets): no fixed micro-geometry \rightarrow $\epsilon \approx 0$ by construction, no 6Ω predicted

We treat " ϵ not exponentially suppressed" (i.e., regime 1) as an auxiliary physical assumption (A_ϵ) required for amplitude-based falsification. This is not a weakness but a clarification: the structural prediction (6Ω vs 4Ω) remains ϵ -independent, while amplitude predictions require specifying the ϵ regime.

For the suppression order p , RG and symmetry arguments suggest $p \geq 4$. The most conservative assumption is $p = 4$.

8.4 Experimental Channel: Rotating Optical Cavities

Modern Michelson-Morley-style experiments compare resonance frequencies of orthogonal optical cavities on a rotating turntable. Cavity resonance: $f \approx mc_{\text{eff}}/(2L)$, so $\delta f/f \approx \delta c/c$.

For rotation at angular speed Ω :

$$\theta(t) = \Omega t + \theta_0$$

The six-fold signature becomes:

$$\delta f/f(t) \approx \epsilon (ka)^p \cos(6\Omega t + \phi)$$

This predicts a **spectral line at 6Ω** in the beat-frequency time series.

8.5 Current Experimental Status

Leading constraints on light propagation anisotropy:

Experiment	Sensitivity ($\Delta c/c$)	Reference
Herrmann et al. (2009)	$\sim 10^{-17}$	Phys. Rev. D 80, 105011
Zhang et al. (2021)	$\sim 10^{-17}$	Phys. Lett. A 416, 127666

Existing analyses focus on 2Ω (two orthogonal cavities repeat every 180°). A targeted search for 6Ω has not been reported.

8.6 Implied Bounds on Interface Scale

Using $\delta c/c < 10^{-17}$ with optical wavelength $\lambda = 1064 \text{ nm}$ ($k \simeq 5.9 \times 10^6 \text{ m}^{-1}$):

Throughout we treat $p = 4$ as the conservative baseline motivated by symmetry + RG suppression; $p = 6$ is shown to illustrate the stronger-suppression regime and the resulting scale sensitivity.

p Bound on a (m) for $\varepsilon = 1$ Energy scale $\hbar c/a$ (eV)

2	5.4×10^{-16}	3.7×10^8
4	9.5×10^{-12}	2.1×10^4
6	2.5×10^{-10}	7.9×10^2

Note on the $p = 6$ case: The energy scale $\sim 800 \text{ eV}$ is surprisingly low—well within reach of X-ray physics. However, this does not immediately conflict with high-energy tests of Lorentz invariance for several reasons:

(i) **Different SME coefficients:** In the Standard-Model Extension (SME) photon sector, astrophysical birefringence tests primarily constrain the $\tilde{\kappa}_{\{o+\}}$ (odd-parity, CPT-odd) coefficients, while time-of-flight tests constrain $\tilde{\kappa}_{\{tr\}}$. Rotating cavity experiments—and our 6Ω prediction—constrain the even-parity $\tilde{\kappa}_{\{e-\}}$ coefficients related to spatial anisotropy. These are independent parameters in the SME; strong bounds on one do not directly constrain the others. See Kostelecký & Russell [3] Tables D12–D15 for current coefficient bounds.

(ii) **Angular integration:** Our prediction involves a specific angular signature ($\cos 6\theta$) that integrates to zero over all directions. Isotropic high-energy tests—which average over source directions or propagation paths—would not detect such anisotropy even if present. Many high-energy constraints are effectively direction-averaged or probe different dispersion/birefringence channels; they do not directly constrain a pure $\cos(6\theta)$ cavity-anisotropy term.

(iia) **Collider physics:** High-energy collider experiments probe reaction rates, cross-sections, and decay rates rather than the specific angular-anisotropy channel addressed here. The $\cos(6\theta)$ signature averages to zero in isotropic collision environments and does not affect the observables typically measured at particle accelerators. The $\sim 10 \text{ keV}$ ($p = 4$) or $\sim 800 \text{ eV}$ ($p = 6$) scales are therefore not directly constrained by collider data.

The most relevant existing constraints come from the rotating cavity experiments cited above, which directly probe angular anisotropy. Future X-ray cavity experiments or direction-resolved astrophysical polarimetry might provide complementary constraints at higher effective energies.

(iii) **X-ray and keV-scale physics:** One might ask: if the interface scale is at $\sim 800 \text{ eV}$ for $p = 6$, shouldn't effects appear in X-ray crystallography, synchrotron experiments, or keV-scale atomic physics? The answer is that these experiments typically probe absorption, scattering, or spectroscopic transitions rather than the angular-anisotropy channel measured by rotating optical cavities. Direct constraints on $\cos(6\theta)$ anisotropy at X-ray frequencies would require purpose-

built rotating cavity experiments at those wavelengths, which do not currently exist. The ~ 800 eV scale is therefore not immediately excluded by existing X-ray physics.

Parameter space implication: However, this cuts both ways. If $p = 6$ and $\varepsilon \approx 1$, the framework predicts effects at ~ 800 eV that would likely have been seen in *precision spectroscopy* of atomic transitions or X-ray interferometry if a suitable experiment had been performed. The absence of anomalous reports from the X-ray community, while not a direct constraint, suggests that either (i) $p < 6$ (making predictions at higher energies), (ii) $\varepsilon \ll 1$ (suppressing the amplitude), or (iii) the interface scale a is smaller than the table indicates. This constrains the (p, ε, a) parameter space more tightly than rotating optical cavities alone.

ε -scaling of bounds: The bounds above assume $\varepsilon = 1$. For general ε , the inferred length bound scales as:

$$a_{\max}(\varepsilon) = a_{\max}(\varepsilon = 1) \cdot \varepsilon^{(-1/p)}$$

For example, $\varepsilon = 10^{-3}$ increases a_{\max} by a factor of $10^{(3/p)}$. With $p = 4$, this is a factor of ~ 5.6 ; with $p = 6$, a factor of ~ 3.2 .

For $p = 4$ (most conservative), the interface scale is bounded above ~ 10 keV—well below the Planck scale but potentially accessible to high-energy probes.

8.6a Auxiliary Constraints from Independent Channels

A framework with three undetermined parameters (ε, p, a) could be accused of excessive flexibility. We note several independent channels that constrain the parameter space from different directions, strengthening falsifiability:

1. High-energy astrophysics (GZK cutoff, gamma-ray propagation)

If the interface scale a is sufficiently large (energy scale $\hbar c/a$ sufficiently small), photon propagation over cosmological distances could accumulate detectable dispersion. The absence of anomalous time delays in gamma-ray bursts (GRBs) and the observed GZK cutoff constrain dispersion effects at high energies. These tests primarily constrain $p = 1$ (linear) or $p = 2$ (quadratic) dispersion, which our framework forbids by symmetry for a hexagonal substrate. For $p \geq 4$, the suppression is strong enough that current astrophysical observations do not constrain the interface scale.

2. Precision atomic physics

For $p = 4$ and $\varepsilon \sim 1$, the interface energy scale $\hbar c/a \sim 20$ keV falls within reach of precision X-ray spectroscopy. Anomalous angular dependence in atomic transition frequencies or fine-structure splitting could constrain $\varepsilon(ka)^4$ independently of cavity experiments. To our knowledge, no targeted search for 6Ω angular dependence has been performed in atomic physics.

3. Photon propagation in strong magnetic fields (QED tests)

High-precision QED tests (anomalous magnetic moment, Delbrück scattering, vacuum birefringence) probe the photon sector at energies up to $\sim \text{MeV}$. While these tests are generally isotropic (averaging over propagation directions), direction-resolved measurements in magnetar environments or laboratory strong-field setups could provide independent constraints.

4. Multi-wavelength rotating cavity experiments

As discussed in Section 8.9, measuring R_6 at multiple optical wavelengths extracts p independently. If p is determined, then ε and a are constrained separately rather than through their combination $\varepsilon(ka)^p$. This converts a three-parameter family into separate constraints.

Summary: While the 6Ω cavity measurement is the primary prediction, auxiliary channels provide independent constraints. Detection of 6Ω combined with null results in auxiliary channels would overdetermine (ε, p, a) and test internal consistency. Conversely, tension between channels would indicate physics beyond the minimal interface model.

8.7 Proposed Experimental Protocol

1. **Data acquisition:** Collect beat-frequency time series $y(t) = \delta f/f(t)$ from rotating cavity experiment
2. **Harmonic decomposition:** Fit:

$$y(t) = \sum_n [A_n \cos(n\Omega t) + B_n \sin(n\Omega t)] + \text{sidereal sidebands} + \text{noise}$$

3. **Extract 6Ω amplitude:** $R_6 = \sqrt{A_6^2 + B_6^2}$
4. **Systematic control:** Repeat at multiple Ω ; verify any signal doesn't scale with mechanical artefacts
5. **Interpretation:**
 - o $R_6 > \text{threshold}$ with correct Ω -independence \rightarrow Detection
 - o R_6 consistent with noise \rightarrow Bound: $|\varepsilon|(ka)^p < R_6_max$
 - o $R_4 > 0$ but $R_6 = 0 \rightarrow$ Falsifies hexagonal selection

8.8 Falsification Logic and the Role of ε

The 6Ω prediction has two logically distinct layers, and clarity about their status prevents confusion about what constitutes falsification.

(I) Structural falsifiers (ε -independent)

These are hard falsifiers that do not depend on the amplitude parameter ε :

1. **Detection of 4Ω without 6Ω :** Falsifies hexagonal selection regardless of amplitude. A square lattice predicts 4Ω as the lowest harmonic; hexagonal predicts 6Ω . This is a geometry discriminator, not an amplitude test.
2. **Wrong wavelength scaling:** If a 6Ω signal appears but does not scale as $(ka)^p$ with consistent p across wavelengths, the RG suppression mechanism is falsified.

3. **Wrong symmetry class:** Detection of dominant 3Ω , 5Ω , or other harmonics inconsistent with hexagonal symmetry would falsify the framework.

These structural tests are clean: they depend only on which harmonic appears, not on its strength.

(II) Amplitude falsifiers (require auxiliary assumption)

A null result—no 6Ω detected at some sensitivity $R_{6,max}$ —bounds the combination $|\epsilon|(ka)^p$ but is not strictly falsifying unless one adopts an auxiliary physical assumption:

Assumption A_ϵ (No perfect UV symmetry-protection): ϵ cannot be arbitrarily small; there exists a physically motivated lower bound $\epsilon \geq \epsilon_{min}$.

Under this assumption, a null becomes falsifying when:

$$R_{6,max} < \epsilon_{min} \cdot (ka)^p$$

Without A_ϵ , one can always "save" the framework by postulating smaller ϵ . This is not a defect unique to this framework—it is generic to any theory with undetermined coupling constants.

ϵ taxonomy: We distinguish three regimes:

Regime	Expected ϵ	Testability
Generic interface (no UV error-correction): imperfections, strain, defects, zero-point roughness	$\epsilon \sim 10^{-6}$ to 1	Amplitude tests feasible
Symmetry-protected / error-corrected interface: strong UV protection mechanisms	ϵ exponentially small	Amplitude tests impractical; but detection would rule out such protection
No fixed micro-geometry (e.g., causal sets, random sprinkling)	$\epsilon \approx 0$ by construction	Model predicts no 6Ω ; structural test only

For the generic-interface regime, the statistical mechanics arguments in Section 8.3 suggest $\epsilon \gtrsim 10^{-6}$ is physically reasonable. Exact $\epsilon = 0$ would require either infinite stiffness or topological protection—strong theoretical commitments.

Conditional falsification criterion: If one assumes (i) $a \gtrsim \ell_{Planck}$ (or some BCB-motivated micro-length), (ii) $p \leq 4$ (conservative baseline), and (iii) $\epsilon \geq 10^{-6}$ (generic interface), then a null at sensitivity $R_{6,max}$ falsifies the model when $R_{6,max} < 10^{-6} \cdot (ka)^4$. Current optical cavity experiments are not yet at this threshold, but future improvements could reach it.

Summary: Null results should be interpreted as parameter constraints unless accompanied by the auxiliary assumption A_ϵ . The structural prediction (6Ω vs 4Ω) remains the sharper discriminator because it is ϵ -independent.

Falsification summary table:

Prediction type	Confirms	Falsifies	Depends on ϵ ?
Structural	Detect 6Ω with correct symmetry class	Detect 4Ω without 6Ω ; wrong harmonic	No
Amplitude	Detect 6Ω at predicted $(ka)^p$ scaling	Null below $\epsilon_{\min}(ka)^p$ under A_ϵ	Yes

This ends the "unfalsifiable" concern: the structural prediction is clean and ϵ -independent; amplitude predictions require auxiliary assumptions that are explicitly stated.

One-sentence summary for experimentalists: Detection of 4Ω without 6Ω falsifies the framework unconditionally; detection of 6Ω confirms it; null results constrain parameters under assumption A_ϵ .

8.9 Determining the Suppression Order p

In this framework, p is not a free parameter but the order of the first nonvanishing symmetry-allowed anisotropic term in the long-wavelength expansion of the effective dispersion relation.

Theoretical expectations: For a generic local, translation-invariant propagation operator on a triangular/hex lattice, the first anisotropic harmonic is $\cos(6\theta)$ and appears at order $|k|^6$; lower-order $|k|^4$ anisotropy is symmetry-forbidden by the six-fold rotation group. This gives $p = 6$ as the **symmetry baseline** for an ideal hexagonal interface.

However, if additional microphysical symmetries exist, or if there are fine-tuned cancellations in the effective operator coefficients, the leading anisotropy could in principle be pushed to higher order ($p = 8, 10, \dots$). This constitutes an additional structural assumption beyond hexagonal symmetry alone.

When can $p = 4$ appear despite a hexagonal substrate? In an ideal, infinite, translation-invariant hexagonal interface with symmetric local couplings, the leading anisotropy is six-fold and enters at order $|k|^6$. Apparent lower-order harmonics can nevertheless arise if additional symmetry-breaking occurs outside the ideal interface dynamics:

- (i) **Apparatus coupling symmetry:** Orthogonal cavity geometries introduce strong $2\Omega/4\Omega$ systematic channels that must be distinguished from substrate effects. Even if the substrate is C_6 , the coupling to the cavity can pick out orthogonal axes (two cavities at 90°), producing effective 4Ω response in the measured signal.
- (ii) **Domain/grain structure:** Grain boundaries can create domains with different orientations; averaging over domains with preferred alignment relative to the apparatus can produce lower-order harmonics in the observed signal (especially if domain distribution is not isotropic in the lab frame).

(iii) **Interface-to-bulk reconstruction anisotropy:** If the 3D reconstruction map is anisotropic at leading order (e.g., due to boundary conditions, finite-size effects, or coarse-graining kernel), it can leak lower-order harmonics not present in the interface dispersion.

We therefore treat $p = 6$ as the **symmetry baseline** and $p = 4$ as a **conservative bounding assumption** that allows for such additional symmetry-breaking. A measured value of p is itself an empirical discriminator between these scenarios.

Primary prediction: $p = 6$: To be explicit about the framework's commitment: the primary theoretical prediction is $p = 6$, corresponding to an ideal hexagonal interface with no additional symmetry-breaking. Detection of 6Ω at the predicted $(ka)^6$ scaling would be strong confirmation. Detection of 6Ω with $(ka)^4$ scaling would indicate additional symmetry-breaking mechanisms (domain averaging, apparatus coupling, reconstruction effects) requiring follow-up investigation—it would support the hexagonal substrate while revealing additional physics not captured by the minimal model. The key discriminator remains 6Ω vs 4Ω : any $p \geq 4$ with six-fold angular dependence supports the hexagonal substrate over square or isotropic alternatives.

Empirical determination of p : The suppression order p can be extracted experimentally by repeating the 6Ω harmonic measurement at multiple optical wavelengths λ . Since:

$$\delta c/c \propto \varepsilon(ka)^p = \varepsilon(2\pi a/\lambda)^p$$

the 6Ω amplitude scales as $R_6 \propto \lambda^{-p}$ at fixed a and ε . A multi-wavelength rotating cavity programme would therefore:

1. Measure R_6 at wavelengths $\lambda_1, \lambda_2, \lambda_3, \dots$ (e.g., different laser lines)
2. Fit $\log(R_6)$ vs $\log(\lambda)$ to extract the slope $-p$
3. Convert the present bound on $\varepsilon(ka)^p$ into separate constraints on ε and a once p is known

This converts p from a theoretical uncertainty into a measurable property, and transforms amplitude bounds into direct constraints on interface microphysics.

Important caveat: This protocol requires detecting a 6Ω signal first. If no 6Ω signal is detected at any wavelength, the protocol cannot extract p separately; it can only bound the combination $\varepsilon(ka)^p$. The multi-wavelength approach becomes informative only after a positive detection.

9. Anticipated Objections and Responses

For the general reader: Any bold theoretical claim invites skepticism—and it should. This section addresses the most natural objections head-on. If you're thinking "but wait, what about...?", we've probably anticipated your concern.

Objection 1: "Maxwell's equations already describe local propagation in 3D. Why is an interface needed?"

Response: Maxwell's equations describe emergent continuum behaviour with extraordinary accuracy. The interface framework addresses a different question: what is the *ontological substrate* from which this behaviour emerges? Under BCB/admissibility, primitive 3D propagation encounters a consistency problem (the dilution trilemma). The interface provides a consistent substrate; Maxwell theory emerges as its effective description.

Objection 2: "Why can't distinguishability dilute like energy?"

Response: See Section 2.3. The key difference is the resolution threshold. Energy can be arbitrarily diluted while remaining physically meaningful. Distinguishability below threshold resolves no facts—it is not "weak information" but *no information*. This discrete/continuous distinction is not ad hoc but reflects the nature of binary facts.

Objection 3: "Isn't this just holography in different language?"

Response: Related but distinct. Standard holography asserts bulk/boundary *duality*—the bulk exists and is encoded on the boundary. Our framework makes a stronger claim: the interface is *fundamental*, and the bulk is *reconstructed*. There is no separate bulk reality being encoded; rather, bulk geometry emerges from interface correlation structure.

Objection 4: "Why 2D? Why not 1D or fractal?"

Response:

- 1D cannot support isotropic propagation (only two directions: forward/backward)
- 3D fails the admissibility analysis (dilution trilemma)
- 2D is the minimum dimension supporting isotropy while avoiding volumetric scaling problems
- Fractal substrates could be considered but would need effective 2D scaling at large distances to preserve admissibility

Objection 5: "Does this introduce preferred frames or directions?"

Response: No. The interface update rate defines a universal speed limit, but no preferred *spatial* direction arises. Lorentz symmetry emerges as the continuum limit of isotropic discrete dynamics. Residual anisotropy (the 6Ω signal) is a higher-order effect, not a preferred frame.

Objection 6: "Is hexagonal geometry assumed or derived?"

Response: Derived. Four independent selection channels (dispersion symmetry, RG flow, information theory, BCB admissibility) each select hexagonal coordination. This convergence

from independent arguments is the strongest evidence that hexagonal geometry is required rather than chosen.

Objection 7: "How does this relate to quantum mechanics?"

Response: The framework is currently classical—it addresses field propagation structure. Extension to quantum dynamics is future work. However, the BCB principle has natural connections to quantum information theory (conservation of distinguishability \leftrightarrow unitarity), and the interface picture may offer new perspectives on measurement and decoherence.

Objection 8: "What about General Relativity?"

Response: GR remains the correct effective theory of emergent geometry. The framework is compatible with interpretations where geometry arises from causal/entropic structure (as in various quantum gravity approaches). Deriving Einstein's equations from interface dynamics is open work, though the correlation-based metric construction (Section 6) provides a starting point.

Objection 9: "Why accept BCB as fundamental?"

Response: BCB is not an extra dynamical postulate. It formalises what "finite information processing in a local causal runtime" means:

- Each update has bounded capacity
- Committing a distinction consumes capacity
- Consistent runtimes cannot allow the same fact to be committed at path-dependent cost without enabling arbitrage

Any physically realisable computation—classical or quantum—operates with bounded local processing per step. BCB simply requires that fundamental propagation respects this constraint at the fact-commitment level.

Rejecting BCB requires accepting either infinite information processing per update, inconsistent fact-accounting (path-dependent costs), or unexplained anisotropy. Each alternative has severe theoretical costs. BCB is "physics-as-runtime-consistency," not a new force law.

10. Open Questions and Future Directions

For the general reader: No theory is complete on day one. This section honestly acknowledges what we haven't yet worked out—and sketches where the research goes next. Incompleteness isn't a weakness; it's a sign that this is a genuine scientific programme with a research agenda, not a closed philosophical system.

10.1 Matter Coupling

The framework addresses photon propagation. Extension to matter fields requires:

- Characterising matter as localised interface excitations (topological defects? solitons?)
- Deriving fermionic statistics from interface topology
- Reproducing Standard Model particle spectrum

10.2 Gravitational Dynamics

Deriving Einstein's equations from interface dynamics is a major open problem. Promising directions:

- Interpreting spacetime curvature as variation in interface correlation structure
- Connecting to entropic gravity proposals
- Relating bulk reconstruction to tensor network models of AdS/CFT

10.3 Quantum Foundations

Potential connections to quantum mechanics:

- BCB conservation \leftrightarrow unitary evolution
- Interface update discreteness \leftrightarrow quantum measurement
- Emergent bulk \leftrightarrow quantum contextuality

10.4 Cosmological Implications

Interface-based propagation may affect:

- Early-universe physics (interface formation/phase transitions)
- Cosmological light propagation (cumulative dispersion over Gpc distances)
- Gamma-ray burst timing (high-energy dispersion tests)

10.5 Detailed Parameter Determination

Theoretical work needed to:

- Derive ε from first-principles interface microphysics
 - Determine whether $p = 4, 6$, or higher
 - Connect interface scale a to other fundamental scales
-

11. Conclusion

For the general reader: We've argued that the three-dimensional space we experience is not fundamental—it emerges from a two-dimensional causal interface with hexagonal structure. This isn't idle speculation: it follows from requiring that information processing be finite. The framework makes a specific, testable prediction (a six-fold pattern in precision light-speed measurements) that existing technology can probe. If we're right, it changes how we understand the relationship between information, causality, and the geometry of space. If we're wrong, the experiments will tell us.

We have presented a framework in which light propagation is fundamentally two-dimensional, occurring on a causal interface from which three-dimensional space emerges. The framework rests on a single deep principle: **reality has finite information capacity**. This is not an arbitrary assumption but an empirically motivated constraint (Bekenstein bounds, finite entropy, finite causal speeds).

The core insight: Finite capacity and perfect reversibility are incompatible. If capacity is finite, some distinctions must be discarded at some step. Discarding distinctions *is* irreversibility. Irreversibility *is* fact-commitment. BCB is not a new postulate—it is what finite capacity means. The alternative—insisting on fundamental reversibility via exact unitarity—requires infinite information capacity. That is the extraordinary claim, not BCB.

Principal results:

1. **Non-admissibility theorem:** Primitive 3D propagation violates consistency under finite-information constraints. The failure mode (dilution below resolution threshold) is specific to distinguishability and does not apply to energy.
2. **Hexagonal selection:** Four convergent arguments (dispersion symmetry, RG flow, information theory, BCB) point to hexagonal coordination as the unique admissible discrete geometry.
3. **Emergent field theory:** Lorentz symmetry and Maxwell electrodynamics emerge as continuum limits of interface dynamics, preserving all empirical successes of standard physics.
4. **Bulk reconstruction:** Three-dimensional space emerges from interface correlations across coarse-graining scales, with flat isotropic geometry at leading order.
5. **Falsifiable prediction:** A distinctive six-fold angular harmonic (6Ω) in rotating optical cavity experiments, distinguishable from competing geometries (4Ω for square lattice).

The framework explains nine structural features of light that standard physics describes but does not derive. It is not speculative metaphysics but a falsifiable physical proposal.

The logical chain: The argument proceeds through a closed chain: (1) The physical world exhibits finite information capacity (Bekenstein bounds, finite entropy). (2) Finite capacity requires discarding distinctions. (3) Discarding distinctions *is* irreversibility. (4) Irreversibility *is* fact-commitment. (5) Fact-commitment is subject to BCB constraints. (6) Under BCB, primitive

3D propagation is non-admissible. (7) A 2D interface with hexagonal symmetry is uniquely selected. (8) Residual micro-geometry should be observable as a 6Ω harmonic.

An alternative route (same destination): If time is emergent rather than fundamental—a position most quantum gravity approaches take—then facts are already presupposed. Emergent time *is* fact-commitment viewed from another angle: time emerges from sequential closure, from "this happened, then that happened," from the accumulation of settled matters. If nothing ever became settled, there would be no sequence, no before/after, no time. Thus: emergent time \rightarrow facts \rightarrow BCB constrains their structure \rightarrow finite capacity gives rate limits \rightarrow 2D hexagonal. Both routes arrive at the same destination. The framework is overdetermined by independent philosophical commitments that most physicists already accept.

What we are not claiming: We are not defending an unusual ontological choice. We are taking seriously what everyone already implicitly assumes when they say "the experiment showed X." The framework that requires defense is the one claiming infinite capacity—that the universal wavefunction preserves all phase correlations forever at perfect precision. BCB takes finitude seriously and follows the consequences.

The bet this framework makes: Detection of 6Ω would vindicate finite-capacity fact-commitment and constrain interface parameters. Continued null results at increasing sensitivity would either push the framework toward smaller scales / stronger suppression, or shift probability toward infinite-capacity amplitude-primitive interpretations. Either outcome advances our understanding of the relationship between information, causality, and spacetime geometry.

Remaining work lies in matter coupling, gravitational dynamics, and detailed parameter determination. The foundational structure is complete and internally consistent.

References

1. S. Herrmann et al., "Rotating optical cavity experiment testing Lorentz invariance at the 10^{-17} level," *Phys. Rev. D* 80, 105011 (2009); arXiv:1002.1284
2. T. Zhang et al., "Test of Lorentz invariance using rotating ultra-stable optical cavities," *Physics Letters A* 416, 127666 (2021)
3. V.A. Kostelecký and N. Russell, "Data tables for Lorentz and CPT violation," *Rev. Mod. Phys.* 83, 11 (2011); updated regularly at arXiv:0801.0287 (most recent: January 2025)
4. R.D. Sorkin, "Causal sets: Discrete gravity," in *Lectures on Quantum Gravity*, eds. A. Gomberoff and D. Marolf (Springer, 2005); arXiv:gr-qc/0309009
5. C. Rovelli, "Loop quantum gravity," *Living Rev. Relativ.* 11, 5 (2008)
6. S. Wolfram, "A class of models with the potential to represent fundamental physics," arXiv:2004.08210 (2020)
7. J. Maldacena, "The large N limit of superconformal field theories and supergravity," *Adv. Theor. Math. Phys.* 2, 231 (1998); arXiv:hep-th/9711200

8. B. Swingle, "Entanglement renormalization and holography," Phys. Rev. D 86, 065007 (2012); arXiv:0905.1317
9. K.G. Wilson, "Confinement of quarks," Phys. Rev. D 10, 2445 (1974)
10. J. Kogut and L. Susskind, "Hamiltonian formulation of Wilson's lattice gauge theories," Phys. Rev. D 11, 395 (1975)
11. T. Jacobson, "Thermodynamics of spacetime: The Einstein equation of state," Phys. Rev. Lett. 75, 1260 (1995); arXiv:gr-qc/9504004
12. E. Verlinde, "On the origin of gravity and the laws of Newton," JHEP 04, 029 (2011); arXiv:1001.0785
13. S. Liberati, "Tests of Lorentz invariance: A 2013 update," Class. Quantum Grav. 30, 133001 (2013); arXiv:1304.5795
14. M. Creutz, *Quarks, Gluons and Lattices* (Cambridge University Press, 1983)
15. J. Kogut, "The lattice gauge theory approach to quantum chromodynamics," Rev. Mod. Phys. 55, 775 (1983)
16. R. Landauer, "Irreversibility and heat generation in the computing process," IBM J. Res. Dev. 5, 183 (1961)

Appendix A: Glossary of Key Terms

Term	Definition
Admissibility	Property of propagation processes satisfying finite budget, locality, path-independence, and isotropy
BCB	Bit Conservation & Balance—the principle that distinguishability is conserved per causal update
Distinguishability	The capacity to resolve whether one state of affairs obtains rather than another
Interface	The 2D causal substrate on which fundamental propagation occurs
Resolution threshold	Minimum distinguishability (Δi_{\min}) required to resolve any binary fact
Tick	Minimal causal update step with bounded processing capacity

Appendix B: Selection Channel Summary

Channel	Analysis method	What it shows	Selected geometry
Dispersion symmetry	Fourier expansion of discrete Laplacian	Hex pushes anisotropy to 6th order	Hexagonal
Renormalization flow	Scaling of anisotropy operators	Hex: $g \sim b^{-4}$ vs Square: $g \sim b^{-2}$	Hexagonal

Channel	Analysis method	What it shows	Selected geometry
Information theory	Angular sampling analysis	Hex: 60° gaps vs Square: 90° gaps	Hexagonal
BCB/Admissibility	Directional variance calculation	Hex minimises variance	Hexagonal

Appendix C: Experimental Bounds

Current sensitivity: $\Delta c/c \sim 10^{-17}$ (Herrmann et al. 2009; Zhang et al. 2021)

Implied bounds on interface scale (optical wavelength $\lambda = 1064$ nm, assuming $\epsilon = 1$):

Suppression order p	Maximum a (m)	Energy scale hc/a (eV)
2	5.4×10^{-16}	3.7×10^8
4	9.5×10^{-12}	2.1×10^4
6	2.5×10^{-10}	7.9×10^2

Appendix D: Rejection Costs and Alternative Ontologies

This appendix clarifies the scope of the non-admissibility theorem and addresses common strategies by which one might attempt to reject its conclusions. In each case, we show that the move does not refute the argument but instead constitutes a change of ontological framework with distinct methodological or empirical consequences.

The purpose of this appendix is not to rule out alternative ontologies, but to make explicit what must be given up—and what new commitments are incurred—if one rejects the assumptions under which the theorem applies.

D.1 Rejecting Counterfactual Causation

Claim: One may reject counterfactual dependence ("if A were different, B would differ") as part of what causation means, treating causality instead as mere temporal ordering or a static mathematical relation.

Response: The framework adopts counterfactual dependence as a minimal operational criterion for physically meaningful causation. This criterion is not philosophical embellishment; it is implicit in experimental practice. Physical causal claims are validated through interventions: rotating a cavity, shielding a detector, changing a source configuration, randomizing settings. If varying A under controlled conditions never alters the distribution of B, then A does not cause B in the sense required for explanation, prediction, or experimental control.

Cost: Rejecting counterfactual causation avoids the record-necessity argument, but at the price of abandoning the operational meaning of causal explanation. This is not a refutation of the theorem; it is an exit from its domain of applicability.

D.2 Allowing Nonlocal, Retrocausal, or Superdeterministic Coordination

Claim: One may allow distinguishability resources to be "pre-positioned" at future wavefront locations via nonlocal coordination, retrocausality, or superdeterministic correlations, thereby evading the budget-scaling argument.

Response:

Nonlocal pre-positioning: Pre-allocating distinguishability to distant registers in anticipation of a future emission requires those registers to depend on the source's future behavior before any causal influence has reached them. This is operationally equivalent to superluminal broadcast and explicitly violates locality.

Retrocausality: Retrocausal frameworks can coordinate present resources using future boundary conditions. Such models are internally consistent but must explain why macroscopic interventions behave as if settings are free, why effective causal arrows appear forward-directed, and why no observable signaling artifacts arise. This constitutes a distinct causal ontology with separate empirical commitments.

Superdeterminism: Superdeterministic models evade all experimental discriminators by postulating hidden correlations between settings and hidden variables. While logically possible, this weakens ordinary statistical inference unless supplemented with independent testable constraints.

Cost: Allowing nonlocal, retrocausal, or superdeterministic coordination avoids the non-admissibility theorem only by rejecting locality or standard experimental independence. This is a shift to a non-local or non-interventional ontology whose empirical status must be assessed independently.

D.3 Amplitude-Primitive Ontologies (Everettian or Pilot-Wave)

Claim: Fundamental dynamics consists solely of continuous amplitude evolution (unitary quantum mechanics or pilot-wave dynamics), and "facts" or "records" are merely emergent descriptions. In this view, Δi_{\min} is unnecessary.

Response: Amplitude-primitive frameworks do not eliminate records; they relocate them:

- Everettian approaches yield stable records as decohered pointer states with redundancy across many degrees of freedom.
- Pilot-wave theories yield stable records via an actual configuration that selects outcomes.

In both cases, causal explanation at the level of observed physics is mediated by robust, record-like structures. Such records generically require redundancy, error suppression, or stability criteria—implying effective thresholds even if no fundamental constant is introduced.

The reversibility problem (see Section 3.1): A deeper issue is that purely amplitude-primitive dynamics are fundamentally reversible—time-reversal invariant, unitary, information-preserving. But the world exhibits ontological irreversibility: detectors click and cannot un-click; experiments have outcomes, not just correlations. Decoherence explains why interference becomes unobservable, but decoherence is still unitary and does not select one outcome as *the* outcome. If the world contains real irreversible facts, fact-commitment must be realised somewhere. Amplitude-primitive frameworks defer this problem to emergence; they do not remove it.

The key distinction is ontological:

- Amplitude-primitive frameworks may allow exact Lorentz symmetry (or exact erasure of microscopic residue) in the ultraviolet.
- Record-primitive frameworks permit symmetry-protected residual structure from the causal substrate.

Cost: Adopting an amplitude-primitive ontology avoids treating Δi_{\min} as fundamental, but incurs the obligation to explain why emergent record structures erase all microscopic lattice residue without leaving symmetry-protected signatures. The predicted *absence* of a 6Ω harmonic is therefore a nontrivial empirical claim of amplitude-primitive approaches, not a neutral default.

Direct engagement with Everettian many-worlds: Everettian many-worlds relocates facts to branch-relative statements: within each branch, records exist and observers experience definite outcomes. The present framework does not claim Everett is internally inconsistent—it is a coherent position. Rather, the two approaches make different empirical predictions:

Framework	Ontology	Micro-geometry	6Ω prediction
Everettian	Amplitude-primitive, exact unitarity	No residual lattice structure	No 6Ω (exact Lorentz in UV)
Interface (this paper)	Fact-primitive, discrete substrate	Symmetry-protected hexagonal residue	6Ω at $(ka)^p$

Sharpening the disagreement: The distinction is not merely about where facts "live." In Everett, there literally is no fact of the matter about which outcome occurred—only relative facts (what observer O records in branch B). All branches are equally real; none is privileged. Our framework asserts something stronger: that committed facts are substrate-level primitives with *definite truth values* independent of observers or branches. The causal bit Δi_{\min} is not branch-relative—it is a property of the fundamental substrate.

Everett's implicit commitment to infinite capacity: The Everettian insistence on exact unitarity—no information ever discarded, all branches preserved with all phase correlations at perfect precision forever—is an implicit commitment to infinite information capacity. The universal

wavefunction must contain every branching possibility with every micro-correlation intact, recoverable in principle at arbitrary future times. This is not a neutral default position; it is an extraordinary claim about the information-processing capacity of reality.

Our framework makes the opposite bet: reality has finite capacity. Given finite capacity, irreversibility follows (Section 2.2). Given irreversibility, fact-commitment follows. The question is not "should we accept facts as primitive?" but "can we coherently deny facts while accepting finite capacity?" We claim not.

This is a sharper disagreement than "relocation of facts." Everett denies that any global fact-of-the-matter exists *and* implicitly claims infinite capacity to make this coherent. We assert that capacity is finite, that irreversibility follows, and that fact-commitment leaves observable residue. Detection of 6Ω would be evidence that the universe commits facts rather than merely superposing branches at infinite precision.

Detection of 6Ω would strongly disfavour Everett and other amplitude-primitive completions; continued null results at increasing sensitivity would shift probability toward amplitude-primitive interpretations (or toward smaller interface scales / stronger suppression). The choice between frameworks is therefore empirical, not purely philosophical. We are not claiming to refute Everett by argument; we are claiming the two frameworks make different predictions that experiment can distinguish.

D.4 Summary of Rejection Costs

Rejection move	What it avoids	What it commits you to
Reject counterfactual causation	Record necessity	Causality becomes non-operational
Allow nonlocal pre-positioning	Budget scaling	Superluminal/global coordination
Allow retrocausality	Budget scaling	Nonstandard intervention logic, arrow-of-time explanation
Invoke superdeterminism	Any discriminator	Weakened statistical inference
Adopt amplitude-primitive ontology	Δi_{\min} as primitive	Emergent records must erase all UV residue

D.5 Scope Clarification

The non-admissibility theorem applies to the class of theories in which:

1. Causal relations are operationally meaningful (counterfactual)
2. Update structure is forward-local
3. Facts are primitive constraints on future evolution

Within this class, primitive three-dimensional propagation is inconsistent and a two-dimensional causal substrate is required.

The purpose of the framework is not to exclude alternative ontologies, but to show that within this well-defined and widely used causal class, the interface picture is forced and leaves a symmetry-protected experimental discriminator.

Appendix E: Assumptions Ledger

For clarity and to facilitate critical evaluation, we explicitly categorise the assumptions underlying this framework.

E.1 Core Assumptions (Required for Main Results)

These assumptions are necessary for the non-admissibility theorem and hexagonal selection:

Assumption	Statement	Section
BCB	Distinguishability is conserved per causal update with finite capacity	§2.2
Locality	Causal updates affect only immediately adjacent registers	§2.5
Path-independence	Resolution cost is independent of the path taken	§2.5
Causal bit	The minimum recordable distinction is one binary constraint	§2.4.1
Counterfactual causation	Causal relations support interventionist counterfactuals	Box after §2.4.1

Rejecting any core assumption exits the framework's domain of applicability (see Appendix D for rejection costs).

E.2 Auxiliary Assumptions (Required for Specific Predictions)

These assumptions are not required for the main theorems but are needed for particular empirical predictions:

Assumption	Statement	What depends on it
A_ε	ε is not exponentially suppressed (generic substrate)	Amplitude-based falsification
Finite-range operator	Propagation operator has finite-range coupling	Harmonic order selection

Assumption	Statement	What depends on it
Coarse-grained wavefront	Substrate admits well-defined effective wavefront scaling (Definition 9, §3.2)	Non-admissibility scope
U(1) link variables	Gauge structure is U(1) on interface links	Maxwell limit, gauge redundancy

Auxiliary assumptions can be relaxed or modified without invalidating core results.

E.3 Speculative Components (Research Directions)

These components are included to indicate explanatory directions but are not claimed as derivations:

Component	Status	Section
Bulk reconstruction	Proposal only; rigorous mechanism not yet established	§6
Emergent 3D geometry	Conditional on bulk reconstruction	§7.8
Light-geometry link	Conditional on bulk reconstruction	§7.9
Gauge group selection	U(1) assumed, not derived	§5.2

Failure of speculative components does not affect core theorems or the 6Ω prediction.

E.4 What Is Not Assumed

For clarity, we note what the framework explicitly does *not* assume:

- **No pre-existing spacetime:** Geometry is derived, not postulated
- **No continuous symmetry:** Lorentz symmetry emerges from discrete dynamics
- **No specific interface scale:** The scale a is constrained by experiment, not assumed
- **No rigid lattice:** The theorem applies to any substrate with effective wavefront scaling

Appendix F: Overlap Uniqueness from BCB Geometry (Sketch)

This appendix provides a brief derivation sketch for the claim in Section 6.7 that the Fisher–Bhattacharyya overlap with phase modulation is the unique overlap functional compatible with BCB constraints.

Setup: Consider a space of distinguishability distributions over a finite set of outcomes, equipped with a notion of "overlap" or similarity $O(p, q)$ between distributions p and q .

Constraint 1 (Monotonicity under coarse-graining): If we partition outcomes into coarser bins, the overlap of the coarse-grained distributions must not exceed the original overlap. Formally, for any coarse-graining map π :

$$O(\pi(p), \pi(q)) \geq O(p, q)$$

This reflects the information-theoretic intuition that coarse-graining cannot increase distinguishability.

Constraint 2 (Composition stability): For independent systems, the overlap of product distributions must equal the product of overlaps:

$$O(p_1 \otimes p_2, q_1 \otimes q_2) = O(p_1, q_1) \cdot O(p_2, q_2)$$

This ensures consistent behaviour under system combination.

Constraint 3 (Gauge invariance): If the distinguishability manifold carries a compact $U(1)$ fibre (phase degree of freedom), the overlap must be invariant under global phase rotations of either argument.

Result: The unique overlap functional satisfying all three constraints is the Bhattacharyya coefficient with phase:

$$O(p, q; \varphi) = |\sum_i \sqrt{p_i q_i} e^{i\varphi_i}|^2$$

In the continuum limit with a Fisher metric on the base manifold, this reduces to the standard quantum inner product structure $\langle \psi | \varphi \rangle$.

Proof sketch: Monotonicity and composition stability together force the overlap to have the form of an f-divergence or Rényi-type functional. The composition requirement specifically selects the geometric mean (Bhattacharyya form). The $U(1)$ gauge constraint then forces the phase-modulated version. Full details are developed in companion work on BCB geometry foundations.

This uniqueness result implies that correlation functions on the interface are not free parameters but are constrained by BCB geometry to have quantum-like structure.

Appendix G: A Minimal Toy Model of Finite-Capacity Fact-Commitment Propagation

This appendix provides an explicit toy model in which the central ideas of the paper can be instantiated with minimal assumptions: (i) local, finite-capacity updates ("BCB runtime"), (ii) binary fact-commitment with a threshold, and (iii) an emergent long-wavelength propagation law

whose anisotropy is determined by lattice point-group symmetry. The purpose is not realism but *witness*: to show that the non-admissibility theorem and the 6Ω discriminator arise in a concrete system.

G.1 Registers, Ticks, and the Causal Bit

Let the substrate be a graph $G = (V, E)$. Each vertex $v \in V$ is an addressable causal register capable of holding a binary record $R_v \in \{0, 1\}$. A tick is a discrete global update step $t \rightarrow t + 1$ in which registers may change state.

Causal bit threshold: A register can commit at most one record per tick, and committing a record requires at least one unit of "commitment resource." We therefore take:

$$\Delta i_{\min} = 1 \text{ (one causal bit per committed record)}$$

This is a logical unit; the toy model represents the implementation cost as a limited per-tick budget (below).

G.2 Finite Budget and Locality (BCB-Runtime)

Introduce a per-source per-tick budget:

$$\Delta I_{\max} = B < \infty$$

interpreted as the maximum number of causal-bit commitments a localised emission event can enable per tick, consistent with finite capacity.

Locality constraint: Resource can move only along edges of G . No register may receive resource from a vertex not connected by an edge chain that has had time to propagate.

This defines a conservative "flow" picture: resource is neither created nor destroyed except at sources (injection) and sinks (commitment). The model's "BCB" is simply finite resource throughput per tick + local transport.

G.3 Propagation Rule (Minimal)

We model causal influence as a resource-carrying wavefront emanating from a localised source s . Let $I_t(v) \geq 0$ be the commitment resource stored at register v at tick t . The update is:

(1) Injection: At each tick t , the source injects B units:

$$I_t(s) \leftarrow I_t(s) + B$$

(2) Local transport: Each vertex splits its resource equally among neighbours:

$$I_{t+1}(u) \leftarrow \sum_{\{v: (v,u) \in E\}} I_t(v) / \deg(v)$$

(Any finite-range local transport rule works; equal splitting is the simplest.)

(3) Fact-commitment (thresholding): A register commits a fact only if it receives at least one unit in that tick:

$$R_{t+1}(v) = 1 \text{ if } I_{t+1}(v) \geq 1; \text{ otherwise } R_t(v)$$

Optionally, one can also consume one unit upon commitment:

$$I_{t+1}(v) \leftarrow I_{t+1}(v) - 1 \text{ if } I_{t+1}(v) \geq 1$$

to enforce that commitments use budget.

This is a literal implementation of "fact-commitment threshold": sub-threshold influence may exist (resource can be nonzero), but no record is created unless threshold is met.

G.4 Dimensional Scaling and the 3D Failure

Define the reachable shell at tick r as the set of registers at graph distance r from the source:

$$S(r) = \{v \in V : d(s, v) = r\}$$

Let $N(r) = |S(r)|$ be the number of registers on the wavefront at radius r .

For any substrate whose coarse-grained propagation is effectively D -dimensional, we have the standard scaling:

$$N(r) \propto r^{D-1}$$

Upper bound on total available committed bits by radius r : By locality, only budget injected along causal chains from the source can reach $S(r)$. Since the source injects at most B per tick, after r ticks the total available budget that could have reached the shell is at most:

$$I_{\text{acc}}(r) \leq B \cdot r$$

(This inequality is conservative: it assumes perfect transport and zero loss.)

Necessary condition for committing facts across the whole shell: To commit at least one record on each register in $S(r)$, the shell must receive at least $N(r)$ causal bits:

$$I_{\text{acc}}(r) \geq N(r)$$

Therefore a necessary condition is:

$$B \cdot r \gtrsim r^{D-1} \Rightarrow B \gtrsim r^{D-2}$$

For $D = 3$ this becomes $B \gtrsim r$, which cannot hold for all r if B is finite. Thus there exists a failure radius $r^* \sim B$ beyond which the wavefront cannot commit one causal bit per register across the sphere. The failure is not gradual: it is categorical because commitment requires crossing a discrete threshold.

This is the toy-model instantiation of the paper's non-admissibility claim: **finite per-tick budget + local transport + binary commitment \Rightarrow primitive 3D wavefront commitment fails at large radius.**

G.5 Why 2D Survives (and Why 1D Fails Isotropy)

For $D = 2$ we have $N(r) \propto r$, so the condition becomes:

$$B \cdot r \gtrsim r \Rightarrow B \gtrsim \text{constant}$$

Thus a finite budget B can maintain one-bit-per-register commitment indefinitely in 2D.

For $D = 1$, $N(r) = 2$ does not grow, so commitment is easy, but isotropy is not meaningful (only two directions).

Hence, within the toy-model class: $D \leq 1$ fails isotropy, $D \geq 3$ fails budget, leaving **$D = 2$ as the maximal admissible dimension.**

G.6 Hexagonal Micro-Geometry and the 6Ω Residue

Now choose G to be the triangular/hexagonal lattice in the plane (6 nearest neighbours, point group C_6). Consider the linear transport operator (step (2)) in Fourier space. For a translation-invariant local operator with finite range R ,

$$(L\psi)(x) = \sum_{\{r \in R\}} w_r [\psi(x+r) - \psi(x)]$$

the Fourier symbol is:

$$\Lambda(k) = \sum_{\{r \in R\}} w_r [1 - \cos(k \cdot r)]$$

Expanding at small $|k|$ gives an isotropic leading term plus symmetry-allowed anisotropic corrections. The key fact is purely group-theoretic:

- On a **square lattice** (C_4), the lowest nontrivial angular harmonic compatible with the symmetry is $\cos(4\theta)$, so anisotropy generically appears at $|k|^4 \cos(4\theta)$.
- On a **triangular/hex lattice** (C_6), harmonics with $m = 1, 2, 3, 4, 5$ are forbidden. The first allowed anisotropic harmonic is $\cos(6\theta)$, and for generic symmetric couplings it appears at order $|k|^6 \cos(6\theta)$.

Thus the long-wavelength dispersion takes the form:

$$\omega^2(\mathbf{k}) = c^2|\mathbf{k}|^2 + \alpha|\mathbf{k}|^4 + \beta|\mathbf{k}|^6 + \gamma|\mathbf{k}|^6 \cos(6\theta) + O(|\mathbf{k}|^8)$$

where θ is the direction of \mathbf{k} relative to a fixed lattice axis.

Interpreting the directional dependence as a small fractional modulation of effective propagation speed:

$$\delta c/c(\theta) \propto (ka)^p \cos(6\theta)$$

with baseline $p = 6$ for the ideal hex-symmetric operator.

Under rotation at angular rate Ω , $\theta(t) = \Omega t + \theta_0$, so a rotating-cavity experiment would see a harmonic at 6Ω :

$$\delta f/f(t) \sim A_6 \cos(6\Omega t + \varphi)$$

This is the toy-model origin of the paper's core discriminator: $C_6 \Rightarrow 6\Omega$ vs $C_4 \Rightarrow 4\Omega$.

G.7 Minimal Numerical Experiment (Recipe)

The toy model is directly simulatable without adding further physics:

1. Construct a finite patch of a triangular lattice with periodic boundary conditions (to avoid edge effects).
2. Initialise $I_0(\mathbf{v}) = 0$ for all \mathbf{v} ; choose a source s .
3. Iterate the update rule (injection \rightarrow transport \rightarrow threshold commitment).
4. Track (i) how many registers on the shell $S(r)$ commit by tick r , and (ii) the anisotropy of propagation by measuring arrival times or commitment probabilities as a function of angle.

A simple metric is the angular Fourier decomposition of arrival/commitment:

$$A_m(r) = |S(r)|^{-1} \left| \sum_{\mathbf{v} \in S(r)} e^{im\theta_{\mathbf{v}}} \chi_{\mathbf{v}}(r) \right|$$

where $\chi_{\mathbf{v}}(r) \in \{0, 1\}$ indicates whether \mathbf{v} has committed by tick r , and $\theta_{\mathbf{v}}$ is the angle of \mathbf{v} relative to the source.

Prediction: $A_6 \neq 0$ with $A_4 \approx 0$ (ideal hex case), and the reverse for square-lattice substrates.

G.8 What the Toy Model Does and Does Not Claim

This toy model does not claim to be electromagnetism, QFT, or a realistic quantum gravity microtheory. It exists to show three structural points with explicit rules:

1. **Finite per-tick capacity + local transport + thresholded record creation implies a dimensional admissibility constraint** ($D = 2$ survives; $D = 3$ fails).

2. **The "register counting" argument is not mere metaphor**; it is a strict necessary condition in an explicit system.
3. **Hexagonal symmetry implies a specific six-fold angular harmonic** as the leading symmetry-allowed anisotropy residue.

These are exactly the three load-bearing steps the main paper relies on. The toy model serves as an existence proof that such systems are coherent and that the predictions follow from the stated assumptions.

Appendix H: A Lattice-Field Toy Model with Explicit Correlation Functions

This appendix provides a second toy model, complementary to Appendix G, in which correlations are explicitly computable using standard lattice field theory tools. The goal is not to reproduce full electromagnetism but to demonstrate that, once one specifies an action on a 2D hexagonal interface, correlation functions $\langle O(x)O(y) \rangle$ follow in a controlled and conventional way. This addresses the common referee concern: "show me a model where you can actually compute correlations."

H.1 Why a Gauge+Matter Toy Model

A technical subtlety matters here: pure compact $U(1)$ gauge theory in two Euclidean dimensions has no local propagating photon degrees of freedom (it is essentially topological; its excitations are global/holonomy-like). Consequently, if one wants nontrivial local two-point functions that behave like field propagation and generate a correlation geometry, one must either:

- Include matter (a charged field), or
- Include an additional phase/ XY field that couples to the gauge link, or
- Study non-compact gauge variables with a gauge-fixing choice

For a minimal, fully standard construction, we use compact $U(1)$ link variables plus a charged phase field. This yields an explicit, gauge-invariant two-point correlator with calculable decay.

H.2 Interface Degrees of Freedom and Action

Let the interface be a triangular/hexagonal lattice with sites x and oriented nearest-neighbour links $\ell = (x \rightarrow x + \mu)$. Define:

Link variables: $U_\ell \in U(1)$, $U_{\{x,\mu\}} = e^{iA_{\{x,\mu\}}}$

Site phase field ("matter"): $\varphi_x \in U(1)$, $\varphi_x = e^{i\theta_x}$

We define the Euclidean partition function:

$$Z = \int \prod_x d\phi_x \prod_\ell dU_\ell \exp(-S[U, \phi])$$

A minimal gauge+matter action is:

$$S[U, \phi] = S_{\text{gauge}}[U] + S_{\text{matter}}[U, \phi]$$

with:

(i) Gauge (Wilson plaquette) term

$$S_{\text{gauge}}[U] = \beta \sum_p (1 - \text{Re } U_p)$$

where p runs over elementary plaquettes (hexagons on the dual / or minimal loops on the triangular lattice), and $U_p = \prod_{\ell \in \partial p} U_\ell$.

(ii) Matter (gauged XY) term

$$S_{\text{matter}}[U, \phi] = -\kappa \sum_{\langle x,y \rangle} \text{Re}(\phi_x^* U_{\{xy\}} \phi_y)$$

where $\langle x, y \rangle$ are nearest neighbours and $U_{\{xy\}}$ is the link variable from x to y .

This is a standard lattice construction with exact local $U(1)$ gauge invariance under:

$$\phi_x \rightarrow e^{i\alpha_x} \phi_x, U_{\{xy\}} \rightarrow e^{i\alpha_x} U_{\{xy\}} e^{-i\alpha_y}$$

H.3 A Gauge-Invariant Two-Point Correlation Function

A natural gauge-invariant correlator is the "string-dressed" two-point function:

$$G(x, y) = \langle \phi_x^* (\prod_{\ell \in \Gamma_{\{x \rightarrow y\}}} U_\ell) \phi_y \rangle$$

where $\Gamma_{\{x \rightarrow y\}}$ is any lattice path from x to y . Gauge invariance is manifest because the phase factors at internal nodes cancel.

Path-independence criterion: In a confining or highly disordered phase, $G(x, y)$ depends strongly on the chosen path Γ . In a deconfined / weak-fluctuation regime, $G(x, y)$ becomes (approximately) path-independent up to small fluctuations, and its decay depends primarily on separation $|x - y|$. This provides a concrete toy instantiation of the paper's "path-independence" requirement.

H.4 Correlation Decay and Correlation Length

Even in this simple model, correlation behaviour is analytically tractable in two canonical limits.

H.4.1 Strong-coupling / disordered regime (small β, κ)

At small β (strong gauge fluctuations) and small κ (weak matter coupling), one can expand the Boltzmann weight in characters. The result is that gauge-invariant correlators decay exponentially with separation:

$$G(x, y) \sim \exp(-|x - y|/\xi)$$

with a correlation length ξ controlled by κ (and suppressed further by strong gauge disorder). This gives a direct, conventional mechanism by which finite-range correlations emerge from an explicit action.

Interpretation: This regime corresponds to a "high-noise" interface in which local commitments do not propagate far.

H.4.2 Weak-coupling / spin-wave regime (large β , sufficiently large κ)

For large β , link fluctuations are small and one may write $U_{xy} \approx e^{iA_{xy}}$ with A_{xy} slowly varying. For sufficiently ordered matter (κ large), one obtains a continuum Gaussian effective action for the phase field (a "spin-wave" approximation):

$$S_{\text{eff}} \approx (K/2) \int d^2x (\nabla\theta - A)^2 + (1/4g^2) \int d^2x F^2 + \dots$$

where $F = \nabla \times A$. In this regime, two-point functions exhibit long-distance behaviour controlled by standard continuum field theory; depending on whether vortices proliferate, decay can be power-law (quasi-long-range order) or exponential (mass gap). In either case, the model yields an explicit, computable correlation structure.

Interpretation: This regime instantiates an "ordered interface" consistent with macroscopic isotropy and long-distance propagation.

H.5 Hexagonal Symmetry and the 6-Fold Residue in Correlations

Because the lattice point group is C_6 , any continuum expansion of correlation kernels or dispersion relations inherits symmetry selection rules. In particular, the leading allowed anisotropic correction to isotropic behaviour is six-fold:

$$G(x, y) = G_0(r)[1 + \varepsilon(r) \cos(6\theta) + \dots]$$

where $r = |x - y|$ and θ is the direction of $y - x$ relative to a fixed lattice axis. Lower harmonics ($\cos(2\theta)$, $\cos(4\theta)$) are forbidden unless additional symmetry-breaking (boundaries, domains, anisotropic couplings) is introduced.

Thus, even at the level of correlation functions—not only at the level of dispersion operators—the hexagonal substrate predicts the same qualitative discriminator emphasised in the main text: **six-fold angular structure is the leading symmetry-allowed residue.**

H.6 Connecting to Bulk Reconstruction

Once $G(x, y)$ is defined and computed (in either regime), one can define a correlation distance exactly as in Section 6:

$$d(x, y) = -\ell \log[G(x, y)/G(x, x)]$$

This gives a fully explicit route from:

$$S[U, \phi] \Rightarrow G(x, y) \Rightarrow d(x, y) \Rightarrow \text{emergent correlation geometry}$$

The main paper does not claim that this particular action is "the" microphysics of the interface. Rather, Appendix H shows that if the interface admits an ordinary local lattice action (as many emergent-spacetime programmes assume), then correlations and induced geometry are computable in the standard way. Appendix G provides the complementary route: correlations defined directly from fact-commitment ancestry without importing amplitude dynamics.

H.7 What This Toy Model Does and Does Not Claim

Does:

- Provide an explicit $S[U, \phi]$, a standard path integral, and a gauge-invariant two-point correlator $G(x, y)$ with controllable decay regimes and symmetry constraints
- Demonstrate that hexagonal symmetry implies a leading $\cos(6\theta)$ angular residue in correlation structure
- Show that correlations are computable using standard lattice field theory methods

Does not:

- Claim to reproduce 3+1D Maxwell theory (in 2D Euclidean space pure $U(1)$ gauge has no propagating photon; matter is included to generate meaningful local correlators)
- Fix the physical parameter mapping to (ϵ, a, p) of the main text; it provides a computable laboratory for how such parameters arise from explicit microphysics

H.8 Complementarity of Appendices G and H

Aspect	Appendix G	Appendix H
Framework	Fact-commitment / resource flow	Lattice field theory / path integral
Correlations from	Common causal ancestry	Action $S[U, \phi]$
Threshold mechanism	Explicit commitment resource	Implicit in gauge coupling
Demonstrates	BCB \Rightarrow dimensional constraint	Standard methods yield correlations
Key result	3D fails, 2D survives	$G(x, y)$ computable, $C_6 \Rightarrow \cos(6\theta)$

Together, the two toy models show that the framework's claims are instantiable from both directions: from an information-theoretic / fact-commitment perspective (G) and from a

conventional field-theoretic perspective (H). Neither claims to be the actual microphysics; both serve as existence proofs that the relevant structures are coherent.

Appendix I: Scope, Ontological Forks, and Limiting Assumptions

This appendix clarifies the scope of the non-admissibility theorem, addresses several deep ontological tensions raised by alternative frameworks, and delineates precisely which conclusions are **derived**, **conditional**, or **interpretive**. Its purpose is not to eliminate philosophical disagreement, but to make explicit where disagreement must occur and what empirical consequences follow.

I.1 The Scope of BCB and the Status of Causal Sets

The non-admissibility theorem applies to a specific but broad class of substrates: **deterministic or quasi-deterministic causal runtimes that admit a well-defined coarse-grained wavefront structure** (Definition 9). Within this class, primitive three-dimensional propagation is inconsistent with finite information capacity and locality.

Causal set theory (CST) lies largely *outside* this class by construction. Its defining features—random Poisson sprinkling, lack of fixed adjacency, and ensemble-level Lorentz invariance—mean that there is no deterministic wavefront whose cardinality scales as r^{D-1} in individual realisations. As such, CST avoids the dilution argument not by solving it, but by **declining to instantiate the conditions under which it arises**.

A natural challenge follows:

If random geometry substrates evade the dilution problem, why not prefer them? What principle forces fixed or quasi-fixed geometry rather than stochastic structure?

Our response is twofold.

First, the present work does not claim that *all* physically viable substrates must fall within the BCB-admissible class analysed here. Rather, it claims that **within this class**, dimensionality and geometry are sharply constrained. CST represents a different ontological choice with different empirical commitments.

Second, the choice between fixed-geometry and stochastic substrates is not neutral. Random substrates achieve Lorentz invariance statistically, but at the cost of:

- Sacrificing exact locality at the micro-level (adjacency is probabilistic)
- Replacing path-independence with ensemble-averaged behaviour
- Eliminating symmetry-protected residue in individual realisations

These differences have empirical implications. In particular:

- Fixed-geometry substrates generically predict suppressed but structured residual anisotropy (here, a 6Ω harmonic).
- Stochastic substrates predict exact Lorentz invariance with no such residue.

The present framework therefore **does not refute causal sets**; it **sets up a clean empirical fork**. Detection of a symmetry-protected angular harmonic would disfavour random-geometry approaches; continued null results at increasing sensitivity favour them. The role of BCB is to show that *if* one insists on deterministic, local, finite-capacity runtimes, then random geometry is not optional—it is the only way to evade the dilution problem without reducing dimensionality.

I.2 BCB vs Everett: Capacity, Unitarity, and Empirical Forks

The claim that amplitude-primitive (Everettian) ontologies are “committed to infinite information capacity” is intentionally strong and risks being overstated if interpreted incautiously. A more precise formulation is as follows.

Everettian quantum mechanics preserves **exact unitarity** and **exact reversibility** of the universal wavefunction. In any finite region, the von Neumann entropy is finite, and holographic bounds apply to quantum states as well as classical configurations. In this sense, Everettian frameworks are not committed to infinite *entropy density*.

However, exact unitarity also implies that:

- No information is ever discarded, even in principle
- Phase relations between branches are preserved indefinitely
- All branch structure remains recoverable at arbitrary future times *in principle*

This is not a statement about practical accessibility, but about **ontological recoverability**. Decoherence renders information *operationally inaccessible*, but does not destroy it.

BCB, by contrast, encodes a different commitment: that **some distinctions are irreversibly discarded at the substrate level**, not merely hidden. This is the sense in which BCB enforces finite *effective* information capacity per causal update.

It is therefore more accurate to say:

BCB and amplitude-primitive ontologies make different ontological commitments about irreversibility, and therefore make different empirical predictions.

The present paper does **not** claim to refute Everettian quantum mechanics by argument alone. Instead, it claims that:

- Amplitude-primitive frameworks naturally predict *exact* erasure of microscopic substrate residue.
- Record-primitive (BCB) frameworks generically predict *suppressed but symmetry-protected* residue.

The 6Ω prediction is precisely the empirical discriminator between these commitments. We therefore soften the burden-of-proof language as follows:

Rather than asserting that amplitude-primitive frameworks are inconsistent, we assert that BCB-based frameworks provide a more natural account of irreversibility and yield distinct, testable predictions. Experiment, not philosophy, must ultimately decide between them.

I.3 Path-Independence, Isotropy, and Renormalisation Scale

Theorem 1 establishes that **exact path-independence at a given scale implies isotropy at that scale**. At the same time, the framework predicts residual anisotropy (the 6Ω signal). This is not a contradiction, but it requires careful scale-dependence clarification.

The correct interpretation is:

- **Path-independence is a scale-relative condition**, defined with respect to a fixed coarse-graining scale.
- At any finite coarse-graining scale ℓ , admissibility requires approximate isotropy of resolution cost at that scale.
- Under renormalisation, anisotropy coefficients flow toward zero in the infrared but are not strictly forbidden in the ultraviolet.

Thus:

- Path-independence holds **within the effective description at each scale**.
- The *parameters* governing cost-per-distance (e.g. effective propagation speed) become scale-dependent, converging toward isotropy in the IR.
- Residual anisotropy appears as irrelevant operators suppressed by RG flow, consistent with the predicted b^{-4} decay for hexagonal substrates.

The framework therefore predicts **asymptotic isotropy**, not exact isotropy at all scales. The 6Ω harmonic is the leading irrelevant operator consistent with C_6 symmetry.

I.4 Tick Invariance and Observer Frames

A tick is defined as a **primitive causal update step** of the substrate, not as a unit of coordinate time. The concern that different observers might count different numbers of ticks along the same worldline mirrors familiar issues in causal set theory and discrete causal models.

The resolution is:

- Tick count is invariant **along causal chains**, not across arbitrary foliations.
- Different observers may assign different coordinate times to the same sequence of ticks (time dilation), but this does not imply different numbers of substrate updates occurred.
- The tick count is closer to a *proper causal length* than to coordinate time.

Operationally, Δi_{\min} is defined **per substrate update**, not per observer-measured second. The framework therefore does not require observers to agree on tick counts directly, only that causal influence chains have invariant update counts. This is analogous to how causal order is invariant even though simultaneity is not.

I.5 The ϵ Floor and Practical Falsifiability

The ϵ parameter controls how strongly micro-geometry leaks into macroscopic propagation. The paper explicitly acknowledges that **null results alone do not falsify the framework without auxiliary assumptions**.

This is not unique to the present framework. Many effective field theories contain undetermined couplings whose smallness can only be bounded, not eliminated, by experiment.

The framework therefore distinguishes:

- **Structural falsification:** Detection of 4Ω without 6Ω (ϵ -independent, clean)
- **Amplitude falsification:** Null results below a physically motivated ϵ -floor (requires A_{ϵ})

Is there an experiment that distinguishes “ ϵ is small” from “the framework is wrong”? Only indirectly:

- Multi-wavelength measurements can determine the suppression order p , reducing parameter degeneracy.
- Independent constraints from atomic physics, X-ray anisotropy, or astrophysical direction-resolved tests can over-constrain (ϵ, p, a) .
- Detection of *any* symmetry-protected residue immediately falsifies amplitude-primitive models with exact Lorentz invariance.

Thus, while ϵ introduces conditionality, the framework remains empirically meaningful and increasingly constrained as sensitivity improves.

I.6 The Record Necessity Theorem: Conditional, Not Universal

The Record Necessity Theorem (§2.4.1) shows that **within an interventionist, counterfactual notion of causation**, stable records are required. It does *not* claim that all possible notions of causation must take this form.

Amplitude-primitive frameworks define causation differently: as correlation propagation under unitary dynamics rather than counterfactual constraint mediated by records. From within such frameworks, the theorem does not bite.

The correct logical status is therefore:

The theorem establishes internal consistency: if one adopts an interventionist, counterfactual notion of causation (the one used implicitly in experimental physics), then stable records and a nonzero commitment threshold are unavoidable.

The paper does not claim this forces fact-primitive ontology universally. It claims that **once one adopts the causal language already used in practice**, the threshold structure follows. This clarification aligns the theorem with Appendix D.1 and avoids over-statement.

I.7 Collective Modes, Vacuum Entanglement, and What “Propagates”

Finally, the collective-modes objection notes that quantum fields exhibit vacuum entanglement across arbitrarily large distances. If the vacuum is already globally correlated, what is “propagating” in the BCB picture?

The answer is: **not correlation, but constraint**.

- Vacuum entanglement represents *potential* correlations.
- Fact-commitment represents *actualised* constraints on future evolution.

BCB does not claim correlations must propagate from scratch. It claims that **the act of committing a definite outcome at a location requires causal resources to reach that location**. Pre-existing entanglement does not eliminate this requirement; it only shapes the statistical structure of outcomes once commitment occurs.

Thus, in the BCB picture:

- The vacuum may be globally entangled.
- But committing the fact “this detector clicked” at a distant location still requires local threshold crossing.
- The budget scaling problem concerns *actualisation*, not correlation.

This distinction mirrors standard quantum practice: entanglement enables correlations, but classical communication (limited by light speed) is still required to establish usable facts. BCB formalises this distinction at the substrate level.

I.8 Summary of Clarifications

- The non-admissibility theorem applies to fixed-geometry, finite-capacity runtimes; causal sets represent a different ontological class with different predictions.
- BCB does not refute Everettian quantum mechanics by logic; it proposes a competing ontology with distinct empirical consequences.
- Isotropy is scale-relative; residual anisotropy is compatible with RG flow and path-independence at fixed scales.
- Tick invariance is causal, not coordinate-based.
- ϵ introduces conditional falsifiability but does not trivialise the framework.
- The record necessity theorem is conditional on interventionist causation.
- Collective quantum correlations do not remove the need for causal resources to commit facts.

The framework’s central claim remains intact: **within a large and well-defined class of finite-capacity causal substrates, dimensionality and geometry are sharply constrained, and those constraints lead to a concrete experimental discriminator.**

Good idea. Section 9 already has "Anticipated Objections and Responses," but those are pitched at a general level. What you need is a more technical, referee-focused section that preempts the specific attacks a PRD reviewer would mount. Appendix I addresses the deep philosophical issues, but a referee might not read that far before forming their opinion.

Here's a draft of a new section (or a replacement/expansion of Section 9) that's more technically aggressive:

Appendix J: Technical Objections and Responses

This section addresses objections likely to arise from specialists in quantum gravity, quantum foundations, and precision tests of Lorentz invariance.

Objection 1: "The non-admissibility theorem excludes causal sets by construction. This is a selection effect, not a result."

Response: Correct, and this is by design. The theorem applies to substrates with well-defined coarse-grained wavefronts (Definition 9). Causal sets intentionally lack such structure—their Lorentz invariance is achieved statistically via random sprinkling, not geometrically.

This is not a defect but a *classification*. The theorem identifies a clean fork:

- **Fixed-geometry substrates** (within BCB scope): Dimensionality constrained; residual anisotropy predicted (6Ω)
- **Random-geometry substrates** (outside BCB scope): No such constraint; exact Lorentz invariance; no 6Ω

The frameworks make different empirical predictions. Detection of 6Ω disfavors random-geometry approaches; continued nulls favor them. We do not claim to refute causal sets—we claim to set up an experimental discriminator.

Objection 2: "Everettian QM doesn't require infinite information capacity. Finite regions have finite entropy."

Response: The distinction is between *entropy density* and *ontological recoverability*. Everettian QM preserves exact unitarity: no information is ever discarded, all branch structure remains recoverable in principle, and phase relations are preserved indefinitely. This is not infinite entropy but infinite *precision*—the universal wavefunction encodes all branches with all correlations at arbitrary accuracy forever.

BCB makes a different commitment: some distinctions are irreversibly discarded at the substrate level, not merely rendered operationally inaccessible. This is finite *effective* capacity per update.

The disagreement is ontological, not logical. We do not claim Everett is inconsistent—only that it predicts exact erasure of substrate residue while BCB predicts suppressed but detectable residue. The 6Ω harmonic is the empirical test. See Appendix I.2 for extended discussion.

Objection 3: "You require path-independence (implying isotropy) but predict anisotropy (6Ω). This is inconsistent."

Response: Path-independence is scale-relative. At any fixed coarse-graining scale ℓ , admissibility requires that resolution cost be approximately direction-independent at that scale. Under RG flow, anisotropy coefficients decay as b^{-4} for hexagonal substrates—they are irrelevant operators, suppressed but not forbidden.

The correct statement is: path-independence holds *within the effective description at each scale*. The parameters governing cost-per-distance become scale-dependent, converging toward isotropy in the IR. Residual anisotropy is the leading irrelevant operator consistent with C_6 symmetry.

This is standard EFT logic: exact symmetry in the IR limit, symmetry-breaking corrections suppressed by powers of (a/ℓ) . See Appendix I.3.

Objection 4: "How is the 'tick' defined across reference frames? Different observers measure different proper times."

Response: A tick is a primitive causal update step, not a coordinate time interval. Tick count is invariant *along causal chains*, not across arbitrary foliations. Different observers assign different coordinate times to the same tick sequence (time dilation), but the number of substrate updates is frame-independent in the same sense that causal order is invariant.

The threshold Δi_{\min} is defined per substrate update. Observers need not agree on tick rates measured in their coordinates; they must only agree that causal influence chains have well-defined update counts. This parallels how causal sets define temporal order via the causal partial order without pre-existing time.

For electromagnetic propagation in approximately flat spacetime (the regime of current experimental tests), tick count along any lightlike path is well-defined. Extension to curved spacetime requires coupling the interface to gravitational dynamics (Section 10.2).

Objection 5: "The parameter ε can always be made small enough to evade any null result. The framework is unfalsifiable."

Response: Amplitude predictions are conditional on auxiliary assumption A_ε (ε not exponentially suppressed). This is standard for EFTs with undetermined couplings—one can always postulate smaller couplings.

However, the *structural* prediction is ϵ -independent: the lowest allowed angular harmonic is $\cos(6\theta)$, not $\cos(4\theta)$. Detection of 4Ω without 6Ω falsifies the hexagonal selection unconditionally. Detection of 6Ω at any amplitude confirms the framework and disfavors exact-Lorentz alternatives.

Additionally:

- Multi-wavelength measurements extract p independently, reducing parameter degeneracy
- Independent channels (atomic physics, X-ray anisotropy, astrophysical polarimetry) over-constrain (ϵ, p, a)
- The heuristic $\epsilon \gtrsim 10^{-6}$ follows from generic substrate arguments (thermal/quantum fluctuations, defect densities)

The framework is falsifiable via the structural prediction; amplitude predictions become falsifiable under physically motivated assumptions. See Section 8.8.

Objection 6: "The Record Necessity Theorem assumes counterfactual causation, which amplitude-primitive frameworks reject. The argument is circular."

Response: The theorem is conditional, not universal. It shows that *if* one adopts interventionist, counterfactual causation—the notion implicit in experimental physics ("if we had set the apparatus differently, the outcome distribution would differ")—then stable records and a nonzero commitment threshold follow.

Amplitude-primitive frameworks define causation as correlation propagation under unitary dynamics, not counterfactual constraint. From within such frameworks, the theorem does not apply.

This is not a bug but a feature: the theorem delineates exactly where ontological commitments diverge. The 6Ω prediction then provides the empirical fork. We do not claim to refute amplitude-primitive ontologies by logic; we claim they make different predictions. See Appendix I.6.

Objection 7: "QFT modes are collective and delocalized. Your local budget constraint doesn't apply to field theory."

Response: Collective modes relocate the problem rather than solving it. A record instantiated across N correlated degrees of freedom requires N -way coordination. The budget for achieving such correlation must still propagate from the source.

More precisely: vacuum entanglement represents *potential* correlations; fact-commitment represents *actualized* constraints. BCB concerns the latter. Pre-existing entanglement shapes the statistical structure of outcomes but does not pre-commit which outcome occurs.

Consider EPR: joint outcome statistics are encoded in the entangled state, but *which* outcome occurs at each detector must still be committed locally. Entanglement enables correlations; it does not eliminate the need for local fact-commitment resources.

The 3D scaling problem reappears for collective records: correlation budget grows as r (ticks \times rate), while degrees of freedom requiring correlation grow as r^2 (wavefront) or r^3 (volume). The deficit persists. See Section 3.7 and Appendix I.7.

Objection 8: "Bulk reconstruction (Section 6) is speculative. You haven't shown it yields Euclidean 3D geometry."

Response: Correct. Section 6 is explicitly labeled as a "speculative outline" and "research direction," not a completed derivation. The core results—non-admissibility theorem, hexagonal selection, 6Ω prediction—do not depend on bulk reconstruction being correct.

A complete treatment would require: (i) explicit interface microdynamics, (ii) derived correlation functions, (iii) proof that induced geometry is approximately Euclidean, (iv) observer consistency. These remain open problems.

The paper's load-bearing claims are in Sections 2–4 and 8. Section 6 indicates how the framework *could* connect to emergent spacetime; failure of this specific mechanism would not affect the experimental prediction.

Objection 9: "The U(1) gauge group is assumed, not derived. This undermines the claim to explain electromagnetism."

Response: Correct. BCB and hexagonal selection constrain the *geometry* of the interface; they do not determine which fields propagate on that geometry. U(1) link variables are an input justified by phenomenological success, not first-principles derivation.

The framework explains:

- Why gauge redundancy is *natural* given link-based degrees of freedom
- Why Maxwell structure emerges in the continuum limit given U(1)
- Why two polarizations arise from 2D transverse modes

It does not explain:

- Why U(1) rather than SU(2) or other groups
- The origin of charge quantization
- Matter coupling and the full Standard Model

Gauge group selection is listed as an open problem (Section 10). The present paper addresses photon propagation geometry, not the complete theory of fundamental interactions.

Objection 10: "How does 6Ω map onto Standard Model Extension (SME) coefficients? What are the actual bounds?"

Response: The 6Ω signature corresponds to even-parity spatial anisotropy in the photon sector—specifically, the $\tilde{\kappa}_{\{e\}}$ coefficients in SME notation. These are independent of:

- $\tilde{\kappa}_{\{o+\}}$ (odd-parity, CPT-odd): constrained by astrophysical birefringence
- $\tilde{\kappa}_{\{tr\}}$: constrained by time-of-flight tests

Existing rotating cavity experiments (Herrmann et al. 2009, Zhang et al. 2021) achieve $\Delta c/c \sim 10^{-17}$ but focus on 2Ω systematics. A targeted 6Ω search in existing or future data would directly constrain the framework.

The $\cos(6\theta)$ signature integrates to zero over all directions, so isotropic tests (direction-averaged astrophysics, collider experiments) do not constrain it. The most relevant bounds come from direction-resolved cavity experiments.

For $p = 4$, $\varepsilon = 1$: the interface scale is bounded above ~ 10 keV ($a < 10^{-11}$ m). This is well below the Planck scale but above current direct probes. See Section 8.6 and Appendix C for detailed parameter space.

Objection 11: "Why should we believe the interface is hexagonal rather than some other geometry we haven't thought of?"

Response: Four independent selection channels converge on hexagonal coordination:

1. **Dispersion symmetry:** Hexagonal pushes anisotropy to 6th order; square only to 4th
2. **RG flow:** Anisotropy decays as b^{-4} (hex) vs b^{-2} (square)
3. **Information theory:** 60° angular gaps (hex) vs 90° (square)
4. **BCB admissibility:** Hexagonal minimizes directional variance

This convergence from independent arguments is the evidence. Among regular, translation-invariant planar tilings with finite-range local couplings, hexagonal is uniquely selected.

Irregular or aperiodic tilings (quasicrystals) could in principle satisfy some constraints but violate translation invariance, which BCB requires for strict path-independence. Whether any non-periodic tiling satisfies full admissibility remains an open problem. If such a structure exists and predicts (say) 10Ω , it would be a competing hypothesis distinguishable by harmonic order.

Objection 12: "This is just lattice QCD / lattice gauge theory rebranded. What's new?"

Response: Lattice gauge theory is a *computational tool* for QCD—a regularization scheme where the lattice is taken to zero spacing at the end. The continuum limit is the physics; the lattice is scaffolding.

This framework makes a different claim: the discrete interface is *physical*, not computational. The lattice spacing a is a real length scale, not a regulator to be removed. Consequently:

- Discreteness effects are suppressed but not eliminated
- Residual anisotropy is a physical prediction, not a lattice artifact
- The continuum limit is an approximation, not the exact theory

This is closer to causal set theory or loop quantum gravity (discrete structure is fundamental) than to lattice QCD (discrete structure is a tool). The experimental discriminator is whether symmetry-protected residue exists at all.