

Why Two Dimensions Are Not Emergent: The Tick–Bit Asymmetry and the Minimal Geometry of Distinguishability

Abstract (General Audience)

Why does space have the structure it does? This paper argues that different spatial dimensions serve fundamentally different purposes, and therefore cannot all emerge the same way from some deeper level of reality.

The key insight comes from distinguishing two types of change. A *tick* is a reversible change—like a pendulum swinging—where no information is permanently lost. A *bit* is an irreversible commitment—like writing in ink or making a measurement—where information is permanently discarded and a fact is created.

We show that these two types of change require different geometric structures. Reversible change can occur on a simple one-dimensional chain. But irreversible commitment requires at least two dimensions. The reason: in one dimension, there's no way to contain information loss locally—it spreads to affect everything. In two dimensions, loops and enclosed regions can trap and stabilize information loss, allowing facts to exist in one place without disrupting the whole system.

This means two-dimensional structure is special: it's the minimum geometry needed for facts to exist at all. The third dimension isn't required for making facts—it emerges as an index labeling *which scale* or *which effective description* facts are defined at. It doesn't store what was lost; it parameterizes where you are in the hierarchy of descriptions. These ideas help explain why information in black holes lives on surfaces, and why many theories treat one "dimension" as really being about scale rather than space.

Abstract (Technical)

We argue that spatial dimensions are not functionally equivalent and cannot emerge uniformly from a common substrate. By distinguishing reversible transitions (ticks) from irreversible commitments (bits), we demonstrate that different dimensions satisfy different informational requirements. Using a separator-based characterization of effective dimension, we prove that constant-separator substrates (those with bounded region boundaries) suffice for ticks but cannot support scalable localized bit formation. Two-dimensional substrates (boundaries scaling as $\sqrt{|R|}$) constitute the minimal geometry for irreversible distinguishability through boundary entropy

absorption. We prove that any emergent third dimension cannot encode recoverable information about discarded alternatives (No-Storage Constraint) and conjecture that it must satisfy boundary–scale consistency with area-law bounds. This framework provides structural foundations for holographic principles and offers falsifiable criteria for dimensional emergence.

Table of Contents

1. [Introduction](#)
 2. [Definitions and Conceptual Primitives](#)
 3. [Reversibility and the Nature of a Tick](#)
 4. [Constant-Separator Substrates: Ticks Without Bits](#)
 5. [Irreversibility and the Nature of a Bit](#)
 6. [Two-Dimensional Substrates as Minimal Bit Geometry](#)
 7. [The Tick–Bit Asymmetry Theorem](#)
 8. [Asymmetry and the Emergence of the Third Dimension](#)
 9. [Relation to Existing Frameworks](#)
 10. [Implications and Predictions](#)
 11. [Anticipated Objections](#)
 12. [Discussion](#)
 13. [Conclusion](#)
 14. [References](#)
-

1. Introduction

For General Readers

Imagine trying to write a permanent record using only a single line of dominos. You can make patterns flow back and forth, but any "message" you create can easily be erased by the same processes that created it. There's nowhere to "lock in" information—every part of the line is exposed to disturbance from both directions.

Now imagine a two-dimensional grid of tiles. Here you can create enclosed regions—loops that separate an interior from an exterior. A pattern inside a loop can be protected from what happens outside. Facts can be localized and preserved.

This simple geometric difference has profound implications. This paper argues that the distinction between one and two dimensions isn't just quantitative but qualitative: two dimensions are the minimum structure required for permanent facts to exist.

Technical Introduction

The emergence of spatial dimensions is often discussed as a uniform process, with all dimensions treated as equally derivative or equally fundamental. In this paper we argue that this assumption is incorrect. Spatial dimensionality is constrained not by symmetry alone, but by the informational roles different geometries can support.

We introduce a sharp distinction between reversible change (ticks) and irreversible distinction (bits). A tick is a reversible transition that preserves information; a bit is an irreversible commitment that discards alternatives and produces a fact. This distinction leads to a minimality result: constant-separator structures can support ticks but cannot support scalable localized bits. Two-dimensional structures are the minimal geometry capable of hosting localized, irreversible distinguishability.

We show that two dimensions are therefore not emergent in the same sense as the third. Instead, the third dimension arises from the existence of multiple levels of effective description at different scales. Depth parameterizes which description level is in use—it does not store discarded information but indexes where facts are defined.

The structure of the paper is as follows. Section 2 establishes definitions, including a separator-based characterization of effective dimension. Sections 3–4 characterize ticks and show that constant-separator substrates suffice for reversible dynamics. Sections 5–6 characterize bits and prove that two-dimensional geometry is minimal for irreversible distinguishability. Section 7 states the Tick–Bit Asymmetry Theorem with rigorous proof. Section 8 analyzes the emergent status of the third dimension. Sections 9–11 address connections to existing frameworks, implications, and objections. Section 12 concludes.

2. Definitions and Conceptual Primitives

For General Readers

Before diving into the argument, we need precise definitions. In everyday language, words like "information," "reversible," and "dimension" are used loosely. Here we give them exact meanings.

The most important distinction is between two kinds of change:

- **Tick:** A change that could be undone without losing anything. Like shuffling cards in a way where you could unshuffle them to get back exactly where you started.
- **Bit:** A change that permanently eliminates possibilities. Like shuffling cards and then burning half the deck. Something is decided; alternatives are gone forever.

The paper's central question is: what kind of geometric structure do you need to support each type of change?

Technical Definitions

Definition 2.1 (Tick). A *tick* is a reversible transition between states. Let S denote a state space and $f: S \rightarrow S$ an update rule. A tick corresponds to f being injective: for all $s_1, s_2 \in S$, $f(s_1) = f(s_2)$ implies $s_1 = s_2$. Ticks represent ordered change without irreversible commitment.

Definition 2.2 (Bit). A *bit* is an irreversible distinction between alternatives. A bit corresponds to a many-to-one mapping $g: S \rightarrow S'$ where $|S| > |S'|$. Multiple possible prior states collapse into a single outcome, discarding information and producing a fact.

Convention. We use "bit" for the physical act of irreversible commitment; "fact" for the semantic content carried by stable bits. A bit is the mechanism; a fact is what results when bits are stably encoded.

Definition 2.3 (Reversibility and Irreversibility). A dynamical system is *reversible* if its evolution operator is invertible. *Irreversibility* refers to information loss, typically associated with entropy production or coarse-graining.

Definition 2.4 (Physical Distinguishability). *Distinguishability* refers to stable, local, and persistent distinction maintained under system dynamics. A distinction is physically meaningful only if it can be stored without requiring global coordination.

Definition 2.5 (Locality and Adjacency). *Locality* denotes that interactions occur between neighboring elements. *Adjacency* refers to the graph-theoretic structure $A \subseteq V \times V$ defining neighbors. In quantum systems, locality is enforced by Lieb-Robinson bounds (Lieb & Robinson, 1972), which limit the speed of information propagation; our framework assumes analogous locality constraints in the discrete setting.

Definition 2.6 (Separator-Based Effective Dimension). For a locally finite adjacency graph $G = (V, E)$, let $R \subset V$ be a finite connected region and ∂R the vertex boundary (vertices in $V \setminus R$ adjacent to at least one vertex in R).

- The substrate is *constant-separator* (effectively 1D for localization) if there exists a constant c such that for all connected regions R , $|\partial R| \leq c$. No matter how large a connected region, its boundary never exceeds a fixed size.
- The substrate is *2D-like* (effectively two-dimensional) if $|\partial R|$ grows unboundedly with $|R|$ and, for lattice-like regions, satisfies the isoperimetric scaling $|\partial R| = \Omega(\sqrt{|R|})$.

Remark. The constant-separator class isolates the bottleneck phenomenon relevant to localization. Branching graphs (trees) may have $|\partial R|$ that grows with $|R|$, but still lack generic enclosure and fail bit-localizability for complementary reasons discussed in Section 11. In trees, the absence of cycles implies unique-path connectivity, so local perturbations cannot be rerouted; this prevents robust enclosure-based localization even when $|\partial R|$ grows. Standard references on separators include Chung (1997), Hoory, Linial & Wigderson (2006), and Lipton & Tarjan (1979) on planar separator theorems.

Definition 2.7 (Generic Local Cycle). A *generic local cycle* is a cycle in the adjacency graph such that (i) all vertices have bounded degree, (ii) cycle length is bounded independent of system size, and (iii) the cycle exists for generic vertex positions. In constant-separator substrates, cycles are either absent (paths) or global (single ring). In 2D-like substrates, local cycles of fixed size tile the substrate.

Definition 2.8 (Scalable Fact-Storage). A *scalable fact-storage scheme* on a substrate G is a construction such that for any $m \in \mathbb{N}$, the substrate can support m independent stable records in disjoint regions $\{R_i\}_{i=1}^m$ where:

- (i) Each region R_i has bounded size independent of m
- (ii) Each record predicate B_i satisfies the Bit Localizability Requirement (Section 5.3)
- (iii) The records are independent: the joint state (B_1, \dots, B_m) can take any of the 2^m possible values

Scalability captures the ability to store arbitrarily many facts without requiring unbounded region sizes or global coordination.

3. Reversibility and the Nature of a Tick

For General Readers

Think of a tick as the universe's "heartbeat"—a single step of change that doesn't destroy any information. If you knew the rules and current state perfectly, you could figure out exactly what the previous state was.

This is how physics at the deepest level seems to work. Quantum mechanics evolves by "unitary" operations that are perfectly reversible. Classical mechanics, tracked precisely, runs equally well forward or backward.

What's remarkable is how little structure this requires. A simple chain of linked elements—a one-dimensional line—is enough to support reversible change.

Technical Content

A tick represents the most elementary unit of change. Crucially, a tick is reversible. Without reversibility, there would be no principled way to distinguish genuine state transition from information loss.

Consider a system with state space S and dynamics $f: S \rightarrow S$. If f is injective, then $s \rightarrow f(s)$ constitutes a tick. The inverse f^{-1} exists on the image of f .

This aligns with foundational results:

- In classical mechanics, Hamiltonian evolution preserves phase-space volume (Liouville's theorem).
- In quantum mechanics, unitary evolution U satisfies $U^\dagger U = I$.
- In reversible computation, logical reversibility enables thermodynamic efficiency (Bennett, 1973).

Ticks encode order but not record. They establish sequence—before and after—but do not produce facts. Any apparent memory in a reversible system is encoded reversibly and therefore erasable without thermodynamic cost.

Because ticks require only ordered adjacency and reversible updates, they place minimal demands on spatial structure. A constant-separator substrate suffices.

4. Constant-Separator Substrates: Ticks Without Bits

For General Readers

A constant-separator substrate is like a string of beads: each element connects to at most a fixed number of neighbors. This is the simplest spatial structure beyond a single point.

Such chains can support reversible change beautifully. Waves travel along the chain, patterns shift left or right. But here's the crucial limitation: there's nowhere to hide.

If you want to permanently record something, you need to isolate it from disturbance. In a constant-separator structure, every region connects to the outside through only a few points—a bottleneck. Any attempt to make a permanent mark either can be undone by local processes, or disrupts the entire system globally.

Technical Content

4.1 Capabilities

In constant-separator substrates, reversible dynamics propagate information without loss. Consider a chain of N sites with states $\sigma_i \in \{0, 1\}$ and shift rule $\sigma_i(t+1) = \sigma_{i-1}(t)$. This is perfectly reversible.

4.2 Structural Limitations

Lemma 4.1. Let G be a constant-separator substrate with boundary constant c . For any connected region R :

- (i) $|\partial R| \leq c$
- (ii) Information transfer between R and exterior is bottlenecked at $O(c)$ channels

Example (Boundary Bottleneck). In constant-separator graphs, the exterior sees R through $O(1)$ sites regardless of $|R|$. In 2D grids, the exterior sees an $n \times n$ region through $O(n) = O(\sqrt{|R|})$ boundary sites. This difference determines whether the thermodynamic cost of irreversibility can remain localized.

Lemma 4.2 (Impossibility of Bit Localization). In a constant-separator substrate, any irreversible operation on R either (i) fails to produce a stable local record, or (ii) renders global dynamics non-invertible.

Argument. Let R have $|R| = n$ sites with alphabet Σ . The state space S_r has cardinality $|\Sigma|^n$. Let $g: S_r \rightarrow S'_r$ be irreversible with information loss $\Delta I(g)$.

The boundary ∂R has $|\partial R| \leq c$ sites, providing $O(1)$ entropy capacity.

For the irreversible operation to be localized, the thermodynamic cost (entropy production $\Delta S \geq kT \ln 2 \cdot \Delta I$) must be absorbed at the boundary without propagating into the exterior. But with $O(1)$ boundary capacity, this dissipation cannot remain confined as region size or information loss grows.

Therefore: either the entropy production propagates globally (rendering exterior dynamics non-invertible), or the record fails to stabilize against exterior perturbations. ■

4.3 The Scalability Problem

Even if a single operation g discards only 1 bit, the architecture cannot scale:

Lemma 4.3 (No Scalable Fact-Storage in Constant-Separator Substrates). Let $G = (V, E)$ be a constant-separator substrate with boundary constant c , and assume local dynamics in the sense of Definition 2.5. Then G does not support scalable fact-storage (Definition 2.8).

Proof. Suppose for contradiction that G supports a scalable fact-storage scheme with m disjoint record regions $\{R_i\}_{i=1}^m$ and stable record predicates $\{B_i\}_{i=1}^m$.

Step 1 (Boundary constraint). By the constant-separator property, $|\partial R_i| \leq c$ for all i . Under local dynamics, the only degrees of freedom through which the exterior couples to R_i are those in ∂R_i .

Step 2 (Independence implies boundary independence). For records to be independent (Definition 2.8(iii)), the value of B_i must be determinable from the state of R_i and its boundary ∂R_i alone, without reference to other regions. Under local dynamics, the boundary state of R_i can only depend on the history of R_i and its immediate neighbors—not on distant regions R_j .

Therefore, if records B_1, \dots, B_m are independent, their boundary configurations must be independently specifiable.

Step 3 (Capacity counting). Each boundary ∂R_i has at most c sites with alphabet Σ , hence at most $|\Sigma|^c$ possible boundary configurations. To support 2^m distinguishable joint record states (B_1, \dots, B_m) , we need the product of boundary configuration spaces to have cardinality at least 2^m :

$$\prod_{i=1}^m |\Sigma|^c = |\Sigma|^{\{mc\}} \geq 2^m$$

This inequality holds for any finite m , but the constraint tightens when we require *stability*: each boundary configuration must reliably encode its record value against perturbation from exterior dynamics.

Step 4 (Stability obstruction). Stability requires that the distinction $B_i = 0$ vs $B_i = 1$ persists under local dynamics on $V \setminus R_i$. But the boundaries $\{\partial R_i\}$ share a common exterior. Under local exterior dynamics, boundary sites of different regions can interact through paths in the exterior.

For the m records to remain independently stable, the exterior dynamics must preserve 2^m distinct equivalence classes of global states—one for each joint record value. But the exterior "sees" each region only through $O(1)$ boundary sites. As $m \rightarrow \infty$, maintaining 2^m stable equivalence classes with only $O(1)$ distinguishing information per class requires either:

- Nonlocal coordination: exterior dynamics that enforce global constraints across distant boundaries (violating locality), or
- External resources: a reservoir that absorbs the entropy of maintaining distinctions (violating the pure substrate assumption).

Why specially-constructed dynamics cannot help: Could one design exterior dynamics that preserve all 2^m equivalence classes without nonlocal coordination? No: such dynamics would need to "know" which global equivalence class the system is in to avoid mixing classes. But under locality, the dynamics at any exterior site can only depend on that site's neighborhood. To check global class membership from local information would require correlations propagating across all m boundaries—which is precisely nonlocal coordination. Any purely local dynamics will eventually mix some equivalence classes as m grows, destroying at least some records.

Therefore, constant-separator substrates cannot support scalable fact-storage. ■

Comment. The obstruction is not that a single 1-bit record is impossible, but that arbitrarily many independent stable records cannot coexist under purely local dynamics with $O(1)$ interface per record.

Error-correction perspective. Viewed through the lens of stability under local perturbations, scalable fact-storage is an error-correction requirement: maintaining m independent stable bits requires protection against local disturbances without global coordination. In coding-theoretic terms, this requires codes with macroscopic distance—the ability to detect and correct errors affecting any bounded subset of sites. Constant-separator substrates lack the necessary geometric redundancy because each record interfaces with the exterior through $O(1)$ channels. Without the enclosure and path redundancy provided by 2D structure, generic local dynamics will induce

mixing between record classes. The two conditions of Theorem 7.1(B)—boundary scaling and generic cycles—are precisely what enable error-resilient fact-storage.

5. Irreversibility and the Nature of a Bit

For General Readers

A bit is fundamentally different from a tick. Where a tick preserves all possibilities, a bit eliminates possibilities permanently.

Before you flip a coin, both heads and tails are possible. After you flip and look, one outcome becomes actual and the other is gone forever. That's a bit being formed—an irreversible commitment.

For facts to exist—for history to be real rather than just eternal reversible flow—irreversible commitments must occur. The deep question is: what geometry can support such commitments?

Technical Content

5.1 Formal Characterization

Let $g: S \rightarrow S'$ with $|g(S)| < |S|$. Define information loss:

$$\Delta I(g) = \log|S| - \log|g(S)|$$

For g to constitute a bit, $\Delta I(g) \geq \log 2$.

5.2 Physical Necessity

Irreversibility is unavoidable for records:

- **Measurement** requires irreversible apparatus-system correlation.
- **Memory** requires records independent of recorded system's evolution.
- **Landauer's principle:** erasing one bit dissipates $\geq kT \ln 2$ energy.

5.3 Localizability Requirement

Requirement (Bit Localizability). A substrate supports localized bits if there exists region R and operation g such that:

- (i) g is irreversible: $|g(S_r)| < |S_r|$
- (ii) g is local: acts only on R and immediate neighbors

- (iii) Stability: there exists record predicate $B(\sigma_r)$ invariant under subsequent local dynamics on $V \setminus R$
- (iv) Global consistency: exterior evolution remains well-defined

6. Two-Dimensional Substrates as Minimal Bit Geometry

For General Readers

What changes in two dimensions? The key is *loops* and *enclosed regions*.

In two dimensions, you can draw a circle separating inside from outside. You can now *isolate* a region. The consequences of irreversible change inside don't have to spread everywhere—they can be contained.

Moreover, in two dimensions, the boundary of a region grows with its size. A bigger enclosed area has a longer perimeter. This growing boundary provides sufficient capacity to absorb the thermodynamic cost of irreversibility (the entropy that must be produced when information is genuinely discarded).

This is why two dimensions are minimum for stable facts: loops create isolation; growing boundaries provide dissipation capacity.

Technical Content

6.1 Structural Features

Consider \mathbb{Z}^2 with nearest-neighbor adjacency. The two necessary conditions (Theorem 7.1B) are satisfied:

Condition (i) — Boundary scaling:

- **Isoperimetric scaling:** $|\partial R| = \Omega(\sqrt{|R|})$ for connected R . An $n \times n$ region has boundary $4n = \Theta(\sqrt{\{n^2\}})$.
- This provides boundary capacity that grows with region size, avoiding the bottleneck obstruction.

Condition (ii) — Generic local cycles:

- **Local cycles:** Every vertex participates in bounded-length cycles (e.g., 4-cycles/plaquettes on the square lattice).
- **Enclosure:** Cycles separate interior from exterior (Jordan curve theorem in the planar embedding). Perturbations outside cannot reach inside without crossing the boundary.
- **Path redundancy:** Multiple independent paths exist between nearby points. Local damage can be "routed around."

6.2 Mechanism of Bit Localization

Proposition 6.1 (Bit Localization in 2D). In a 2D-like substrate, for any local irreversible map g on a region R with information loss $\Delta I(g) \leq |\partial R| \cdot \log|\Sigma|$, there exists a choice of local reversible exterior dynamics such that g satisfies the Bit Localizability Requirement.

Quantifier note. This is an existence result: it asserts that for any local irreversible map with bounded information loss, there exists a choice of local environment coupling and exterior dynamics under which irreversibility can be confined. It does not claim that such confinement occurs for arbitrary exterior dynamics—only that 2D geometry makes confinement *possible*, whereas constant-separator geometry makes it *impossible* (Lemma 4.3).

Proof. Let G be a 2D lattice with alphabet Σ . Let R be an $n \times n$ block, so $|R| = n^2$ and $|\partial R| = 4n$. Let $g: \Sigma^R \rightarrow \Sigma^R$ be a local irreversible map with information loss $\Delta I(g)$.

Step 1 (Partition by output). The map g partitions the input space Σ^R into equivalence classes, where $\sigma \sim \sigma'$ iff $g(\sigma) = g(\sigma')$. Let $\{C_a\}$ be these classes, indexed by output $a = g(\sigma)$ for $\sigma \in C_a$. The number of classes is $|g(\Sigma^R)|$, and

$$\Delta I(g) = \log|\Sigma^R| - \log|g(\Sigma^R)| = \log(|\Sigma|^{n^2} / |g(\Sigma^R)|)$$

Step 2 (Boundary dissipation without microstate recovery). Localization does *not* mean that discarded distinctions are stored at the boundary. If the boundary encoded which element of the preimage fiber C_a was the input, the combined map (interior output, boundary state) would be injective—contradicting Definition 2.2 and eliminating genuine information loss.

Instead, localization means that the thermodynamic cost of the many-to-one map—entropy production required by Landauer's principle—is confined to degrees of freedom at or near ∂R , while the exterior dynamics remain well-defined and unaffected.

Concretely: when g collapses multiple interior microstates to a single output a , this irreversible operation produces entropy $\Delta S \geq kT \ln 2 \cdot \Delta I(g)$ (Landauer, 1961). In 2D-like substrates, this entropy can be deposited into boundary/environment degrees of freedom that are driven into high-entropy states *whose distribution is independent of which element of the preimage fiber occurred*. The boundary absorbs the thermodynamic cost without encoding recoverable information about the lost microstate.

The exterior can retain a stable record of the output macrostate $a = g(\sigma)$ —which fact was created—without encoding the discarded microstate identity. This is the sense in which 2D boundaries "absorb" irreversibility: they provide sufficient capacity for entropy production to remain localized, not for lost information to be secretly preserved.

Capacity requirement. For this dissipation to remain localized, the boundary must have sufficient entropy capacity. An $n \times n$ region has boundary $|\partial R| = 4n$ sites. The entropy capacity of the boundary scales as $4n \cdot \log|\Sigma|$. Specifically, entropy production $\Delta S \leq kT \ln 2 \cdot \Delta I(g)$ must be absorbed by boundary degrees of freedom with entropy capacity approximately $|\partial R| \cdot kT \ln|\Sigma|$,

which is satisfied when $\Delta I(g) \leq |\partial R| \cdot \log|\Sigma|$. For irreversible operations with bounded per-site information loss ($\Delta I = O(1)$ per operation), this capacity condition holds for any region size. For extensive loss ($\Delta I = \Theta(n^2)$), larger boundaries or staged coarse-graining is required.

Step 3 (Exterior dynamics). Define exterior dynamics that:

- Act reversibly on $V \setminus R$
- Couple to the boundary in a way that allows entropy absorption without correlating with preimage identity
- Preserve the macrostate record (which output a occurred) while remaining agnostic to microstate history

Step 4 (Record predicate). Define $B = 1$ iff $g(\sigma_r)$ belongs to a designated subset of outputs (e.g., the "reset" states). This record predicate depends only on the *output* macrostate, not on the lost microstate information.

Why boundary capacity matters for B. The record predicate B does not read the boundary state; the boundary's role is to absorb the entropy cost of irreversibility, enabling exterior dynamics to remain well-defined and the macrostate distinction to persist. Without sufficient boundary capacity, the thermodynamic disturbance would propagate into the exterior, corrupting global dynamics and destabilizing the record.

Under the constructed dynamics:

- The macrostate record B is determined by $g(\sigma_r)$ alone
- Boundary entropy absorption is localized
- Exterior dynamics remain well-defined and reversible
- Therefore B is invariant under subsequent local dynamics on $V \setminus R$

Step 5 (Global consistency). The global evolution is genuinely many-to-one: multiple initial states (differing in R) map to the same final state. The discarded information is truly lost—not relocated to the boundary. The boundary's role is to absorb the thermodynamic cost of this loss, keeping it localized rather than propagating globally.

All four conditions of Bit Localizability are satisfied. ■

Remark. The critical distinction is between (a) storing which microstate was lost (which would make the map invertible) and (b) absorbing the entropy cost of loss (which keeps the map genuinely irreversible but localized). 2D boundaries enable (b) without requiring (a).

6.3 Concrete Example: Bit Formation and Failure

Failure in 1D: Consider a 1D chain of 10 sites. Let $R = \{\text{sites } 4, 5, 6\}$. Define g : reset to all-0 if $\text{sum} \geq 2$, else preserve. The boundary $\partial R = \{\text{sites } 3, 7\}$ has only 2 sites.

The problem is not that "lost information has nowhere to go"—that framing would incorrectly suggest information is being stored somewhere. Rather, the entropy production from the irreversible operation cannot be confined to a local region. With only $O(1)$ boundary sites, the thermodynamic disturbance propagates into the exterior, affecting global dynamics. In a constant-separator substrate, any attempt at localized irreversibility either (a) corrupts exterior dynamics, or (b) requires nonlocal coordination to maintain consistency—violating locality. Extending R to more sites increases the entropy cost while boundary capacity remains fixed at $O(1)$.

Success in 2D: Consider a 2D square lattice. Let R be a 3×3 block of 9 sites with binary states. Define:

$$g(\sigma_r) = (0,0,0,0,0,0,0,0,0) \text{ if } \sum_{i \in R} \sigma_i \geq 5, \text{ else } g(\sigma_r) = \sigma_r$$

This map sends all configurations with $\text{sum} \geq 5$ to a single reset state, leaving others unchanged. The image has size $1 + N_{\{<5\}}$, where $N_{\{<5\}}$ is the number of configurations with $\text{sum} < 5$. Information loss: $\Delta I = \log 512 - \log(1 + N_{\{<5\}})$.

The boundary ∂R has 12 sites (the perimeter of a 3×3 block). This provides sufficient entropy capacity for the thermodynamic cost of the irreversible operation to remain localized—the boundary absorbs the entropy production without encoding which of the $2^{\{n^2\}} - N_{\{<5\}} - 1$ collapsed microstates was the input.

Define record: $B = 1$ iff R is in the reset state (all zeros). This predicate depends only on the *output macrostate*, not on which input microstate led to the reset. The distinction "reset occurred" vs "reset did not occur" is the fact being recorded.

Under local exterior dynamics that respect the boundary as a dissipative interface, B is stable: the interior macrostate is unchanged by exterior evolution (locality), and the boundary's role is to absorb entropy, not to encode recoverable information. The fact "R underwent reset" is genuinely recorded and localized.

Proposition 6.2 (Minimality of 2D). Two-dimensional-like geometry is minimal for scalable fact-storage, being the lowest dimension satisfying both boundary scaling and enclosure.

Proof. Constant-separator substrates fail boundary scaling \rightarrow no scalable fact-storage (Lemma 4.3). Trees fail enclosure \rightarrow no robust isolation (Section 11.3). 2D-like substrates satisfy both \rightarrow scalable fact-storage exists (Proposition 6.1). See Theorem 7.1 for the complete two-condition statement. ■

7. The Tick–Bit Asymmetry Theorem

For General Readers

This section states the paper's main result:

1. **For reversible change (ticks):** A simple chain is enough.
2. **For irreversible commitment (bits):** You need at least two dimensions, and for two distinct reasons:
 - **Boundary scaling:** In one dimension, the boundary between a region and outside stays small no matter how big the region. You can't fit enough bookkeeping through the bottleneck.
 - **Enclosure:** You need loops to create "inside" and "outside." In trees (which have growing boundaries but no loops), there's no way to protect a record from perturbations—every path is unique, so disturbances can't be routed around.

Two dimensions is the minimum geometry satisfying both requirements. This is why facts—permanent records, definite outcomes—require at least two-dimensional structure.

Technical Statement and Proof

Theorem 7.1 (Tick–Bit Asymmetry). Let $G = (V, E)$ be a substrate with local dynamics. Then:

(A) Sufficiency of constant-separator for ticks. If the substrate need only support reversible transitions (ticks), then constant-separator adjacency is sufficient.

(B) Necessity of 2D-like for bits. If the substrate must support scalable fact-storage (Definition 2.8), then the substrate must be 2D-like, requiring both:

- (i) **Boundary scaling:** $|\partial R| = \Omega(\sqrt{|R|})$ for connected regions R , and
- (ii) **Generic local cycles:** enabling enclosure and path redundancy.

Condition (i) is necessary to avoid the boundary bottleneck (Lemma 4.3). Condition (ii) is necessary to enable robust enclosure-based localization (see proof and Section 11.3 on trees).

Proof of (A). Construct a path or cycle graph G . Define local reversible dynamics (e.g., shift map). The system supports ticks, and $|\partial R| \leq 2$ for all connected R . ■

Proof of (B). We prove necessity of each condition.

Necessity of (i): Suppose G has bounded separators: $|\partial R| \leq c$ for all connected R . By Lemma 4.3, G cannot support scalable fact-storage. Contrapositive: scalable fact-storage requires $|\partial R| \rightarrow \infty$ as $|R| \rightarrow \infty$.

More precisely, for m independent records in regions of bounded size k , the total boundary must support 2^m stable distinctions. If each $|\partial R_i| \leq c$, Step 4 of Lemma 4.3 shows stability fails. For stability, we need boundary capacity that grows with the number of records. In 2D-like substrates, $|\partial R| = \Omega(\sqrt{|R|})$, so an $n \times n$ region supports $\Theta(n)$ boundary capacity, enabling $\Theta(n)$ independent $O(1)$ -bit records (Proposition 6.1).

Necessity of (ii): Suppose G has $|\partial R| = \Omega(\sqrt{|R|})$ but lacks generic local cycles (e.g., trees). Then:

- Every path between points is unique
- Local perturbations cannot be rerouted around damaged paths
- Records lack enclosure: there is no "inside" protected from "outside"

In such substrates, a record at region R can be corrupted by perturbations along the unique paths connecting R to the rest of the system. No rerouting is possible. Even with growing boundaries, records are not robustly isolated. A formal argument showing that trees cannot support stable records under arbitrary local dynamics—specifically, that unique-path connectivity allows perturbations to reach any record unavoidably—is given in Section 11, Objection 2.

Sufficiency of 2D-like: When both (i) and (ii) hold, Proposition 6.1 shows bit localization succeeds, and scalable fact-storage follows by tiling the substrate with bounded record regions, each with adequate boundary capacity. ■

Corollary 7.2. Two-dimensional geometry is not emergent in the same sense as higher dimensions. It is the minimal substrate for facts, being the lowest-dimensional geometry satisfying both boundary scaling and enclosure conditions.

Corollary 7.3. Apparent emergence of 2D from lower-dimensional substrates must import structure that supplies the missing condition:

- For constant-separator substrates (missing (i)): nonlocal interactions or hidden degrees of freedom to increase effective boundary capacity
- For trees (missing (ii)): additional edges creating cycles and enclosure
- For either: external reservoirs to offload entropy

Corollary 7.4. Dimensions $d \geq 3$ are not required for distinguishability; they satisfy both conditions but with excess capacity. Additional dimensions serve other purposes (Section 8).

8. Asymmetry and the Emergence of the Third Dimension

Sections 3–7 establish that two-dimensional structure is necessary and sufficient for localized irreversible commitments. Any additional dimension must therefore serve a role orthogonal to distinguishability. In this section we derive constraints on what such an additional dimension can and cannot represent.

For General Readers

If two dimensions suffice for facts, why does our world have three spatial dimensions?

The third dimension serves a different purpose: it indexes *which level of description* you're working at. When you zoom out from atoms to molecules to cells to organisms, you're moving

through levels of effective description—each level has its own facts, defined at that scale. The third dimension emerges as a parameter labeling these levels.

Crucially, depth does *not* store what was lost during coarse-graining. The information discarded when moving from atoms to molecules is genuinely gone—that's what makes molecular-level facts possible. Depth simply tells you which effective description you're using, not what you integrated out to get there.

Think of it like zoom level on a map. The zoom parameter tells you what scale you're viewing, but it doesn't contain the details that disappear when you zoom out. Those details are simply not represented at coarser scales.

Technical Content

8.1 The No-Storage Constraint

We begin with a structural result that constrains what any emergent dimension beyond two can represent.

Proposition 8.1 (No-Storage Constraint on Depth). Any emergent dimension beyond the minimal two required for localized bits cannot encode recoverable information about discarded alternatives without violating irreversibility. Therefore, any additional dimension must parameterize which effective description is in use, not what information was lost to reach it.

Proof. Suppose an emergent dimension $d > 2$ encoded recoverable information about alternatives discarded during bit formation. Then the transition from "pre-commitment" to "post-commitment" states would be invertible: given the post-commitment state plus the information stored in the extra dimension, one could reconstruct the pre-commitment state. But this contradicts the definition of a bit (Definition 2.2), which requires genuine information loss—a many-to-one mapping with no recovery possible.

Therefore, if irreversibility is to be maintained, any additional dimension cannot store discarded information. It can only label *which* effective description is in use (parameterizing the coarse-graining level) without retaining *what* was coarse-grained away. ■

This proposition elevates the interpretation of depth from a heuristic to a necessity. The RG analogy is not merely suggestive—it is structurally forced.

8.2 Scale Index: Formal Definition

Definition 8.1 (Scale Index). Let $\{\pi_k: S_k \rightarrow S_{k+1}\}$ be a sequence of coarse-graining maps between state spaces, where each π_k is surjective and many-to-one (irreversible). The index $k \in \mathbb{N}$ labels the effective description level at which distinctions are defined. The set $\{S_k\}$ forms a partially ordered hierarchy under refinement (S_k refines S_{k+1}), and k is referred to as the *scale index*.

Key properties of the scale index:

- **Non-invertibility:** Given k and a state in S_k , one cannot recover the pre-image in S_{k-1} .
- **Non-retention:** k does not encode which element of the pre-image fiber was the actual prior state.
- **Description-labeling:** k specifies which effective theory applies, not which microstate obtains.

The scale index k is the formal object corresponding to "depth" in geometric language.

8.3 Coarse-Graining and Scale

Consider a sequence of coarse-graining operations $\pi_1, \pi_2, \dots, \pi_n$ relating descriptions at different scales:

$$S_0 \xrightarrow{\pi_1} S_1 \xrightarrow{\pi_2} S_2 \rightarrow \dots \xrightarrow{\pi_n} S_n$$

where $|S_{k+1}| < |S_k|$. Each S_k is an effective description at scale k . The maps π_k are irreversible—information is genuinely lost in each transition, enabling facts to be defined at scale $k+1$.

Geometrically: each S_k is a 2D slice where facts at scale k are localized; the scale index k parameterizes which slice. The "third dimension" in this picture is the scale index direction.

8.4 Boundary–Scale Consistency

The connection between 2D fact-localization and scale hierarchy is not merely analogical—it should obey a scaling constraint. We state this as a conjecture with supporting argument.

Conjecture 8.2 (Boundary–Scale Consistency). In a 2D-like substrate supporting scalable localized bits, any hierarchy of effective descriptions indexed by scale k must satisfy: the number of independent stable facts representable at level k scales at most with the boundary measure of regions at that level.

Here "independent stable facts" means a set of record predicates $\{B_i\}$ supported on disjoint bounded regions at level k whose joint value realizes 2^m distinguishable macrostates, exactly as in Definition 2.8 (Scalable Fact-Storage). This ties the conjecture directly to the operational notion of fact-storage developed in Sections 4–6.

Supporting argument. At each scale k , facts are localized on 2D slices S_k . By Proposition 6.1, the capacity for independent stable bits in a region $R \subset S_k$ is bounded by the boundary's entropy-absorption capacity, which scales as $|\partial R| \cdot \log|\Sigma|$. Coarse-graining to scale $k+1$ cannot create new facts exceeding the boundary capacity of regions in S_{k+1} . Since coarse-graining reduces resolution ($|S_{k+1}| < |S_k|$), boundary measures at coarser scales are typically smaller.

The gap between this argument and a full proof lies in making precise how boundary capacity "inherits" through coarse-graining. A complete proof would require specifying the class of admissible coarse-graining maps and showing that fact-count inheritance respects boundary scaling uniformly. We leave this formalization for future work, noting that the conjecture is

consistent with all known instances (Bekenstein bounds, Ryu-Takayanagi formula, tensor network area laws).

This conjecture connects:

- 2D boundary scaling (Theorem 7.1)
- RG scale hierarchies
- Area laws for entropy

The Bekenstein bound and holographic entropy formulas (Ryu & Takayanagi, 2006) are instances of this general pattern.

8.5 Why Depth Behaves Geometrically

Physical constraints on coarse-graining induce geometric regularity in the scale index:

1. **Locality preservation:** Coarse-graining typically maps local structures to local structures, inducing continuity between adjacent description levels.
2. **Monotonicity:** Resolution decreases monotonically with coarse-graining, providing natural ordering.
3. **Dimensional consistency:** If each description level S_k is two-dimensional and interlevel maps π_k are local, the combined structure admits a three-dimensional embedding with consistent adjacency.

The metric properties of depth inherit from the regularity of physical coarse-graining. A complete derivation of emergent metric structure from coarse-graining constraints remains open but is constrained by Proposition 8.1 and Conjecture 8.2.

8.6 Connections to Established Frameworks

- **Renormalization group:** The flow parameter μ is precisely the scale index. It labels which effective theory is in use. It does not encode what degrees of freedom were integrated out—those are genuinely gone (Proposition 8.1).
- **Tensor networks (MERA):** The radial direction is the scale index, parameterizing entanglement renormalization level. The "holographic" reconstruction from boundary to bulk is exactly the refinement hierarchy $\{S_k\}$. The connection between tensor networks and holography is developed in Swingle (2012).
- **AdS/CFT:** The bulk radial coordinate z corresponds to RG scale in the boundary theory. Conjecture 8.2 suggests why bulk reconstruction respects area laws. The Ryu-Takayanagi formula (Ryu & Takayanagi, 2006) for holographic entanglement entropy is a specific instance of the boundary–scale pattern.

8.7 Dimensional Reduction Under Reversibility

Prediction: when irreversibility is suppressed, the distinction between description levels collapses. Systems approaching perfect reversibility have no coarse-graining hierarchy—all scales are equivalent. Effective dimensionality reduces toward the minimal 2D substrate.

This aligns with holographic behavior in highly coherent quantum systems, where bulk degrees of freedom become redundant descriptions of boundary physics.

8.8 Falsifiability

The claims in this section would be falsified if:

- An emergent dimension were shown to retain invertible (recoverable) information about discarded degrees of freedom, contradicting Proposition 8.1.
- A physical system exhibited stable fact-storage at some scale exceeding the boundary capacity at that scale, contradicting Conjecture 8.2.
- Coarse-graining produced geometric structure (metric, smooth embedding) without locality preservation, contradicting Section 8.5's explanation.
- A third spatial dimension were shown necessary for bit localization itself, contradicting the minimality of 2D (Theorem 7.1).

8.9 Independence of Core Results

Importantly, none of the results in Sections 3–7 depend on the interpretation offered here. The Tick–Bit Asymmetry Theorem (7.1) and the minimality of two-dimensional fact-localization stand independently of how one models the emergence of additional dimensions. Even if the scale-index interpretation of depth were replaced by an alternative account, the structural necessity of 2D for localized bits would remain.

9. Relation to Existing Frameworks

For General Readers

This framework explains patterns across physics:

Holographic Principle: Information in a region scales with surface area, not volume. Our framework explains: bits require 2D for localization, so information lives on surfaces.

Renormalization: Ignoring microscopic details while preserving large-scale predictions involves "flow" through scales resembling an extra dimension. That dimension labels which effective description you're using—not what was thrown away.

Tensor Networks: Mathematical structures with stacked layers where the "radial" direction means scale. Our framework: that direction parameterizes the level of description, indexing where facts are defined.

Technical Content

9.1 Holographic Principles

Bulk physics encoded on boundaries. Present framework adds necessity: facts *must* reside on minimal bit-supporting geometries. Area-law entropy (Bekenstein, 1973) reflects 2D localization.

9.2 Renormalization Group

Scale parameter μ as emergent depth—labeling which effective theory is in use, not encoding integrated-out modes (Wilson, 1971).

9.3 Tensor Networks

MERA separates local entanglement from layered depth (Vidal, 2007). Radial direction = scale parameter indexing description level.

9.4 Measurement and Decoherence

Irreversibility must be localizable. 2D provides minimal setting for stable local records from decoherence.

10. Implications and Predictions

For General Readers

1. **Information lives on surfaces.** Holographic behavior isn't special to gravity—it's geometric necessity.
2. **Time and space differ fundamentally.** Time (tick sequence) is 1D. Space beyond 1D supports commitment, not just evolution.
3. **Coherent systems lose a dimension.** Approaching reversibility, the third dimension becomes irrelevant.
4. **1D memory requires cheating.** Real 1D storage uses 3D structure, external reservoirs, or nonlocal control.

Technical Content

Implication 1: Surface encoding explains area-scaling laws structurally.

Implication 2: Temporal and spatial degrees of freedom should not be treated symmetrically.

Note: Full treatment of tick-ordering and thermodynamic arrows is outside scope; this paper isolates minimal structural requirements for localized records.

Implication 3: Dimensional reduction under reversibility—coherent quantum systems, error-corrected systems, near-critical systems.

Implication 4: Robust 1D storage requires importations violating pure locality (Corollary 7.3).

Implication 5: Layered structures diagnose a hierarchy of effective descriptions, not a ledger of discarded information.

11. Anticipated Objections

This section addresses the most substantive challenges likely to be raised against the framework. We treat each objection seriously and explain either why it does not apply or how the framework accommodates it.

Objection 1: "You defined dimension to make the theorem true."

The Challenge. Any structural necessity claim invites the suspicion that definitions were chosen to guarantee the desired result. A skeptic might argue: "You defined 'constant-separator' and '2D-like' precisely so that your theorem would hold. With different definitions, the result might fail."

Response. The separator-based definition is not chosen for convenience but is *forced* by the question we are asking: "Can irreversibility be localized?"

The capacity of a region's boundary to absorb the thermodynamic cost of irreversibility (entropy production) is the quantity that determines localization. If boundary capacity is bounded while region size grows, the entropy cost of irreversible operations cannot remain localized and will corrupt global dynamics (Lemma 4.3). If boundary capacity grows with region size, localization becomes possible (Proposition 6.1). This is not a matter of geometric aesthetics but of thermodynamic necessity.

Other notions of dimension—Hausdorff dimension, embedding dimension, topological dimension in the covering sense—are irrelevant to the localization question because they do not track boundary capacity. A fractal dust can have fractional Hausdorff dimension but no notion of "boundary" relevant to information flow. A space can embed in \mathbb{R}^3 while having constant-separator structure.

The key point: We define effective dimension operationally by boundary scaling because this is the quantity that constrains localized irreversibility. Any definition that ignores boundary growth

cannot adjudicate the question posed. The definition is justified by the capacity arguments in Lemmas 4.2–4.3 and Proposition 6.1, not by geometric convention.

Objection 2: "Trees and branching graphs have growing boundaries—aren't they counterexamples?"

The Challenge. Trees can have $|\partial R| = \Omega(|R|)$ for certain regions—even faster growth than 2D lattices. If boundary scaling is the key, shouldn't trees support bit localization?

Response. This objection correctly identifies that boundary scaling alone is insufficient. That is precisely why Theorem 7.1(B) requires *two* conditions:

- (i) Boundary scaling: $|\partial R| = \Omega(\sqrt{|R|})$
- (ii) Generic local cycles enabling enclosure

Trees satisfy (i) in some cases but categorically fail (ii). The consequences are severe:

Unique-path vulnerability. In a tree, exactly one path connects any two vertices. A perturbation propagating toward a record region R follows this unique path and cannot be rerouted. There is no alternative.

Formal argument: Let R be a record region in a tree T , and let $v \in \partial R$ be any boundary vertex. Let w be any vertex in $T \setminus R$. Since T is a tree, there exists exactly one path P from w to v . Under local dynamics, a perturbation at w can propagate along P , reaching v in $|P|$ steps. Once at v , the perturbation can enter R . Since this holds for *any* choice of w and v , there is no way to shield R from perturbations originating anywhere in the tree. The growing boundary provides many entry points but no protection mechanism—perturbations can reach *any* entry point unavoidably.

No enclosure. Cycles create topological separation: interior vs. exterior. In 2D, a perturbation outside a cycle C cannot reach inside without crossing C . Trees have no cycles, so every region is "exposed" to the entire graph. There is no protected interior.

Failure of robust stability. Even with large boundaries, a record in a tree can be corrupted by a perturbation entering through *any* boundary vertex and propagating inward. The boundary provides capacity but not protection.

Summary: Boundary growth alone is insufficient; enclosure via cycles is required to prevent perturbations from propagating unavoidably. Trees fail localization for this complementary reason, not because of boundary capacity. This is why 2D is minimal: it is the lowest dimension satisfying *both* conditions.

Objection 3: "Topologically protected states exist in 1D—don't Majorana chains and symmetry-protected phases contradict your result?"

The Challenge. Condensed matter physics has identified robust 1D systems: Majorana zero modes at chain endpoints, symmetry-protected topological phases, and related constructions. These appear to store information stably in 1D structures.

Response. We acknowledge the sophistication of these systems, but they do not constitute counterexamples because they violate the locality assumptions explicitly stated in our framework.

Global constraints. Topological protection in 1D relies on constraints that span the entire system: fermion parity conservation, boundary conditions at both ends, or symmetry requirements that must be enforced globally.

Symmetry protection. Symmetry-protected phases are robust only against perturbations respecting the protecting symmetry. A generic local perturbation (as in our Definition 2.5) can break the symmetry and destroy the protected state.

Nonlocal observables. The "information" stored in these systems is often encoded in nonlocal observables—string operators, endpoint correlations, or topological invariants—rather than in local, independently readable records.

Scope clarification. Our framework concerns local fact formation: the ability to create and maintain localized records under generic local dynamics without global coordination. Topologically protected 1D states achieve stability through global structure, which is precisely the "import" identified in Corollary 7.3(i). They do not demonstrate local bit formation in the sense of Definition 2.8.

Objection 4: "Quantum mechanics is fundamentally reversible—doesn't unitary evolution undercut the entire notion of bits?"

The Challenge. If the universe evolves unitarily at the fundamental level, there is no genuine irreversibility. All "information loss" is merely apparent, arising from coarse-graining or entanglement with the environment. How can bits exist if microdynamics are reversible?

Response. This objection is based on a misunderstanding of the framework's scope.

We agree: Unitary quantum evolution constitutes ticks. The Schrödinger equation, applied to a closed system, is perfectly reversible. No information is lost at the fundamental level.

The question is different: Given that measurement, decoherence, and record formation *do* occur—wherever the irreversibility comes from—what geometric structure is required for these processes to be localized?

Our framework is agnostic about the *origin* of irreversibility. Whether irreversibility is:

- Fundamental (objective collapse theories)
- Emergent from coarse-graining (statistical mechanics)
- Perspectival (decoherence relative to an observer)

...the geometric constraints on *localization* remain the same. A measurement outcome must be recorded somewhere. That record must be stable against subsequent local dynamics. This requires 2D-like geometry regardless of one's interpretation of quantum mechanics.

Framing: The framework complements, rather than challenges, unitary quantum mechanics by explaining the geometric constraints on record formation. It addresses where irreversibility can *live*, not whether microdynamics are reversible.

Objection 5: "Isn't this just holography in disguise? What's new here?"

The Challenge. Holographic principles (AdS/CFT, Bekenstein bounds) already establish that information scales with area, not volume. Is this paper merely restating known results in different language?

Response. The relationship to holography is substantive but not one of derivation or redundancy.

Holography states that boundary encoding works. AdS/CFT demonstrates that bulk physics can be reconstructed from boundary data. The Bekenstein bound shows that entropy scales with area. These are powerful results, but they are typically presented as surprising properties of specific theories (gravity, black holes) or as conjectures requiring verification case by case.

We explain why it must work. The tick-bit asymmetry provides a *structural* reason for holographic behavior: localized irreversibility requires 2D geometry, so information (= stable records = localized bits) inherently resides on surfaces. This is not a property of gravity or specific field theories; it follows from the geometry of distinguishability itself.

The framework applies beyond gravity. Holography in its original sense concerns gravitational systems. Our result applies to any substrate with local dynamics and fixed adjacency. It explains why tensor networks, RG flows, and non-gravitational systems exhibit "holographic" features.

Strong but fair summary: Holography is typically presented as a surprising property of certain theories; the present framework shows it is a structural consequence of localized irreversibility. We provide the *why* behind the *what*.

12. Discussion

12.1 Scope and Limitations

The framework applies to substrates with fixed local adjacency and local dynamics. Extensions to the following settings require separate analysis:

- **Dynamical adjacency:** Quantum gravity approaches (causal sets, spin foams) where the adjacency structure itself evolves
- **Fundamentally nonlocal theories:** Theories with action-at-a-distance that cannot be approximated by local dynamics
- **Continuous field theories:** Where "region" and "boundary" require regularization

The present results establish constraints for the discrete, local-dynamics case; generalization is an open problem.

12.2 The Status of Depth

Section 8 establishes that any emergent third dimension is constrained by Proposition 8.1 (No-Storage Constraint, proven) and Conjecture 8.2 (Boundary–Scale Consistency, supported by argument and consistent with known instances). The No-Storage Constraint is a structural necessity: an additional dimension cannot encode recoverable information without violating irreversibility. Boundary–Scale Consistency, while not yet fully proven, is consistent with Bekenstein bounds, Ryu-Takayanagi formulas, and tensor network area laws.

The metric properties of depth—why it admits smooth, geometric description—remain partially open. We conjecture that the regularity of physical coarse-graining (locality preservation, monotonicity) induces metric structure, but a complete derivation is future work. The formal Definition 8.1 (Scale Index) provides the foundation for such a derivation.

12.3 Relation to Other Dimensional Analyses

Prior work has asked "why 3+1 dimensions?" from anthropic, stability, or dynamical perspectives. Our analysis is complementary: we ask what geometric structure is *necessary* for certain informational functions. The answer (2D for facts, with depth as a description-level index) may or may not coincide with other approaches' conclusions, but it addresses a distinct question.

12.4 Experimental and Theoretical Tests

The framework suggests several avenues for testing:

- **Dimensional reduction in coherent systems:** Highly coherent quantum systems should exhibit effective dimensional reduction as the need for coarse-graining hierarchy diminishes.
- **1D memory limitations:** Attempts to build stable memory in strictly 1D substrates (without nonlocal control or external reservoirs) should fail in characteristic ways predicted by Lemma 4.3.

- **Tensor network structure:** The emergence of "radial" directions in tensor network representations of physical systems should correlate with coarse-graining structure, not fundamental geometry.

13. Conclusion

For General Readers

Different spatial dimensions serve different purposes:

- One dimension suffices for orderly, reversible change.
- Two dimensions are minimum for permanent facts—where genuine information loss is localized. This requires both boundary scaling (capacity) and enclosure (protection).
- Three dimensions parameterize the hierarchy of effective descriptions at different scales.

The deep structure of spacetime is about the minimal architecture for reality to contain facts rather than just eternal flow.

Technical Summary

Results:

1. Constant-separator substrates suffice for ticks but cannot support scalable fact-storage due to boundary bottleneck (Theorem 7.1A, Lemma 4.3).
2. 2D-like substrates are minimal for bits, requiring both (i) boundary scaling $|\partial R| = \Omega(\sqrt{|R|})$ and (ii) generic local cycles enabling enclosure (Theorem 7.1B, Proposition 6.1). Trees show that (i) alone is insufficient.
3. Third dimension is constrained to be a scale index parameterizing effective descriptions, not storing discarded information (Proposition 8.1; Conjecture 8.2). The No-Storage Constraint is a structural necessity; Boundary–Scale Consistency is conjectured with strong supporting evidence.

Contributions:

- Separator-based characterization linking dimension to localization capacity
- Two-condition necessity theorem: boundary scaling + enclosure
- Explicit construction showing any bounded-loss irreversible map admits localization in 2D
- No-Storage Constraint (Prop. 8.1) proving emergent dimensions cannot retain discarded information
- Boundary–Scale Consistency (Conj. 8.2) connecting fact-localization to area laws
- Structural explanation for holographic entropy formulas
- Falsifiable criteria for dimensional emergence

Open Questions:

- Deriving depth's metric structure from coarse-graining regularity
- Extension to dynamical adjacency
- Quantitative dimensional reduction thresholds

Dimensionality is constrained by the minimal requirements for making facts—structural foundations independent of specific dynamical laws.

References

- Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D*, 7(8), 2333.
- Bennett, C. H. (1973). Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6), 525–532.
- Chung, F. R. K. (1997). *Spectral Graph Theory*. American Mathematical Society.
- Hoory, S., Linial, N., & Wigderson, A. (2006). Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4), 439–561.
- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183–191.
- Lieb, E. H., & Robinson, D. W. (1972). The finite group velocity of quantum spin systems. *Communications in Mathematical Physics*, 28(3), 251–257.
- Lipton, R. J., & Tarjan, R. E. (1979). A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2), 177–189.
- Maldacena, J. (1999). The large-N limit of superconformal field theories and supergravity. *International Journal of Theoretical Physics*, 38(4), 1113–1133.
- Ryu, S., & Takayanagi, T. (2006). Holographic derivation of entanglement entropy from AdS/CFT. *Physical Review Letters*, 96(18), 181602.
- Swingle, B. (2012). Entanglement renormalization and holography. *Physical Review D*, 86(6), 065007.
- Vidal, G. (2007). Entanglement renormalization. *Physical Review Letters*, 99(22), 220405.
- Wilson, K. G. (1971). Renormalization group and critical phenomena. *Physical Review B*, 4(9), 3174.