

The Topological Threshold for Fact Formation

Plain Language Summary

What this paper is about, in everyday terms:

We ask a deceptively simple question: *How does anything become a fact?*

Think about it this way. Before you flip a coin, both outcomes—heads and tails—are possible. After the coin lands, one outcome becomes real and the other doesn't. But what actually happened at that moment? Where did the "other" outcome go? And why can't you simply unflip the coin?

This paper argues that the answer lies in the *shape* of how information connects to itself—what mathematicians call topology.

Here's the key insight: Imagine information flowing along paths, like water through pipes. In a simple linear arrangement (a straight pipe), any "decision" that splits the flow can always be undone—the streams can rejoin downstream. Nothing is permanent.

But if the pipes form loops, something new becomes possible. One stream can get trapped circulating inside a loop while the other continues onward. Once this happens, the streams can never rejoin without breaking the loop open—which would require action from outside the system.

We prove that this trapping is impossible in one-dimensional structures (straight lines, tree-like branching) but becomes possible the moment you have two-dimensional structure (loops, cycles). This is a sharp threshold, not a gradual transition.

The implications cascade outward:

- **Entropy** (the famous $k_B \ln 2$ of Landauer's principle) emerges because the trapped information still exists—it's just inaccessible.
- **Time's arrow** emerges because trapping is easy but untrapping requires outside help.
- **Facts** are simply trapped distinctions that can no longer be undone.

In short: reality requires loops. Without them, everything remains reversible, provisional, unfixed. The geometry of information itself determines what can become real.

Technical Abstract

We identify the minimal structural condition under which irreversible information—a fact—can first exist. Working within the tick-bit framework, we show that the fixation of a single bit requires a topological transition: the emergence of closed separating structures in the connectivity of information pathways. One-dimensional structures (trees) cannot support local separation; two-dimensional structures can, via cycles that enable alternate routing. The threshold is discrete, the entropy cost is exactly $k_B \ln 2$, and the result provides a structural derivation of Landauer's bound independent of heat-flow assumptions. We formalise these claims using reversible dynamics, coarse-graining maps, and graph homology, proving that if irreversibility is locally constructible from reversible microdynamics, then $\beta_1 \geq 1$ is necessary. This work fixes the entropy scale of the minimal fact; converting this informational scale into physical energy, time, or length scales requires additional substrate-dependent assumptions.

Table of Contents

1. The Problem: How Can a Fact Exist?
 - 1.1 Premises
 - 1.2 Requirements for a Solution
 - 1.3 The Core Idea (General Reader Guide)
2. Formal Framework
 - 2.1 Microstates and Reversible Tick Dynamics
 - 2.2 Bits as Coarse-Grained Facts
 - 2.3 Connectivity and Path Equivalence
 - 2.4 Recombinability
 - 2.5 Graph-Compatible Dynamics
 - 2.6 What This Means (General Reader Guide)
3. Topology Defined Operationally
 - 3.1 What Topology Is Not
 - 3.2 What Topology Is
 - 3.3 Local Modification and Local Traps
 - 3.4 The First Betti Number
 - 3.5 The Threshold Theorem
 - 3.6 Discreteness
 - 3.7 Example
 - 3.8 The Topology of Trapping (General Reader Guide)
4. Dimensional Analysis
 - 4.1 The Void: No Adjacency
 - 4.2 One Dimension: No Local Separation
 - 4.3 Two Dimensions: Local Separation Becomes Possible
 - 4.4 Why Dimensions Matter (General Reader Guide)
5. The Threshold Event
 - 5.1 Mechanism

- 5.2 Where the Discarded Alternative Resides
- 5.3 Why Reversal Requires Nonlocal Action
- 5.4 The Moment of Becoming (General Reader Guide)
- 6. Derivation of Landauer's Bound
 - 6.1 Standard Formulation
 - 6.2 Structural Derivation
 - 6.3 Implications
 - 6.4 Why Erasure Costs Energy (General Reader Guide)
- 7. The Emergence of Time
 - 7.1 The Arrow of Time as Accumulated Irreversibility
 - 7.2 Irreversibility Without Special Initial Conditions
 - 7.3 Time as Accumulated Fact
 - 7.4 Where Time Comes From (General Reader Guide)
 - 7.5 Decoherence as Trap Formation
- 8. Higher Dimensions and Emergent Space
 - 8.1 The Status of Three-Dimensional Space
 - 8.2 Structural Hierarchy
- 9. Related Work
- 10. Objections and Responses
- 11. Summary and Conclusions
 - 11.1 Core Results
 - 11.2 The Foundational Claim
 - 11.3 The Big Picture (General Reader Guide)

Appendix A: Minimal Mathematical Formulation

Appendix B: Notation Summary

Appendix C: Glossary for General Readers

References

1. The Problem: How Can a Fact Exist?

1.1 Premises

The tick-bit framework rests on three premises:

Premise 1. The void may contain reversible information with zero entropy. This is not a claim about physical vacua but a logical stipulation: we admit the possibility of distinction without commitment, variation without record.

Premise 2. Ticks represent reversible variations. A tick is a minimal fluctuation—a distinction that can be made and undone without residue. Ticks may fluctuate freely; their occurrence leaves no permanent trace.

Premise 3. Bits represent irreversible commitments. A bit is a distinction that, once made, cannot be undone locally. Bits do not vary once fixed.

These premises create an immediate problem. If the void contains only reversible information, and ticks are reversible by definition, how does any bit ever become fixed? What mechanism converts provisional distinction into permanent fact?

1.2 Requirements for a Solution

Any viable answer must address three questions:

Q1. Exclusion. How is one alternative excluded permanently? A bit encodes a binary choice: one option is actualised, the other is not. The mechanism must explain how the non-actualised alternative becomes permanently inaccessible.

Q2. Locality. Why can this exclusion not be undone locally? If a local operation could reverse the exclusion, the bit would not be fixed—it would be another tick. The mechanism must explain why reversal requires nonlocal action.

Q3. Ontology. Where does the discarded alternative reside? Information cannot be destroyed without entropy cost. The mechanism must account for the fate of the excluded option.

We will show that all three questions have the same answer: topology.

1.3 The Core Idea (General Reader Guide)

For readers unfamiliar with the technical framework:

Imagine you're at a fork in a path. You can go left or right. Before you choose, both possibilities are "live"—neither is more real than the other. After you choose, one path becomes your actual history and the other becomes... what exactly?

Common sense says the other path simply "didn't happen." But this is philosophically slippery. If information is conserved (a principle supported by physics), the other possibility can't just vanish. It must go *somewhere*.

This paper proposes that the answer depends on the *shape* of the possibility space. In some shapes, every fork eventually rejoins—like streams that split around an island but merge again downstream. In such spaces, no choice is truly permanent; everything remains reversible.

In other shapes, forks can lead to regions that never reconnect—like a side channel that loops back on itself, trapping water that can never return to the main flow. In these spaces, genuine permanence becomes possible. One alternative gets trapped; the other continues. A fact is born.

The remarkable finding: these two types of shapes are separated by a sharp mathematical boundary. Below a certain threshold of complexity, trapping is impossible. Above it, trapping becomes possible. This threshold is defined by whether the space contains *loops*—closed paths that circle back on themselves.

The rest of the paper makes this precise.

2. Formal Framework

2.1 Microstates and Reversible Tick Dynamics

Let Ω be the set of microstates representing all possible informational configurations.

A **tick evolution** is an invertible map:

$$F: \Omega \rightarrow \Omega \text{ (bijection)}$$

If Ω is finite, F is a permutation. Invertibility expresses the premise that ticks are reversible variations: any tick can be undone by applying F^{-1} .

2.2 Bits as Coarse-Grained Facts

A **bit** is a coarse-graining map:

$$b: \Omega \rightarrow \{0, 1\}$$

A bit is **fixed** on a region $A \subseteq \Omega$ if it is forward-invariant under the reversible dynamics:

$$b(F^n(\omega)) = b(\omega) \text{ for all } \omega \in A \text{ and all } n \geq 0$$

This is the formal statement of "once fixed, it does not vary."

Interpretation. A tick is the reversible evolution F . A fact is a property b that becomes dynamically locked-in on accessible states—a coarse-grained distinction that persists despite ongoing microscopic variation.

2.3 Connectivity and Path Equivalence

Let $G = (V, E)$ be a graph encoding adjacency of informational states, where vertices V correspond to microstates and edges E represent permitted single-step transitions. View G as a topological space by treating each edge as a continuous interval glued at vertices; this makes homotopy of paths well-defined.

Path equivalence \sim is homotopy relative to endpoints in this space. For closed paths (loops), this recovers the fundamental group $\pi_1(G)$; its abelianisation gives the first homology group $H_1(G)$.

Two sets of states $A, B \subseteq V$ are **reachable** from each other if there exists a path in G from some $a \in A$ to some $b \in B$.

2.4 Recombinability

Let $\omega_0 \in \Omega$ be an initial microstate, and suppose a branching occurs into sets $A_0, A_1 \subseteq \Omega$ with $b(\omega) = 0$ on A_0 and $b(\omega) = 1$ on A_1 .

The alternatives are **recombinable** if:

$$\exists n, m \geq 0 \text{ such that } F^n(A_0) \cap F^m(A_1) \neq \emptyset$$

The alternatives are **non-recombinable** if:

$$F^n(A_0) \cap F^m(A_1) = \emptyset \text{ for all } n, m \geq 0$$

Non-recombinability is the formal condition for irreversibility: once alternatives separate into non-recombinable sets, they cannot merge back into a common evolution.

When alternatives become non-recombinable, their path classes in π_1 or H_1 become distinct: $[\gamma_1] \neq [\gamma_2]$. A bit corresponds to choosing between two non-equivalent path classes.

2.5 Graph-Compatible Dynamics

Definition 2.5 (Graph-compatible dynamics). Let $G = (V, E)$ be the adjacency graph with $V \cong \Omega$. The reversible dynamics $F: \Omega \rightarrow \Omega$ is G -compatible if for all $\omega \in \Omega$:

$$(\omega, F(\omega)) \in E$$

Equivalently, F maps each state to a graph neighbour in one step. This constraint ensures that the topology of G genuinely governs the dynamics; without it, F could "jump" arbitrarily and G would be irrelevant.

2.6 What This Means (General Reader Guide)

For readers unfamiliar with the mathematical notation:

Think of Ω as the set of all possible "states" the system could be in—like all possible configurations of atoms, or all possible positions on a game board.

The function F describes how the system evolves from one moment to the next. Crucially, F is *reversible*: if you know where the system ended up, you can always figure out where it came from. Nothing is lost at the microscopic level.

A "bit" is what you see when you don't look at all the microscopic detail—just some coarse summary, like "is the particle on the left side or the right side?" The bit is "fixed" if this coarse summary never changes, even as the microscopic details continue to evolve underneath.

The graph G represents which states can transition to which. Two alternatives are "recombinable" if their evolutionary paths can eventually meet up again. They're "non-recombinable" if, no matter how long you wait, they can never rejoin.

The key question: under what conditions can alternatives become non-recombinable? When can the system reach a point of no return?

3. Topology Defined Operationally

3.1 What Topology Is Not

In this framework, topology does not refer to:

- **Curvature.** We make no claims about Riemannian structure.
- **Distance.** There is no metric.
- **Embedding.** The structures need not be embedded in ambient space.
- **Dynamics.** No equations of motion are assumed beyond invertibility.

3.2 What Topology Is

Topology, operationally, is the global connectivity structure that determines whether alternatives can recombine.

A system is **topologically open** (fully reversible) if all branching alternatives are eventually recombinable. No distinction can become permanent.

A system is **topologically closed** (supports irreversibility) if there exist branching alternatives that are non-recombinable. Distinctions can become permanent facts.

3.3 Local Modification and Local Traps

Definition (k-local modification). A local construction modifies G only inside a ball $B_r(x)$ of radius r (bounded number of edges from x), leaving the rest of G unchanged.

Definition (permitted local modifications). A k -local modification is any change to the adjacency structure $G = (V, E)$ obtained by adding or removing edges with both endpoints in $B_k(x)$, while leaving the induced subgraph on $V \setminus B_k(x)$ unchanged and leaving all edges incident to $V \setminus B_k(x)$ unchanged. The reversible dynamics F is held fixed and required to remain G -compatible (Definition 2.5).

Crucially, local modification excludes interventions that alter connectivity along the unique path between alternatives outside the bounded region. This constraint is what makes local separability nontrivial on trees.

Definition (local trap / local separability). A local trap exists if, after a bounded modification confined to $B_r(x)$, there are two alternatives whose future reachability sets become permanently disjoint without changing connectivity outside the ball.

Local separability is the key requirement: bit fixation must be achievable by local operations, not by global restructuring of the entire substrate.

Why locality is assumed. We treat locality as a structural requirement: in physical systems, interactions are mediated by bounded neighbourhoods (finite propagation speed, finite-range couplings, finite control). Accordingly, we restrict "bit fixation" mechanisms to operations that modify connectivity only within a bounded region while leaving the complement unchanged. This is not a claim that global coordination is impossible in principle; rather, Theorem 3.5 establishes a conditional: *if irreversible fact formation is achieved by bounded local construction on an otherwise reversible substrate, then cyclic connectivity ($\beta_1 \geq 1$) is necessary.*

3.4 The First Betti Number

The topological capacity for supporting cycles is captured by the first Betti number:

$$\beta_1(G) = \text{rank } H_1(G)$$

This counts the number of independent cycles in G .

- If G is a tree (acyclic), then $\beta_1(G) = 0$.
- If G contains at least one cycle, then $\beta_1(G) \geq 1$.

3.5 The Threshold Theorem

Lemma 3.5.1 (Two boundary crossings in trees). Let $T = (V, E)$ be a tree and let $B := B_k(x)$ be the ball of radius k around x in graph distance. Let $u, v \in V \setminus B$. Then the unique simple path $P_T(u, v)$ in T satisfies exactly one of:

- (i) $P_T(u, v) \subseteq V \setminus B$; or

(ii) $P_T(u, v) \cap B \neq \emptyset$, and in this case $P_T(u, v)$ intersects the boundary ∂B in exactly two vertices (an "entry" and an "exit"), and the subpath inside B is unique.

Proof. Since T is a tree, there is a unique simple path between any two vertices. Consider the first vertex along $P_T(u, v)$ that lies in B (if any); its predecessor lies outside B and must lie in ∂B . Similarly, the last vertex in B before the path returns to $V \setminus B$ must lie in ∂B . If there were more than two boundary crossings, the path would enter and leave B multiple times, implying two distinct simple routes between some boundary vertices, which would create a cycle—contradiction. \square

Lemma 3.5.2 (Boundary-path rigidity). Let T be a tree and $B = B_k(x)$. For any two boundary vertices $p, q \in \partial B$, the unique path $P_T(p, q)$ is determined by the induced subgraph on $V \setminus B$ together with the set of edges incident to $V \setminus B$. Hence, under permitted local modification, whether p connects to q through $V \setminus B$ cannot change.

Proof. The path $P_T(p, q)$ either lies entirely in $V \setminus B$ (and is thus unchanged by any modification inside B) or passes through B . In the latter case, by Lemma 3.5.1 it enters at some boundary vertex and exits at another. Since permitted local modifications preserve all edges incident to $V \setminus B$, the entry and exit edges remain, and the existence of the path is preserved. \square

Definition (Outside-B recombination witness). A recombination event for A_0, A_1 is *witnessed outside B* if there exist $n, m \geq 0$ and a vertex $w \in V \setminus B$ such that $w \in F^n(A_0) \cap F^m(A_1)$.

Theorem 3.5 (Conditional Cycle Threshold for Locally Constructible Irreversibility). Let G be a connected graph with G -compatible reversible dynamics F (Definition 2.5). Suppose bit fixation requires local separability: there exists a bounded region $B_k(x)$ such that by permitted local modifications confined to $B_k(x)$, one can create two alternatives whose reachable sets become permanently non-intersecting under F , while leaving $G \setminus B_k(x)$ unchanged.

We consider separability of alternatives whose supports are not contained entirely within $B_k(x)$ —that is, alternatives that remain accessible outside the construction region before the modification. Local separability that only traps states already confined to $B_k(x)$ is excluded as trivial.

Then a necessary condition is:

$$\beta_1(G) \geq 1$$

In particular, if G is a tree ($\beta_1(G) = 0$), no bounded modification can produce such a local separability event.

Proof (Tree impossibility). Assume $G = T$ is a tree and let $B := B_k(x)$. Let $A_0, A_1 \subseteq V$ be two alternative sets that are recombinable under F prior to any modification, and assume each has support not contained in B . Let T' be the modified graph obtained by a k -local modification within B as defined in §3.3, and suppose for contradiction that A_0, A_1 become non-recombinable

under the same reversible F while the complement $V \setminus B$ and all edges incident to it are unchanged.

The key observation is that G -compatibility of F links graph structure to dynamics: since F is G -compatible, any recombination event $F^n(a_0) = F^m(a_1)$ requires a walk in G connecting a_0 to $F^n(a_0)$ and another connecting a_1 to the same vertex. Specifically, the sequence $a_0 \rightarrow F(a_0) \rightarrow F^2(a_0) \rightarrow \dots \rightarrow F^n(a_0)$ traces a walk in G , and similarly for a_1 .

Consider any recombination witnessed outside B before modification: some n, m and $w \in V \setminus B$ with $w \in F^n(A_0) \cap F^m(A_1)$. Because F is unchanged and $V \setminus B$ adjacency is unchanged, any such witness w remains a witness after modification. The walks from A_0 to w and from A_1 to w may pass through B , but by Lemma 3.5.1 they enter and exit through boundary vertices, and by Lemma 3.5.2 these boundary connections are preserved.

Therefore, to destroy recombination one must eliminate all recombination witnesses, which requires ensuring no witness exists in $V \setminus B$. But since alternatives have support outside B and F is unchanged, orbit segments in $V \setminus B$ are unaffected. Any previous outside- B witness remains valid after modification.

Hence no k -local modification can render such alternatives non-recombinable on a tree. \square

Remark (Cycles inside B). In the tree case, permitted local modification may create cycles inside B (by adding edges with both endpoints in B). However, since all edges incident to $V \setminus B$ are fixed, such internal cycles cannot produce separability for alternatives with support extending outside B . The cycle exists but provides no "escape route" that could trap external alternatives.

Remark (Modification space on trees). The constraint is actually stronger than the proof requires. On a tree with fixed G -compatible dynamics F , permitted local modifications are limited to removing edges not traversed by any F -orbit segment (adding edges would create cycles inside B , but these don't affect external recombination as noted above). The theorem is therefore nearly vacuous for trees: the modification space is almost empty precisely because tree topology offers no "slack" for rerouting external alternatives.

Remark (Why cycles enable local trapping). If $\beta_1(G) \geq 1$ globally (not just inside B), there exist vertex pairs connected by multiple internally disjoint paths. Bounded modifications can remove or redirect one path locally while leaving another intact, enabling persistent non-recombination without touching the complement. The global cycle provides the "slack" that trees lack.

Corollary. Any physical system exhibiting genuine local irreversibility must admit cyclic information pathways in its effective substrate description.

3.6 Discreteness

The first Betti number is integer-valued:

$$\beta_1(G) \in \mathbb{Z}_{\geq 0}$$

The threshold corresponds to:

$$\beta_1: 0 \rightarrow 1$$

There is no "partial cycle." The discreteness of the topological transition propagates to the informational level: either local separability is possible, or it is not. Either a bit can be fixed locally, or it cannot.

3.7 Example: The Lollipop Graph

Let G be the "lollipop" graph with vertices $\{1, 2, 3, 4, 5\}$ and edges forming a path $1-2-3$ attached to a triangle $3-4-5-3$. This graph has $\beta_1 = 1$ (one independent cycle).

Dynamics. Define the G -compatible bijection F as follows:

- $F(1) = 2, F(2) = 3$ (flow along the path toward the cycle)
- $F(3) = 4, F(4) = 5, F(5) = 3$ (rotation around the triangle)

Note that F is a bijection: the path portion feeds into the cycle, and the cycle rotates. The inverse F^{-1} rotates the cycle backward and pulls states back along the path.

Initial recombability. Consider alternatives $A_0 = \{1\}$ and $A_1 = \{4\}$. Under F :

- $F^2(1) = 3, F^3(1) = 4$, so $F^3(A_0) \cap A_1 \neq \emptyset$
- The alternatives are recombable: starting from vertex 1 and vertex 4, trajectories meet at vertex 4 after 3 steps from vertex 1.

Coarse-graining. Define $b: V \rightarrow \{0, 1\}$ by $b(v) = 0$ for $v \in \{1, 2\}$ and $b(v) = 1$ for $v \in \{3, 4, 5\}$. This distinguishes "path states" from "cycle states."

Local modification. Consider the ball $B_1(4) = \{3, 4, 5\}$ (all vertices within distance 1 of vertex 4). A permitted local modification within $B_1(4)$ could remove edge $3-4$ and add edge $2-4$, creating a new graph G' where the cycle $3-5-4-2-3$ exists but vertex 3 is no longer directly connected to vertex 4.

However, a simpler modification suffices: modify F to F' where $F'(3) = 5$ instead of $F'(3) = 4$, keeping G unchanged. Now:

- $F'^3(1) = F'(F'(F'(1))) = F'(F'(2)) = F'(3) = 5$
- The trajectory from vertex 1 enters the cycle at vertex 5, not vertex 4
- Trajectories starting at vertex 1 cycle through $\{3, 5\}$ while vertex 4 becomes trapped in its own orbit

Result. After modification, alternatives starting outside the ball can be routed into different cycling patterns within the ball. The cycle provides multiple entry points and internal routes, enabling local separation that would be impossible on a tree.

Contrast with a tree. If we removed the edge 4–5 (making G a tree), the only path from vertex 1 to vertex 4 would be 1–2–3–4. Any modification to F that blocks this path would either break G -compatibility or require changing edges incident to vertices outside any bounded ball containing the modification. The tree offers no alternative route.

3.8 The Topology of Trapping (General Reader Guide)

For readers unfamiliar with topology:

Topology is the mathematics of connectivity—what's connected to what, and how. It ignores details like distance or angle; it cares only about the overall "shape" of connections.

The key concept here is the **loop** (or cycle). A loop is a path that returns to its starting point. The question "does this network contain loops?" turns out to have profound consequences.

Networks without loops (trees): Think of a family tree or a corporate org chart. There's exactly one path between any two points. If you want to block communication between two people, you have to cut the single path connecting them—and that cut affects everyone on either side. You can't create isolated pockets without affecting the whole structure.

Networks with loops: Think of a city street grid. There are multiple paths between any two points. If you block one street, traffic can route around. More importantly, you can create isolated zones—close off a few streets and you have a cul-de-sac that traffic enters but never leaves (unless you open the streets again).

The **Betti number** β_1 counts how many independent loops a network has. A tree has $\beta_1 = 0$ (no loops). A figure-eight has $\beta_1 = 2$ (two independent loops). The number must be a whole number—you can't have half a loop.

The theorem says: **you need at least one loop ($\beta_1 \geq 1$) to create a local trap.** Without loops, any attempt to isolate something requires changing the whole network. With loops, you can create cul-de-sacs—regions that capture things permanently through purely local modifications.

This is why the threshold is sharp. Either you have loops or you don't. Either local trapping is possible or it isn't.

4. Dimensional Analysis

4.1 The Void: No Adjacency

The void, by definition, contains no structure. There are no adjacency relations—no sense in which one state is "next to" another.

Without adjacency, trajectories cannot be defined. Without trajectories, there is nothing to separate. The question "can alternatives recombine?" is malformed.

Conclusion. The void cannot host the topological threshold.

4.2 One Dimension: No Local Separation

In \mathbb{R} (or any tree-structured graph), adjacency exists. Trajectories can be defined. But trees have $\beta_1 = 0$: no cycles.

In \mathbb{R} , every connected subset is an interval. Intervals do not separate \mathbb{R} into an "inside" and "outside" the way a closed curve does in \mathbb{R}^2 . Specifically: $\mathbb{R} \setminus \{x\}$ disconnects into two rays $(-\infty, x)$ and (x, ∞) , but no bounded set has an interior separated from an exterior by a closed boundary.

By the theorem above, trees cannot support local separability. Any attempt to quarantine an alternative requires global modification—affecting the unique path that connects regions.

Conclusion. One-dimensional structures ($\beta_1 = 0$) cannot host the topological threshold via local operations.

4.3 Two Dimensions: Local Separation Becomes Possible

Two-dimensional structures introduce cycles: $\beta_1 \geq 1$.

In a pure graph, $\beta_1(G) \geq 1$ guarantees nontrivial loop classes in $H_1(G)$ —alternatives that cannot be deformed into one another within the 1-complex. This is sufficient for non-recombinability.

If the connectivity is additionally realised as a planar embedding or as a 2-complex with faces, then simple closed cycles serve as boundaries of regions. By the **Jordan curve theorem**, a simple closed curve $\gamma \subset \mathbb{R}^2$ divides the plane into exactly two path-connected components (interior and exterior), with γ as their common boundary. Any path from interior to exterior must cross γ . This gives a literal "trap" interpretation: trajectories in the interior are geometrically enclosed.

Key distinction:

- **Algebraic (always available when $\beta_1 \geq 1$):** Nontrivial loop classes enable non-recombinability via alternate routing.
- **Geometric (when embedded in 2D or higher):** Jordan separation provides literal interior/exterior regions.

Both suffice for the threshold. The algebraic condition is more general; the geometric interpretation is more intuitive.

4.4 Why Dimensions Matter (General Reader Guide)

For readers unfamiliar with dimensional analysis:

Why does the number of dimensions matter for whether facts can exist?

One dimension (a line): Imagine ants walking along a wire. If two ants are heading toward each other, they *must* eventually meet—there's only one path. You can't build a trap on a wire because there's no way to route around it. Any barrier affects everything on both sides.

Two dimensions (a surface): Now imagine ants on a tabletop. If two ants are approaching the same point, one can detour around a barrier while the other walks into a corral and gets trapped. The extra dimension provides "room to maneuver"—alternative routes that make local trapping possible.

The jump from one to two dimensions isn't gradual. There's no such thing as "1.5 dimensions" in this context. Either alternative routes exist or they don't. Either trapping is possible or it isn't.

This is why we say the threshold is *topological*: it depends on the overall shape of connections, not on any continuous parameter you could dial up or down.

5. The Threshold Event

5.1 Mechanism

The topological threshold is crossed when a separating structure completes in a way that renders alternatives non-recombinable.

This process requires no external agent. The routing of trajectories into separated classes arises from the reversible microdynamics F whenever the adjacency structure (with $\beta_1 \geq 1$) admits configurations where one alternative enters a cycle and the other does not.

Formally, the threshold occurs when:

1. A local modification creates or activates a cycle-based trap.
2. The state space decomposes into non-recombinable regions under F .
3. At least one alternative becomes confined to the trap.

At this moment, exactly one bit becomes fixed.

5.2 Where the Discarded Alternative Resides

The excluded trajectory does not disappear. Information is conserved.

Lemma (Irreversible coarse-graining requires an environment). If $F: \Omega \rightarrow \Omega$ is bijective but the observed evolution reduces distinguishability—merging alternatives under the coarse-graining map b —then there must exist an auxiliary register E and a bijection:

$$U: \Omega \times E \rightarrow \Omega \times E$$

such that the full evolution remains reversible. The "missing" information is stored in E .

Interpretation. The discarded alternative resides as correlation information in E —the inaccessible degrees of freedom required to keep the underlying evolution bijective.

In topological terms: the cycle-based trap corresponds to the part of state space that stores correlation with E . It cannot be accessed or recombined without operations that cross the separating structure—nonlocal action requiring entropy export.

This answers Q3: the excluded alternative is not annihilated. It is topologically quarantined in an invariant region that maintains microdynamic reversibility while being inaccessible to the coarse-grained observable.

5.3 Why Reversal Requires Nonlocal Action

To reverse bit fixation, the quarantined trajectory must rejoin the accessible region. This requires "opening" the trap—modifying the cycle structure.

But the trap, once formed, is an invariant set under F . Opening it requires coordinated modification that affects the entire cycle. No operation confined to the trap's interior can break the invariance.

Furthermore, merging previously separated alternatives decreases the distinguishability of the accessible system. This requires exporting entropy—at minimum, $k_B \ln 2$.

5.4 The Moment of Becoming (General Reader Guide)

For readers seeking intuitive understanding:

What actually happens at the threshold?

Picture a river delta where water flows through branching channels. In most deltas, all branches eventually reach the sea—the water "recombines" at the ocean. Any split is temporary.

Now imagine one branch leads into a circular lagoon with a one-way valve at the entrance. Water can flow in, but it can never flow back out to rejoin the main channels. Once water enters the lagoon, its fate is sealed.

The "threshold event" is the moment when such a valve closes. Before: both paths are open, both fates are possible. After: one path leads to the ocean, one leads to eternal circulation in the lagoon. A distinction has become permanent.

Where does the trapped water go? It doesn't vanish—it's still there, circling in the lagoon. But it's *inaccessible* to the main flow. From the perspective of the river, that water is gone. From the perspective of the universe, it's merely quarantined.

Why can't you undo it? To release the trapped water, you'd have to open the valve—but the valve is part of the lagoon structure, not something you can reach from inside. You need external intervention.

This is the deep structure of all irreversibility: not destruction, but quarantine. Not disappearance, but inaccessibility. The trapped alternative is as real as the actualised one—just unreachable.

6. Derivation of Landauer's Bound

6.1 Standard Formulation

Landauer's principle asserts that erasing one bit of information requires dissipating at least $k_B T \ln 2$ of energy as heat. The standard derivation proceeds thermodynamically via phase-space compression and the second law.

6.2 Structural Derivation

We derive Landauer's bound from bijective microdynamics and coarse-graining, independent of heat-flow assumptions.

Assume Ω is finite. Suppose a bit-fixing event maps two distinguishable alternatives into one observable outcome (erasure of one bit of distinguishability). The coarse-graining map $b: \Omega \rightarrow \{0, 1\}$ has fibres of size at least 2:

$$\exists y \in \{0, 1\} : |b^{-1}(y)| \geq 2$$

Step 1 (Reversibility constraint). Because the microdynamics F is bijective, the "missing" distinguishability cannot vanish. The joint evolution $U: \Omega \times E \rightarrow \Omega \times E$ must be bijective, with the discarded information stored in E .

Step 2 (Counting). Distinguishing the two preimages of the coarse-grained output requires at least two distinguishable environment states:

$$|E| \geq 2$$

Step 3 (Information). If the environment degree of freedom is inaccessible to the observer (quarantined), then from the observer's perspective there are at least two micro-realizations consistent with the same observed macrostate. The minimal increase in missing information is:

$$\Delta H \geq \ln 2 \text{ (in nats, Shannon)}$$

Step 4 (Entropy). Converting Shannon information to physical entropy uses the standard identification:

$$\Delta S = k_B \Delta H \geq k_B \ln 2$$

This is Landauer's bound in entropy form.

Step 5 (Energy, when applicable). The usual energetic statement $\Delta Q \geq k_B T \ln 2$ follows when the entropy export occurs to a thermal reservoir at temperature T .

6.3 Implications

The counting result $\Delta H \geq \ln 2$ follows from bijective microdynamics plus coarse-graining and is independent of heat-flow assumptions. Translating this missing information into entropy units introduces k_B via the standard bridge $\Delta S = k_B \Delta H$. In this sense, topology forces the counting bound $\Delta H \geq \ln 2$, and thus fixes the numerical constant that appears in Landauer's entropy form.

Any system capable of fixing bits—any system with $\beta_1 \geq 1$ enabling local traps—necessarily incurs the $k_B \ln 2$ cost, because the excluded alternative must be stored, and minimal storage is binary.

Scope of the result. This work fixes the entropy scale of the minimal fact: one bit, corresponding to an irreducible entropy increase of $k_B \ln 2$. Converting this informational scale into a physical energy, time, or length scale requires additional substrate-dependent assumptions (such as temperature, noise floor, or finite distinguishability threshold). The topological framework determines *that* a minimal cost exists and *what* its informational magnitude is; it does not by itself determine the physical units in which that cost manifests for any particular system.

6.4 Why Erasure Costs Energy (General Reader Guide)

For readers unfamiliar with Landauer's principle:

In 1961, physicist Rolf Landauer discovered something remarkable: there's a fundamental minimum energy cost to erasing information. Specifically, erasing one bit costs at least $k_B T \ln 2$ of energy, where T is temperature and k_B is Boltzmann's constant.

At room temperature, this is a tiny amount—about 3×10^{-21} joules. But the principle is profound: it connects information to physics. Information isn't just abstract; it has physical consequences.

Why does erasure cost energy? The standard explanation involves thermodynamics: erasure reduces the number of possible states, which decreases entropy, which must be compensated elsewhere.

Our derivation is different. We show the cost emerges from topology:

1. When a bit gets fixed, one alternative gets trapped.
2. The trapped alternative doesn't vanish—it's stored somewhere inaccessible.
3. That "somewhere" must have at least 2 states (to record which alternative was discarded).
4. Having 2 states you can't see means missing exactly $\ln 2$ nats of information.
5. That's the entropy cost: $k_B \ln 2$.

The energy cost arises when you try to *reset* the trap for reuse. The trapped information has to go somewhere—typically, it gets dumped as heat into the environment.

The punchline: Landauer's bound isn't fundamentally about heat or thermodynamics. It's about the topology of information. Heat is just the most common way the trapped information eventually manifests in physical systems.

7. The Emergence of Time

7.1 The Arrow of Time as Accumulated Irreversibility

Before the threshold, all distinctions are recombining. There is no arrow of time because "forward" and "backward" evolution both preserve all alternatives.

After bits begin to fix, symmetry breaks. Each fixed bit creates an asymmetry:

- **Past:** The state before the bit fixed, when alternatives were recombining.
- **Future:** The state after the bit fixed, when alternatives are non-recombining.

The arrow of time points in the direction of increasing bit count—increasing numbers of topologically quarantined commitments. This is enabled by a substrate with $\beta_1 \geq 1$, but it is the accumulation of fixed bits, not any change in β_1 itself, that constitutes temporal direction.

7.2 Irreversibility Without Special Initial Conditions

Standard physics attributes the arrow of time to special initial conditions (low-entropy Big Bang). This leaves unexplained why initial conditions were special.

The topological account offers an alternative. The arrow of time emerges automatically whenever ticks aggregate into bits, because bit fixation is inherently asymmetric: traps can close but cannot open without nonlocal action and entropy export.

No special initial conditions are required. The asymmetry is structural.

Qualification (content vs asymmetry). The structural claim here concerns the *existence* of temporal asymmetry: once traps can form, closure events create records that are hard to erase. The particular *content* of what becomes trapped at any given event—that is, which alternative is quarantined—depends on the microstate at the moment of trap formation and therefore can depend on initial conditions. The point is that such dependence selects *which* facts occur, not *whether* an arrow of time exists once fact formation is possible.

7.3 Time as Accumulated Fact

On this view, time is not a background dimension. Time is the accumulated record of irreversible distinctions—the count of fixed bits.

Each fixed bit represents one unit of elapsed time: the only time that exists is the time measured by the accumulation of fact.

7.4 Where Time Comes From (General Reader Guide)

For readers interested in the nature of time:

Why does time flow in one direction? Why can we remember the past but not the future? Why can you unscramble an egg?

Physics has long struggled with these questions. The fundamental laws—quantum mechanics, general relativity—work equally well forward and backward. Yet we experience a clear arrow of time.

The standard answer involves the Big Bang: the universe started in a special low-entropy state and has been "unwinding" ever since. But this just pushes the question back: why was the initial state special?

Our answer is structural. Time's arrow emerges from topology, not cosmology.

The key insight: traps are easy to fall into but hard to escape. A one-way valve doesn't need any special setup—it just works asymmetrically by its nature. Once something enters, getting it out requires external help.

If reality is built from irreversible distinctions (bits), and irreversibility requires topological traps, then time's arrow is automatic. Every time a distinction becomes permanent, the universe accumulates one more "fact," and that accumulation defines the direction of time.

Time isn't something that flows. Time is something that *accumulates*. Each trapped bit is a tick of the only clock that matters: the count of things that have become irreversibly true.

This explains why we remember the past: memories are stored bits, and bits can only be fixed, not unfixed (without external entropy cost). The "past" is simply the collection of already-fixed bits. The "future" is where new bits are yet to be fixed.

7.5 Decoherence as Trap Formation

Consider a quantum system S coupled to an environment E . Microscopic evolution is unitary (bijective on the underlying state space), but coarse-grained observables on S exhibit effective irreversibility because which-alternative information becomes encoded in correlations with E .

In the present framework, this is exactly Lemma 5.2 (environment register): coarse-graining that merges alternatives requires an auxiliary register. Decoherence corresponds to the regime in which the environment correlation record becomes dynamically inaccessible—recoherence would require controlling an exponentially large set of environmental degrees of freedom, which is nonlocal in the operational sense of §3.3.

Topologically, the "cycle" condition is not a literal geometric loop in space but the existence of alternate routing channels in the effective information graph: branches in S that would be recombinable in isolation become non-recombinable because their trajectories diverge into disjoint regions of $\Omega \times E$. The pointer basis can be interpreted as the coarse-graining b whose fibres align with these disjoint reachability classes.

Toy example: Two qubits. Consider a system qubit S coupled to an environment qubit E . The joint state space has four basis states: $|00\rangle$, $|01\rangle$, $|10\rangle$, $|11\rangle$. Consider the unitary (bijection) that acts as a controlled-NOT: $F|s,e\rangle = |s, e \oplus s\rangle$.

The effective information graph G has four vertices $\{00, 01, 10, 11\}$ with edges determined by F :

- $00 \rightarrow 00$ (fixed point)
- $01 \rightarrow 01$ (fixed point)
- $10 \rightarrow 11 \rightarrow 10$ (2-cycle)

This graph has $\beta_1 = 1$: the 2-cycle $10 \leftrightarrow 11$ forms a loop. The system states $|0\rangle_S$ and $|1\rangle_S$ correspond to coarse-graining over environment states. Initially, superpositions of $|0\rangle_S$ and $|1\rangle_S$ are recombinable (both map to themselves or cycle together). After entanglement with E , the $|1\rangle_S$ branch enters the cycle while $|0\rangle_S$ remains at fixed points—they become non-recombinable from the perspective of S -only observations.

This does not claim that decoherence requires a spatial loop; it claims that effective irreversibility requires cyclic connectivity in the extended information substrate—that is, $\beta_1 \geq 1$ in the effective transition structure of $S \cup E$. The environment provides the minimal "trap capacity" required for facts (records) to form.

Research direction. Constructing the effective information graph for realistic physical systems (e.g., a qubit coupled to a thermal bath with many modes) and verifying $\beta_1 \geq 1$ is a topic for future work. The exponential growth of Hilbert space dimension makes this computationally challenging, but the framework predicts that any system exhibiting effective irreversibility must have $\beta_1 \geq 1$ in its effective transition structure when the environment is included.

Mapping of roles:

Decoherence concept	Tick-bit framework
System S	Accessible region A
Environment E	Trap/discard region D
Unitary evolution	Bijjective F on $\Omega \times E$
Pointer basis	Coarse-graining b aligned with reachability classes
Loss of coherence	Non-recombinability of alternatives
Einselection	Formation of invariant partition under F

8. Higher Dimensions and Emergent Space

8.1 The Status of Three-Dimensional Space

The threshold requires $\beta_1 \geq 1$, achievable in two-dimensional structures. Physical space appears three-dimensional. What is the relationship?

The tick-bit framework suggests that apparent three-dimensional space is not fundamental. It emerges as systematic organisation over many fixed bits—bookkeeping of layered irreversible distinctions.

A full treatment of spatial emergence is beyond this paper's scope. The key point: the topological threshold does not require three dimensions. It operates in the minimal structure that admits cycles.

If three-dimensional space emerges from bit accumulation, then the threshold is logically prior to space.

8.2 Structural Hierarchy

Structure	β_1	Information Status
Void	Undefined	No trajectories; no distinction
Tree (1D-like)	0	Trajectories exist; no local traps
Graph with cycle (2D-like)	≥ 1	Local traps possible; bits can fix
2-complex with faces	≥ 1	Jordan separation; geometric traps

Structure	β_1	Information Status
Higher complexes	$\beta_1 \geq 1$, possibly $\beta_2 \geq 1$, etc.	Higher-dimensional trap structures (beyond scope)

The critical transition occurs at $\beta_1: 0 \rightarrow 1$, where local separability first becomes possible. Higher Betti numbers (β_2 counting voids, β_3 counting cavities, etc.) may enable richer trap structures, but their information-theoretic significance is not explored in this paper.

Remark. A connected graph is a tree if and only if it is acyclic, equivalently if and only if $|E| = |V| - 1$. See Diestel [10] or Bollobás [9] for standard characterisations.

9. Related Work

This framework sits within a broader tradition of information-theoretic approaches to fundamental physics. Wheeler's "it from bit" program [40] proposed that physical reality emerges from information-theoretic primitives; our work can be seen as identifying the minimal topological conditions under which Wheeler's "bits" can first become fixed. 't Hooft's cellular automaton interpretation of quantum mechanics [41] similarly posits discrete, deterministic dynamics underlying quantum phenomena; our framework shares the emphasis on bijective microdynamics but focuses on the coarse-graining threshold rather than the quantum-classical correspondence.

The connection between topology and computation has been explored in various contexts, notably in topological quantum computing where anyonic braiding provides protected gates [45]. Our use of topology is different: we are not concerned with computational universality but with the more primitive question of when irreversible distinctions can arise at all. The reversible computation literature [30, 31, 32] provides the technical foundation for our treatment of environment registers and Landauer's bound.

Several authors have proposed that time's arrow emerges from information-theoretic considerations rather than cosmological boundary conditions [35, 36]. Our contribution is to identify a specific structural threshold—the $\beta_1: 0 \rightarrow 1$ transition—that enables such emergence. This complements rather than replaces thermodynamic accounts: we derive the *existence* of irreversibility from topology, while thermodynamics describes its *quantitative* behaviour in equilibrium contexts.

10. Objections and Responses

10.1 "This is just a metaphor."

Objection. The language of "traps," "cycles," and "quarantine" is metaphorical.

Response. The Betti number β_1 is a well-defined topological invariant. The theorem relating $\beta_1 \geq 1$ to local separability is a mathematical result about graphs. The claim that information topology constrains fact formation is testable: genuine irreversibility should not occur via local operations in systems with $\beta_1 = 0$.

10.2 "What are the trajectories made of?"

Objection. You never say what microstates or ticks fundamentally are.

Response. A trajectory is a sequence of microstates under F . The ontological status of microstates is a boundary condition of the framework. We explain how ticks become bits, not what ticks are.

10.3 "Where do the two dimensions come from?"

Objection. You show $\beta_1 \geq 1$ is necessary but don't explain why cyclic structures exist.

Response. Correct. The framework explains what is possible given structure; it does not explain the origin of structure.

10.4 "This contradicts quantum mechanics."

Objection. Quantum mechanics allows superposition. Your framework requires exclusion.

Response. Superposition is a feature of ticks (reversible, recombinable alternatives). Quantum coherence persists when no local trap has formed. Decoherence and measurement correspond to the threshold event—formation of a trap rendering alternatives non-recombinable. This aligns with decoherence-based interpretations.

10.5 "Landauer's bound is about energy, not topology."

Objection. Landauer's principle concerns heat. Your derivation involves no heat.

Response. Both derivations yield $k_B \ln 2$. The standard derivation assumes thermodynamics; ours assumes bijective dynamics and coarse-graining. That both yield the same constant suggests topology underlies thermodynamics.

10.6 "You smuggled in k_B and entropy."

Objection. You claim independence from thermodynamics but use Boltzmann's constant.

Response. The structural result is counting: fixing one bit requires storing one binary alternative. The constant k_B enters only when translating counting to entropy units—a definitional bridge, not a physical assumption.

10.7 "How is this testable?"

Objection. What experiment could falsify the framework?

Response. The framework predicts that genuine bit fixation cannot occur via local operations in systems with $\beta_1 = 0$ (trees, linear chains). Demonstrating local irreversibility in a strictly tree-structured information topology would falsify the theory.

In particular, the framework predicts that no system whose effective information graph is provably tree-structured can exhibit genuine bit fixation via bounded local operations, regardless of timescale. Any experimental demonstration of such fixation—not merely apparent irreversibility due to practical limitations, but genuine non-recombinability—would constitute a falsification.

Practical testability. In practice, one tests the claim at the level of an engineered effective connectivity: construct a system whose allowed transitions (or couplings) implement a known graph G , and verify whether locally induced bit fixation occurs. The falsification target is not "a naturally occurring system proven to have $\beta_1 = 0$ at all scales," but an experimentally controlled substrate whose effective transition graph is enforced to be a tree.

Candidate experimental platforms:

1. **Engineered spin chains / 1D circuits:** Nearest-neighbour spin chains or superconducting qubit chains with strictly nearest-neighbour couplings implement a tree-like (path) interaction graph. Test whether local operations can produce persistent, non-recombinable coarse-grained outcomes.
2. **Tree-structured reversible circuits:** Implement reversible gates on a tree wiring graph; test whether local operations can produce persistent bit fixation without introducing hidden cyclic couplings.
3. **Photonic waveguide trees:** Interferometric networks constrained to tree connectivity (no loops) where recombination is enforced by design; test whether "measurement-like" irreversibility can be induced without adding looped ancilla modes.

What would constitute "loop leakage": Stray couplers, uncontrolled environmental modes, or any unintended cyclic pathway would introduce $\beta_1 > 0$ and invalidate the tree constraint. Careful isolation and characterisation of the effective connectivity graph is essential.

Operational verification of $\beta_1 = 0$: In engineered platforms, verifying $\beta_1 = 0$ reduces to verifying that the implemented coupling/transition graph is acyclic (e.g., via tomography of couplings or network reconstruction), up to known error bounds. This makes "provably tree-structured" an experimental criterion rather than a metaphysical one.

10.8 "Any bijection has invariant subsets. Why is β_1 relevant?"

Objection. I can always find invariant sets under a permutation. Your theorem seems to prove too much.

Response. The theorem is about *local* construction of traps. Yes, any bijection has invariant subsets (unions of cycles in the permutation). But the question is whether such subsets can be *locally created* to separate alternatives that were previously recombinable. On trees, any separation requires modifying the unique path between alternatives—a global operation. With $\beta_1 \geq 1$, alternate routes exist, enabling local trapping without global modification.

11. Summary and Conclusions

11.1 Core Results

1. **Topology is connectivity.** Topology determines whether alternative trajectories can recombine, formalised via reachability and path equivalence.
2. **$\beta_1 \geq 1$ is necessary for local separability.** Trees ($\beta_1 = 0$) cannot support local trap formation; any separation requires global modification.
3. **Two-dimensional structure is minimal.** Cycles first appear with $\beta_1 \geq 1$, enabling alternate routing and local traps.
4. **The threshold is discrete.** $\beta_1 \in \mathbb{Z}_{\geq 0}$ admits no partial values. Either local trapping is possible, or it is not.
5. **Discarded alternatives persist.** Irreversibility is inaccessibility, not annihilation. The excluded alternative is stored in an environment register maintaining microdynamic reversibility.
6. **Landauer's bound is structural.** The cost $k_B \ln 2$ follows from bijective dynamics plus coarse-graining, independent of heat-flow assumptions.
7. **Time emerges from bit accumulation.** The arrow of time is the direction of increasing bit count on a substrate with $\beta_1 \geq 1$.

11.2 The Foundational Claim

Fact formation has a topological threshold. Below the threshold ($\beta_1 = 0$), all distinctions are provisional and local trapping is impossible. At the threshold ($\beta_1: 0 \rightarrow 1$), local separation becomes possible, bits can fix, and discarded alternatives become quarantined. Above the threshold, facts accumulate into the layered structure we call physical reality.

Ticks vary; bits do not. This asymmetry, grounded in the topology of cyclic connectivity, is the origin of irreversibility, entropy, and time.

11.3 The Big Picture (General Reader Guide)

For readers seeking the philosophical implications:

What have we actually learned?

Reality requires loops. The ability for anything to become permanently, irreversibly true depends on the shape of possibility space. If that space is too simple (no cycles), everything remains eternally provisional—facts cannot form. Only when cycles exist can some things become trapped, excluded, finalised.

Irreversibility is quarantine, not destruction. When one outcome becomes actual and another becomes "merely possible," the merely possible doesn't vanish. It's still there, stored in degrees of freedom we can't access. This is why information is never destroyed, why entropy always increases, why erasure costs energy.

Time is accumulated fact. The arrow of time isn't imposed from outside or dependent on special initial conditions. It emerges naturally from the asymmetry of trapping: easy to get in, hard to get out. Every irreversible event is a ratchet click forward. Time is the sum of those clicks.

The threshold is sharp. There's no gradual transition between a world where facts are possible and one where they're not. Either loops exist or they don't. Either reality can crystallize or it remains eternally fluid. This suggests that the very possibility of a factual world is a discrete, all-or-nothing matter.

If this framework is correct, the deep structure of reality is topological. The emergence of facts, the arrow of time, the cost of forgetting—all trace back to the simple question of whether information pathways can form closed loops.

Appendix A: Minimal Mathematical Formulation

Definition A1 (Reversible tick dynamics). A tick evolution is a bijection $F: \Omega \rightarrow \Omega$.

Definition A2 (Bit as coarse-graining). A bit is a map $b: \Omega \rightarrow \{0, 1\}$. It is fixed on $A \subseteq \Omega$ if $b(F^n(\omega)) = b(\omega)$ for all $\omega \in A$ and $n \geq 0$.

Definition A3 (Recombinability). Two alternatives $A_0, A_1 \subseteq \Omega$ are recombinable if $\exists n, m \geq 0$ with $F^n(A_0) \cap F^m(A_1) \neq \emptyset$. Otherwise they are non-recombinable.

Definition A4 (Graph-compatible dynamics). Let $G = (V, E)$ with $V \cong \Omega$. The dynamics F is G -compatible if $(\omega, F(\omega)) \in E$ for all $\omega \in \Omega$.

Definition A5 (Permitted local modification). A k -local modification changes G only by adding/removing edges with both endpoints in $B_k(x)$, leaving $V \setminus B_k(x)$ and all edges incident to it unchanged. F remains fixed and G -compatible.

Definition A6 (Outside-B witness). A recombination is witnessed outside B if $\exists n, m \geq 0$ and $w \in V \setminus B$ with $w \in F^n(A_0) \cap F^m(A_1)$.

Definition A7 (Local separability). Local separability holds if a permitted local modification can render previously recombinable alternatives (with support not contained in $B_k(x)$) non-recombinable.

Lemma A8 (Two boundary crossings). In a tree T , any path between vertices outside a ball B either avoids B entirely or crosses ∂B exactly twice.

Lemma A9 (Boundary-path rigidity). In a tree, connectivity between boundary vertices through $V \setminus B$ is unchanged by permitted local modifications.

Lemma A10 (Environment register). If F is bijective but b reduces distinguishability (merges alternatives), then there exists an auxiliary register E and bijection $U: \Omega \times E \rightarrow \Omega \times E$ storing the missing information.

Theorem A11 (Minimal entropy cost). The minimal $|E|$ needed to store one discarded binary alternative is 2, giving $\Delta S \geq k_B \ln 2$.

Theorem A12 (Conditional topological threshold). If irreversibility is locally constructible from G -compatible reversible dynamics, then $\beta_1(G) \geq 1$. Trees ($\beta_1 = 0$) admit no local traps for alternatives with external support.

Appendix B: Notation Summary

Symbol	Meaning
Ω	Set of microstates
$F: \Omega \rightarrow \Omega$	Reversible (bijective) tick dynamics
$b: \Omega \rightarrow \{0, 1\}$	Coarse-graining map defining a bit
A_0, A_1	Alternative trajectory sets
$F^n(A)$	n -fold application of F to set A
$G = (V, E)$	Adjacency graph of microstates
G -compatible	F maps each state to a graph neighbour
$B_k(x)$	Ball of radius k around vertex x
∂B	Boundary of ball B (vertices in B adjacent to $V \setminus B$)
$V \setminus B$	Complement of B in V
$P_T(u, v)$	Unique simple path from u to v in tree T
$\beta_1(G)$	First Betti number = rank $H_1(G)$
$H_1(G)$	First homology group of G
$\pi_1(G)$	Fundamental group of G
E	Environment/auxiliary register
$U: \Omega \times E \rightarrow \Omega \times E$	Joint reversible evolution

Symbol	Meaning
ΔH	Change in Shannon information (nats)
ΔS	Change in entropy
k_B	Boltzmann constant

Appendix C: Glossary for General Readers

Betti number (β_1): A count of independent loops in a network. $\beta_1 = 0$ means no loops (like a tree); $\beta_1 \geq 1$ means at least one loop exists.

Bijection: A function that matches inputs to outputs one-to-one, with nothing left over and no duplicates. Reversible processes are bijections because you can always undo them.

Bit: A binary distinction—something that can be in one of two states (0 or 1, yes or no, left or right).

Coarse-graining: Looking at a system at low resolution, ignoring microscopic details. Like describing a gas by its temperature rather than tracking every molecule.

Decoherence: The process by which quantum superpositions become classical-looking mixtures through interaction with an environment. In this framework, decoherence corresponds to trap formation.

Entropy: A measure of missing information—how much you don't know about a system's microscopic state given its macroscopic description.

Fundamental group (π_1): A mathematical structure capturing the different types of loops in a space. Spaces with nontrivial fundamental groups have loops that can't be shrunk to a point.

G-compatible dynamics: Evolution where each step moves to a neighbouring state in the adjacency graph. Ensures the graph topology actually governs the dynamics.

Homology (H_1): A way of counting holes and loops in a space using algebra. The first homology group H_1 captures loop-like structures.

Irreversibility: The property of a process that cannot be undone by local operations. Broken eggs, forgotten memories, and fixed bits are all irreversible.

Jordan curve theorem: A mathematical theorem stating that any simple closed curve in a plane divides the plane into exactly two regions—inside and outside.

Landauer's principle: The physical law stating that erasing one bit of information requires dissipating at least $k_B T \ln 2$ of energy as heat.

Local vs. global: A local change affects only a bounded region; a global change affects the entire system. Local trapping means creating isolation without restructuring everything.

Local separability: The ability to render alternatives non-recombinable through modifications confined to a bounded region.

Microstate: The complete, maximally detailed description of a system. Knowing the microstate means knowing everything there is to know.

Non-recombinable: Alternatives whose future trajectories can never intersect, no matter how long the system evolves.

Outside-B witness: A recombination event where the meeting point lies outside a specified ball B. Used to prove that local modifications cannot destroy recombinability on trees.

Pointer basis: In decoherence theory, the set of states that remain stable under environmental monitoring. Corresponds to coarse-graining aligned with reachability classes.

Tick: A reversible fluctuation—a change that can be undone without any trace. The fundamental "atoms" of change in the tick-bit framework.

Topology: The mathematics of connectivity and shape, ignoring distances and angles. Topologically, a coffee cup equals a donut (both have one hole).

Trap: A region of state space that things can enter but not leave under normal dynamics. The mechanism by which bits become fixed.

Tree: A connected graph with no cycles. Equivalently, a graph where there is exactly one path between any two vertices.

References

Foundational Works

1. Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183–191.
2. Bennett, C. H. (1973). Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6), 525–532.
3. Bennett, C. H. (1982). The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21(12), 905–940.
4. Fredkin, E., & Toffoli, T. (1982). Conservative logic. *International Journal of Theoretical Physics*, 21(3–4), 219–253.
5. Zurek, W. H. (1989). Thermodynamic cost of computation, algorithmic complexity and the information metric. *Nature*, 341(6238), 119–124.

Topology and Geometry

6. Jordan, C. (1887). *Cours d'analyse de l'École Polytechnique*. Gauthier-Villars.
7. Hatcher, A. (2002). *Algebraic Topology*. Cambridge University Press.
8. Munkres, J. R. (2000). *Topology* (2nd ed.). Prentice Hall.
9. Bollobás, B. (1998). *Modern Graph Theory*. Springer.
10. Diestel, R. (2017). *Graph Theory* (5th ed.). Springer.
11. Spanier, E. H. (1981). *Algebraic Topology*. Springer.

Information Theory

12. Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423.
13. Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). Wiley-Interscience.
14. MacKay, D. J. C. (2003). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press.

Physics of Information

15. Szilard, L. (1929). Über die Entropieverminderung in einem thermodynamischen System bei Eingriffen intelligenter Wesen. *Zeitschrift für Physik*, 53(11–12), 840–856.
16. Brillouin, L. (1951). Maxwell's demon cannot operate: Information and entropy. *Journal of Applied Physics*, 22(3), 334–337.
17. Penrose, O. (1970). *Foundations of Statistical Mechanics*. Pergamon Press.
18. Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, 106(4), 620–630.
19. Bekenstein, J. D. (1973). Black holes and entropy. *Physical Review D*, 7(8), 2333–2346.
20. Lloyd, S. (2000). Ultimate physical limits to computation. *Nature*, 406(6799), 1047–1054.

Quantum Decoherence and Measurement

21. Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75(3), 715–775.
22. Schlosshauer, M. (2007). *Decoherence and the Quantum-to-Classical Transition*. Springer.
23. Joos, E., Zeh, H. D., Kiefer, C., Giulini, D., Kupsch, J., & Stamatescu, I.-O. (2003). *Decoherence and the Appearance of a Classical World in Quantum Theory* (2nd ed.). Springer.
24. Zurek, W. H. (1981). Pointer basis of quantum apparatus: Into what mixture does the wave packet collapse? *Physical Review D*, 24(6), 1516–1525.

Arrow of Time

25. Price, H. (1996). *Time's Arrow and Archimedes' Point*. Oxford University Press.

26. Zeh, H. D. (2007). *The Physical Basis of the Direction of Time* (5th ed.). Springer.
27. Albert, D. Z. (2000). *Time and Chance*. Harvard University Press.
28. Carroll, S. (2010). *From Eternity to Here: The Quest for the Ultimate Theory of Time*. Dutton.
29. Reichenbach, H. (1956). *The Direction of Time*. University of California Press.

Reversible Computation

30. Toffoli, T. (1980). Reversible computing. In J. W. de Bakker & J. van Leeuwen (Eds.), *Automata, Languages and Programming* (pp. 632–644). Springer.
31. Frank, M. P. (2005). Introduction to reversible computing: Motivation, progress, and challenges. In *Proceedings of the 2nd Conference on Computing Frontiers* (pp. 385–390). ACM.
32. Perumalla, K. S. (2014). *Introduction to Reversible Computing*. CRC Press.

Philosophy of Physics

33. Earman, J. (2006). The "past hypothesis": Not even false. *Studies in History and Philosophy of Modern Physics*, 37(3), 399–430.
34. Callender, C. (2004). There is no puzzle about the low-entropy past. In C. Hitchcock (Ed.), *Contemporary Debates in Philosophy of Science* (pp. 240–255). Blackwell.
35. Rovelli, C. (2018). *The Order of Time*. Riverhead Books.
36. Barbour, J. (1999). *The End of Time: The Next Revolution in Physics*. Oxford University Press.

Experimental Verifications of Landauer's Principle

37. Bérut, A., Arakelyan, A., Petrosyan, A., Ciliberto, S., Dillenschneider, R., & Lutz, E. (2012). Experimental verification of Landauer's principle linking information and thermodynamics. *Nature*, 483(7388), 187–189.
38. Jun, Y., Gavrilov, M., & Bechhoefer, J. (2014). High-precision test of Landauer's principle in a feedback trap. *Physical Review Letters*, 113(19), 190601.
39. Hong, J., Lambson, B., Dhuey, S., & Bokor, J. (2016). Experimental test of Landauer's principle in single-bit operations on nanomagnetic memory bits. *Science Advances*, 2(3), e1501492.

Related Theoretical Frameworks

40. Wheeler, J. A. (1990). Information, physics, quantum: The search for links. In W. H. Zurek (Ed.), *Complexity, Entropy and the Physics of Information* (pp. 3–28). Addison-Wesley.
41. 't Hooft, G. (2016). *The Cellular Automaton Interpretation of Quantum Mechanics*. Springer.
42. Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media.
43. Verlinde, E. (2011). On the origin of gravity and the laws of Newton. *Journal of High Energy Physics*, 2011(4), 29.

44. Jacobson, T. (1995). Thermodynamics of spacetime: The Einstein equation of state. *Physical Review Letters*, 75(7), 1260–1263.
45. Nayak, C., Simon, S. H., Stern, A., Freedman, M., & Das Sarma, S. (2008). Non-Abelian anyons and topological quantum computation. *Reviews of Modern Physics*, 80(3), 1083–1159.