Irreducible Reality

Emergent Spacetime from Information-Theoretic Constraints

Christian Halldén

Abstract

In this framework, emergent spacetime, gravity, and the dark sector are built from one keystone, the monotonicity of relative entropy under completely positive maps (Lindblad 1975, Uhlmann 1977; equivalent to strong subadditivity, Lieb-Ruskai 1973), a small set of structural axioms each with a stated cost, and one conjectured parameter. The fundamental object is the thread graph, a network of interaction points on qubits whose shared mutual information defines a metric via geodesic closure. Einstein's field equations follow Jacobson's entanglement-first-law derivation on causal diamonds (1995), with three additions: a concrete graph-to-geometry map via mutual information, the requirement of non-Clifford resource for gravitational backreaction, and an iterative feedback loop that makes the derivation dynamical. Persistent thread tangles, the identities that carry mass, exist only in three dimensions, and the automaton's effective dimension settles low (numerically 2-3), not yet uniquely pinned to three. The same entanglement tensor decomposes into a dark sector: the trace component is dark energy, residual strain from the original bifurcation, fixed to w = -1 by volume-extensivity of the residue; the traceless component is dark matter, directional structure correlated with matter gradients (r = 0.556, preliminary). The tensor coupling at nuclear scales predicts a rank-2, m_l-dependent fifth force (neutron-number and Z^3/n^2 scaling) whose shape, testable via King-plot nonlinearity in isotope shifts, discriminates it from scalar mediators; its magnitude is set by the substrate scale and bounded by current measurements. That parameter is the dilution factor η (conjectured 1-1/e, measured 0.63-0.70). The physics results, gravity, cosmology, and the dark sector, are independent of the consciousness claims, developed in a companion paper.

1. Introduction

1.1 The Master Inequality

A single mathematical fact runs through every derivation that follows: distinguishability cannot increase under physical operations. This is a theorem about information, proved decades ago and never contradicted.

Evaluated at a causal horizon, it becomes gravity. Over cosmological states, it constrains dark energy. Across a network of interacting points, it produces the metric and the arrow of time. Inside a recursively self-referential system, it draws the boundary between self and world, and sets the cost of looking inward.

At the moment an interaction creates mutual information, it sets the conditions for measurement (envariance, Zurek 2003).

The master inequality

$$D(\rho\|\sigma) \geq D(\Phi(\rho)\|\Phi(\sigma)) \quad \text{for any CPTP map } \Phi \tag{1}$$

Proved by Lindblad (1975) and Uhlmann (1977). The monotonicity of relative entropy under completely positive trace-preserving maps.

Set Φ to partial trace: strong subadditivity of von Neumann entropy. From this follow the metric's triangle inequality, the Jacobson derivation of Einstein's equations from entanglement, and the monogamy of squashed entanglement.

Set Φ to the trace map: $D(\rho\|\sigma) \geq 0$ for all states. In equilibrium, the dark-energy value $w = -1$ follows from volume-extensivity of the residue (constant density); off equilibrium the driven dynamics predict a single phantom crossing, consistent with the averaged null energy condition (Faulkner et al. 2016) while locally below the pointwise condition in the phantom phase. However, the framework's spacetime is not an equilibrium structure, it is a **driven dissipative structure** (§7, §11.9) sustained by ongoing genesis against constant dilution. In this driven regime the effective dark-energy equation of state executes a single crossing of the phantom line from below: $w < -1$ in the past (the strain lies below its receding attractor and is rising) and $w > -1$ today ($w_0 > -1$, the strain tracks the falling attractor). The framework thus *predicts* phantom crossing of the Quintom-B type, matching the sign of the current DESI preference (~4σ). Should future surveys falsify phantom crossing, the prediction reverts to the equilibrium bound $w = -1$ (cosmological constant) via the standard ANEC derivation.

Data processing: mutual information cannot increase under local operations. From this follow dilution of pairwise $\mathcal{I}$ between distant subsystems, the Markov blanket self of the companion paper, and the reflexive bound.

One object carries the whole account. The entanglement tensor splits three ways, and each split names a different domain. By symmetry it gives the dark sector, dark energy in the isotropic part and dark matter in the anisotropic part. By computational class, stabilizer against magic, it gives the difference between rigid space and responsive gravity. By partition, a region against its surroundings, it gives the self. One object, three readings: cosmology, dynamics, mind.

Full Derivation: The Keystone

The keystone is the master inequality of §1.1, equation (1): monotonicity of quantum relative entropy under CPTP maps, equivalently strong subadditivity of von Neumann entropy (Lieb-Ruskai 1973; Lindblad 1975; Uhlmann 1977). The chain below derives it from strong subadditivity and shows why every load-bearing result downstream depends on it: the metric's triangle inequality, the fixed-point self, the dilution of distant correlation, the reflexive bound, the monogamy budget, the entanglement first law behind gravity, and the dark-sector energy condition.

Derivation chain

Step 1. Strong subadditivity (Lieb-Ruskai 1973): for any tripartite state,

$$S(ABC) + S(B) \leq S(AB) + S(BC)$$

Step 2. SSA is equivalent to monotonicity of relative entropy under partial trace, and via Stinespring dilation under every physical (CPTP) map:

$$D(\Lambda(\rho)||\Lambda(\sigma)) \leq D(\rho||\sigma)$$

Distinguishability never increases under processing.

Step 3. Everything downstream cites this one inequality: dilution of pairwise I between distant subsystems (monotonicity along the interaction chain), the Lyapunov property of the gravitational fixed point (relative entropy to vacuum non-increasing), the monogamy budget, and the reflexive bound (each modeling layer is a channel). One theorem, one master constraint.

1.2 Five Propositions

Five claims ground the account. Everything that follows is an attempt to derive the rest of physics from them, using the simplest logic from the simplest starting point.

i. Space is geometric relation, information about how interaction points relate to each other from a given perspective. Space extends only as far as a network's relationships reach.

ii. Time is causal relation. Without interaction there is no measure of time, without time there is no interaction. The two are inseparable. An isolated point is timeless.

iii. Physics is subjective experience, what emerges when space and time are experienced together from a specific point. Geometry and causality, rendered from within.

iv. The observer is fundamental. Every interaction implies a perspective from which the interaction is distinguishable. The observer is present from the first interaction onward, at every scale. (This follows from the definition of an interaction: difference implies a vantage from which difference exists. The observer is that vantage, inherent in the interaction structure.)

v. Reality is rendered by subjects. There is no view from nowhere. Every perspective is partial, bounded by its causal history. The manifested world is the intersection of all such perspectives, woven into a mutually consistent whole.

Full Derivation: Axioms

A0 (Arena). A separable Hilbert space H with a preferred factorization into subsystems indexed by a countable set V. Cost: the factorization is chosen, not derived. This is the standard open problem of preferred tensor structure, and the framework inherits it openly.

A1 (Genesis). There exists an isometry B̂: H → H ⊗ H (the bifurcation), intertwining every conserved charge Q. Cost: existence of the operator. Its consequences (conservation, anticorrelation, irreversibility) are then proved, not assumed.

A2 (Interaction). Subsystems that interact acquire mutual information I(i,j) = S(ρ_i) + S(ρ_j) - S(ρ_ij) > 0, and under generic subsequent evolution I never returns exactly to zero. Cost: "generic" is a typicality statement (exact disentanglement is measure-zero), not a theorem for all dynamics.

A2′ (Locality of dynamics). The Hamiltonian couples only along the regular, geometrically local edge structure. Anomalous long-range edges are kinematic, carrying state correlation but no Hamiltonian term.

A3 (Perspective). Each subsystem's accessible information is its own interaction history: the algebra generated by edges incident to it. This is the declared premise about observation. Cost: the premise of the whole project.

A4 (Regularity). The regular edge structure has bounded degree and power-law decay of I with network distance. Cost: needed for manifold emergence; not derived from dynamics.

A5 (Magic). The global state is not a stabilizer state: non-Clifford resource is present. Motivation: stabilizer-only dynamics is classically simulable (Gottesman-Knill) and produces a rigid, response-free entanglement pattern; a world with gravity that responds must compute more than Clifford.

A6 (Minimality of threads). Persistent identities are one-dimensional: chains of events linked by identity edges, each event with a unique successor. Cost: minimality choice. Its payoff is Section 8, where it forces the dimension of space.

Derivation chain: A5 as typicality

Step 1. For n qubits the stabilizer states form a finite set while states form a continuum: the stabilizer sector has measure zero. The Clifford group is finite inside the unitary group; a Haar-generic Hamiltonian is non-Clifford.

Step 2. Magic monotones never increase under stabilizer operations (Veitch et al.). A world that started Clifford and evolved Clifford stays Clifford at every later step.

Step 3. Magic is present now (curvature responds; measurement produces records). Therefore either the dynamics is non-Clifford (the generic case) or magic was injected at genesis. A5 is thereby a typicality statement of the same grade as A2: the question inverts from why magic exists to why it would ever be absent.

1.3 Mathematical Specification

Definition 1 (Thread graph). A thread graph $G = (V, E_{\mathrm{id}}, E_{\mathcal{I}})$ consists of a countable vertex set V representing interaction points, identity edges $E_{\mathrm{id}} \subseteq V \times V$ representing persistence (each vertex has at most one outgoing identity edge, by A6), and interaction edges $E_{\mathcal{I}} \subseteq V \times V$ where (i, j) exists iff the subsystems at i and j have interacted, acquiring mutual information

$\mathcal{I}(i,j) > 0$. Each vertex hosts a qubit (binary distinguishability, consistent with Margolus-Levitin and A2). The global state ρ is a density operator on $(\mathbb{C}^2)^{\otimes|V|}$.

Definition 2 (Mutual information). For any subsystems $A, B \subseteq V$, $\mathcal{I}(A:B) = S(\rho_A) + S(\rho_B) - S(\rho_{AB})$ where $S(\rho) = -\mathrm{Tr}(\rho \log \rho)$.

Definition 3 (Metric). Edge length $\ell(i,j) = -\log(\mathcal{I}(i,j)/\mathcal{I}_{\max})$ where $\mathcal{I}_{\max} = 2\log 2$, the maximal mutual information of a qubit pair (a Bell pair saturates it). The geodesic distance $d_{\text{net}}(i,j) = \min_{\text{paths}} \sum \ell$ over sequences of interaction edges defines a metric on V; because it is a shortest-path length the triangle inequality holds by construction, while strong subadditivity (Lieb-Ruskai 1973) supplies the faithfulness property (iv) below.

Derivation: the metric axioms. Three properties make d_{net} a metric, and a fourth makes it faithful. (i) Non-negativity: $0 < \mathcal{I}(i,j) \le \mathcal{I}_{\max}$ gives $\ell(i,j) \ge 0$, vanishing only at maximal correlation, where the two points are identified. (ii) Symmetry: $\mathcal{I}(i,j) = \mathcal{I}(j,i)$ makes ℓ symmetric. (iii) Triangle inequality: d_{net} is a shortest-path length, so the concatenation of a path $i \to j$ with a path $j \to k$ is itself a path $i \to k$, and the minimum over all $i \to k$ paths cannot exceed that concatenation,

$$d_{\text{net}}(i,k) \;\le\; d_{\text{net}}(i,j) + d_{\text{net}}(j,k).$$

(iv) Faithfulness from monotonicity: along a Markov chain $i - j - k$, data processing (the master inequality (1) with Φ the partial trace) gives $\mathcal{I}(i:k) \le \mathcal{I}(i:j)$, hence $\ell(i,k) \ge \ell(i,j)$. Correlations only weaken down a chain, so the edge lengths respect the geodesic ordering and the shortest-path metric reflects the actual correlation structure, not merely the axioms formally.

Definition 4 (Dynamics). Per time step: (1) compute d_{geo} from the current $\mathcal{I}$ distribution via geodesic closure, (2) pairs interact at the rate-limited, balanced throughput the η floor below assumes (about one slot refresh per step), with selection weight $p \propto \exp(-d_{\text{geo}})$ (interaction range $\alpha = 1$ in Planck units, required for fixed-point stability); the source-free version of this rule decays to the product state and a genesis source sustains the live phase (§11.9), (3) interacting pairs apply an entangling unitary (CNOT + Hadamard + T-gate, universal for quantum computation) raising $\mathcal{I}(i,j)$ toward the set-point $\log 2$ (half the qubit-pair maximum $2\log 2$; real edges range up to $2\log 2$, consistent with the measured mean $\langle \mathcal{I} \rangle \approx 1.17$ nats), (4) all $\mathcal{I}$ values dilute by factor $\eta \approx 0.63{-}0.70$ (measured from automaton saturation; conjectured to be $1 - 1/e$, the occupancy floor: when a saturated qubit refreshes its partner-slots by firing uniformly at mean rate one per slot (balanced throughput) and an unrefreshed slot decays within the step, the surviving fraction is $1 - (1 - 1/k)^k \to 1 - 1/e$ in the large-degree limit. This is the floor, not a generic monogamy value (a budget-rescaling shed gives the cap fraction instead); the measured range sits at the floor plus corrections from partial slot memory and geodesic-biased, non-uniform firing), (5) edges with $\mathcal{I} < \varepsilon$ are removed ($\varepsilon \ll 1$). The automaton has one conjectured parameter (η); $\alpha = 1$ follows from fixed-point stability (conjectured), and η is measured, not tuned.

Definition 5 (Continuum limit). Under A4 (bounded degree, power-law decay of $\mathcal{I}$ with distance), coarse-graining d_{net} yields a Riemannian manifold $(\mathcal{M}, g)$ as $|V| \to \infty$, $\varepsilon \to 0$, in the Gromov-Hausdorff sense. The convergence is supported numerically by the automaton analysis in Section 10 (result M1: the dynamics selects a connected low-dimensional manifold, numerically 2-3 and not yet uniquely pinned) but has not been proved analytically. The universality class of the automaton is conjectured but not yet established.

Adopted results. The following theorems from the quantum information and mathematical physics literature are used as building blocks throughout the derivation: Lieb-Ruskai (strong subadditivity, 1973); Lindblad-Uhlmann (monotonicity of relative entropy, 1975/1977); Lieb-Robinson (finite group velocity, 1972); Margolus-Levitin (orthogonalization bound, 1998); Jacobson (entanglement thermodynamics, 1995); Faulkner-Leigh-Parrikar-Wang (ANEC from relative entropy, 2016); Christandl-Winter (monogamy of squashed entanglement, 2004); Knaster-Tarski (fixed points of monotone maps on complete lattices, 1955; the companion paper refines the self-from-fixed-point into a quantitative Markov blanket formulation, see §4); Calugareanu-White-Fuller (Lk = Tw + Wr); Gromov-Hausdorff convergence theory; the unknotting theorem for 1-spheres in R^n. For alternative emergent gravity approaches, the present Framework shares Verlinde (2010) entropic motivation but constructs the metric differently (mutual information geodesics vs holographic screens). Knot-based dimensional arguments appear independently in Berera et al. (2017) for flux-tube inflation; the present application to thread stability is distinct.

---

2. Genesis

2.1 The First Event of This Manifestation

The first event is a bifurcation, producing balanced entity and anti-entity pairs. This pattern is not invented from scratch. Nature already shows it.

Wheeler's quantum foam (1955): virtual particle-antiparticle pairs appearing and vanishing at the Planck scale. Hawking radiation (1974): a virtual pair separated at the event horizon, one falling in, the other escaping as radiation. The Schwinger effect: pair production in strong electric fields, where virtual pairs become real when the field supplies enough energy for separation. The Casimir effect: virtual particles made detectable through boundary conditions, exerting measurable pressure between conducting plates. In each case, something emerges from the vacuum, balanced by conservation, manifested through separation.

Existence is not treated here as something that begins. What is derived is manifestation: the form this world's first event takes, and what follows from it. Why there is a process rather than none is set aside, since nothing is not a state a process can occupy or leave, and positing it adds nothing. The results below do not depend on any account of an absolute beginning.

Interpretation and derivation. The narrative that follows, bifurcation, irreducibility, the necessity of an inside, is metaphysical interpretation, not derived from the axioms but the layer that motivates them. The mathematics answers "given the axioms, what unfolds" (how). The narrative answers "why these axioms, and why must the process be lived from within" (why). The two are kept distinct throughout: the derivations carry their own status (proven, conditional, conjectured), and the interpretive claims are marked as interpretation. This demarcation follows the standard practice in the foundations of quantum mechanics, where formalism and interpretation (Copenhagen, many-worlds, relational QM, QBism) are held separate.

Situating the interpretation. This "why" joins a lineage of serious attempts to answer the question of existence. Vilenkin's tunneling-from-nothing (1982) and the Hartle-Hawking no-boundary proposal derive the universe from the instability of nothing. Wheeler's participatory universe ("it from bit") places the observer as constitutive: the universe requires an observer to bring it into being. Wolfram's computational irreducibility identifies processes that cannot be shortcut, only run, the same irreducibility that here forces an inside. Tegmark's Mathematical Universe Hypothesis holds that all mathematical structures exist; we inhabit one of them. The distinguishing move is the unification: the same irreducibility that makes an unfolding real rather than readable from outside (Wolfram) is what forces an inside, a vantage (Wheeler's observer); the question of existence (Vilenkin, Hartle-Hawking) is set aside rather than answered by an unstable nothing. A vantage is the bare inside the process carries from the first split; it is not consciousness: an electron has a vantage and no experience. Consciousness is what a vantage becomes only along a climb: a recurrent structure realized in a substrate, organic for organic life, organizes perception into sentience, a system with an internal state reacting adaptively to its surroundings; an internal recursion then folds that structure onto a self-model, from which consciousness and, later, self-awareness emerge. The climb is derived in the companion paper; the bare inside is where it starts, not where it ends. This account ties the inside and the irreducibility of the process to a single root, the irreducibility of the process that began with the bifurcation. This is a philosophical claim, not a derived result.

The first something is a **bifurcation event**. For each virtual something, a virtual anti-something appears.

Bifurcation fits because manifested reality obeys the constraints of conservation. A conserving split is how a distinction appears at zero net cost: the two sides sum to nothing, so the manifestation proceeds without violating any conserved charge. The sum remains balanced. In this manifestation, two entities form from a zero total, balanced so their charges sum to nothing. These entities differ only in relative identity, the information carrying their differentiation from sameness.

Separating the pair lowers the energy relative to the symmetric state, and that differential drives the manifestation forward. What follows derives the rest from this single transformation.

A process that is reducible, whose every state can be known from its starting conditions, is already complete as a logical structure. It has no need to run, no need to take time, no need to occupy space. It exists the way a theorem exists: timelessly, for free. An irreducible process cannot be complete in this way. Its states cannot be computed from the outside in advance. They must be produced, step by step, and that production is what we experience as time, as space, as the unfolding present. The process that may have been set in motion by the bifurcation is of this second kind. Its interacting points had limited perspectives, each seeing only its own interaction history, no point holding the full picture, making the process quantum-irreducible. Operationally, a process is irreducible when its state trajectory cannot be computed from initial conditions by any resource-bounded procedure shorter than running the process itself (Bennett's logical depth). In quantum information terms, this is the distinction between stabilizer states (classically simulable via Gottesman-Knill) and non-Clifford states (which require exponential classical resources or quantum computation). The thread network with non-vanishing magic (A5) falls in the irreducible class: its dynamics cannot be compressed. This connects the philosophical notion of irreducibility to a well-defined computational criterion.

Irreducibility alone does not guarantee continuation. The stronger claim is about existence: an irreducible process cannot be complete from outside, so its states do not exist until they are produced. It requires manifestation, or it cannot exist, even hypothetically, except as a phenomenon, never as a construction.

An objection from the Mathematical Universe Hypothesis. A Tegmark-style reading of the Mathematical Universe Hypothesis resists this conclusion: if a mathematical structure exists timelessly and for free, then an irreducible process, still just a structure, could likewise exist as an *uninstantiated abstract object*, with no need to be run. (This is a position the MUH licenses, not a formal objection Tegmark has pressed against the account.) The answer turns on the structure itself. The process is not an arbitrary mathematical object but a self-modeling region, a quantum Markov blanket in which the boundary renders interior and exterior conditionally independent, $\mathcal{I}(A : C|B) \to 0$, saturating the strong subadditivity axiom. The companion paper formalizes this Markov-blanket self and grounds it in Bennett's logical depth: the self-model's state has no compressed shortcut and exists only by being run, step by step. The irreducibility of the process, that it *must be run*, is carried by the non-Clifford (magic) resource required for responsive geometry, not by a formal diagonal argument. There is then no external standpoint from which the "uninstantiated" version could be written down or pointed to, because the only complete specification of a self-referential Markov blanket is the blanket's own running. For an irreducible self-modeling process, manifestation is not an optional extra over abstract existence; abstract-from-outside existence is precisely what it cannot have. This does not prove the process must continue; it answers why, if it exists at all, it exists only from the inside.

Our manifested reality may itself be the product of a deeper process, a prior layer from which ours arose. The attention and starting point of this framework is the manifested reality that we are aware of.

Recursive cosmogenesis. This acknowledgment has a natural completion the operators already permit. If $\hat{B}$ can fire not only once at the origin but locally wherever the network saturates (Section 6.2: a black-hole horizon is the saturation limit of shortcut density, and the pair production there carries $\hat{B}$'s signature, doubling, conservation through anti-correlation, and separation), then each saturated interior is a fresh start of the genesis cascade $\hat{B}^n$, causally sealed from the net enclosing it because proper time stops at the horizon (eq. 4). Our manifested reality is then one net among such genesis events: a prior layer enclosing it, a successor layer within every black hole it holds. The structure is a one-way succession, not a cycle, and the distinction is forced by Theorem 2.2: $\hat{B}$ has no global inverse, so no successor can return to the net it formed within. A cyclic cosmology would require an invertible genesis; this genesis is irreversible, so the succession of sealed nets runs one way through saturation boundaries and cannot close into a loop. The width of the succession is budgeted by monogamy (companion paper): an enclosing net of total entanglement S seals at most $\sim 2S/q_{\text{succ}}$ full-budget successors. Whether each sealed interior forms a self-sustaining successor, and whether the succession continues without end, are taken up in Section 6.2.

The bifurcation is a partial isometry operator that preserves conservation laws, doubling the Hilbert space. For any conserved quantity, conservation requires anti-correlation between the two branches.

Because the bifurcation is irreversible by simple recombination, the process is not reversible, which is why manifested entities persist rather than immediately annihilating. The energy drops: the split is energetically favored.

Irreversibility. $\hat{B}$ is an isometry, not a unitary: $\hat{B}\hat{B}^\dagger = P$, the projector onto the image subspace. The reverse operation does not exist globally. Under generic local evolution, the state immediately leaves the image. The bifurcation closes a door that cannot be reopened. The network persists because the path back, for almost all states, is not defined.

Full Derivation: Genesis

Theorem 3.1 (Two from a zero total). For any conserved Q with ⟨Q⟩ = 0 before bifurcation, the image state satisfies ⟨Q_A + Q_B⟩ = 0, and for definite initial Q = 0, outcomes in the two factors are exactly anticorrelated: Q_A = -Q_B on the image.

Derivation chain

Theorem 2.1 (the blank). Zero distinctions specifies exactly one configuration: the blank a vantage reads where no distinction is accessible. This is the condition at the first event and, from the other side, inside a saturated interior (Section 6.2).

Theorem 2.2 (no inverse), the door. The bifurcation creates I > 0. Reversing it requires a sequence of local operations driving every pairwise I exactly to zero simultaneously; for Haar-typical states the target is measure zero and monotonicity blocks local erasure of shared correlations. The path back is not forbidden by a force; for almost all states it is not defined.

The bifurcation is an isometry, not a unitary. A unitary has an inverse; an isometry does not. This mathematical distinction has a physical consequence: the first interaction is the first event. Since time is defined as the sequence of distinguishable events along identity edges, and the bifurcation is the first distinguishable event, there is no "before." The question "what happened before the Big Bang?" is not unanswered, it is ill-posed: "before" presupposes a prior distinguishable event, and the bifurcation is the first. The arrow of time is not an emergent statistical tendency; it is built into the structure of the first moment.

2.2 Identity and the Observer

Bifurcation produced identity pairs. Exclusion set them in motion. What follows is the logical cascade from identity to the first observer.

Through bifurcation, balanced pairs form informational entities. The property of difference gives identity. Two opposites that realign into coherence return to nothingness; they no longer retain identity. Identity, in this framework, is the simple fact of being distinguishable from something else.

In isolation, an entity experiences neither time nor space, a point with no relations, no interactions, no way to distinguish anything. The entity remains virtual.

When two distinct entities meet, they occupy different positions by definition. If they occupied the same position, they would not be two entities. Two of the same occupy different positions in what may be called *identity-space*, not yet physical space, but the logical condition that two distinguishable things are distinguishable *somewhere*. This "where" is what the framework calls meta-identity: an entity's position relative to other entities, the seed from which space will grow.

When two entities exclude each other, that exclusion is the first interaction. And the first interaction carries something new.

Difference implies a perspective from which difference exists. There is no asserting that two things differ without a vantage that registers the difference. At the first exclusion, something distinguished itself from something else. That distinction is the first observation.

The first exclusion interaction is the first observer. From this point, the observer is present at every subsequent interaction. The observer does not emerge later, at sufficient complexity. The observer is present from the beginning.

The observer as premise

The observer is adopted as fundamental, a premise the model is built on, not a conclusion derived from it. What follows explores what becomes possible if the premise holds: which connections it unlocks, which phenomena it ties together.

Given the premise, a single-celled organism is an observer. A bacterium navigating a chemical gradient has a perspective, however minimal. Every thread interaction point is an observation point. The question of how these points integrate into richer forms of experience is taken up in the companion work on structured perspective.

---

3. The Network

3.1 Threads

Identity has been established, and the observer has been introduced. The question now is how identity persists across multiple interactions, and what that persistence makes possible.

One interaction is a point. But identities can interact more than once.

When the same identity participates in a second interaction, and a third, the accumulated sequence forms a causal chain of events. This chain is a **thread**.

A thread is identity persistence. The identity A that interacted at event 1 is the same identity A that interacts at event 7, and again at event 12. The thread is not made of anything: it carries no substance or additional information, only a relation: *same identity, across multiple interactions*.

Why one-dimensional

Why threads, and not sheets or volumes? A point suffices for one interaction. A line is the simplest extension that allows an identity to participate in *multiple* interactions in sequence. A sheet would add a dimension with no clear functional role: there is only a sequence of events to connect. A line is the minimal extension that does the work.

This differs from the motivation for strings in string theory, where 1D objects replace 0D particles to resolve ultraviolet divergences. Threads address a different question: what allows an identity to accumulate a causal history across multiple events.

The Thread Network

As threads persist and interact, they form a network. Each thread is a node. Each interaction between threads creates a connection. The result is a graph: a set of nodes linked by edges. Two kinds of edges emerge naturally from the two ways threads relate to each other.

In this network there are two kinds of edges:

Identity edges. Connect successive interaction points of the *same* identity. These are threads. They are the persistence that allows causal chains to form. **Interaction edges.** Connect *different* identities that have interacted. They carry mutual information. Networks of Interaction Edges give rise to geometric relations.

The two edge types together form the full graph. Identity edges supply the persistence. Interaction edges supply the information relationships. The two support each other, and from their combined structure, both time and space begin to take shape, still virtual, but with the architecture that will support physical manifestation.

Meta-identities, the positions entities occupy relative to each other through exclusion, are the beginnings of space. Causal chains, the sequences of interactions along Identity Edges, are the beginnings of time.

Given these three elements, identity, interaction, and persistence, the rest of the framework unfolds. Time is the causal chain built from Identity Edges. Space is the geometric relation built from Interaction Edges. What follows traces these consequences.

If one were to build a universe from these assumptions, three elements would suffice:

- **Identity.** Bifurcation produces distinguishable entities with persistence.
- **Interaction.** Exclusion forces identities to occupy distinct states, producing distinguishability.
- **Persistence.** Identity survives between interactions. This is the thread.

Interaction Persistence

When two identities interact, the interaction does not vanish. It leaves a trace.

Two kinds of edges have been introduced: Identity Edges (threads, connecting the same identity across interactions) and Interaction Edges (connecting different identities that have interacted). The focus now is on what Interaction Edges carry.

When identity A interacts with identity B, their mutual information becomes nonzero, the formal measure of how much the state of A reveals about the state of B, and vice versa. After the interaction, A and B may go their separate ways. They may interact with C, D, and dozens of others. But that shared information does not return to zero. The interaction has left a permanent record in the graph.

This is what quantum mechanics calls entanglement, in these terms. Two entities that have interacted share a connection that persists regardless of what else happens. The connection carries no force between them, no signal, simply the logical residue of having interacted: *these two identities once crossed paths, and the graph remembers*.

Whether this connection appears local or distant is a question that belongs to space, and space has not yet emerged. At this stage in the derivation, the point is simpler: interactions are not ephemeral. They accumulate. The graph grows denser with every event. And that accumulating density is what will eventually give rise to the continuous structure experienced as physical reality.

Interacting Beyond the Speed of Light

The mutual information does not travel through the emerged metric, because the metric has not yet emerged. It exists as an edge in the underlying graph. When quantum mechanics observes that entangled particles appear to communicate faster than light, the framework sees no paradox. The correlation was never transmitted, only embedded in the graph from the moment of interaction. Two observers measuring A and B are not exchanging a signal. They are reading the same edge from two different vantage points. The speed of light, c, limits what can propagate through the emerged manifold. It does not limit what is already woven into its structure.

From Graph to Manifold

The threads and Interaction Edges together form a graph. As this graph grows denser, the discrete begins to look continuous.

A sparse net shows individual threads and knots, each knot a distinct point, each thread a distinct connection. As threads and knots multiply, at some density the individual knots no longer resolve, and the net reads as a surface rather than as points and lines.

The surface is not a new substance, just the same net seen from a distance where the discrete structure blurs into continuity. This is how space emerges. The graph of Identity Edges and Interaction Edges, at sufficient density, approximates a continuous manifold.

Mathematically, this transition is captured by Gromov-Hausdorff convergence: a discrete graph, in the limit of many vertices and edges, converges to a smooth Riemannian manifold. The measure of distance on that manifold is derived directly from mutual information. What is experienced as the geometry of space is the large-scale structure of the interaction graph.

At the scale of individual interactions, the graph is granular. This granularity corresponds to what physics calls the Planck scale. At human scales, the granularity is invisible. This transition is not invented; it is inherited from the mathematics of dense graphs.

From graph to manifold

The construction begins with a separable Hilbert space over a countable index set. Information relationships are quantified by quantum mutual information.

Define edge lengths from mutual information:

$$\ell(i,j) = -\log\left(\frac{\mathcal{I}(i,j)}{\mathcal{I}_{\max}}\right) \tag{2}$$

Network distance is the geodesic closure:

$$d_{\text{net}}(i,j) = \min_{\text{paths}} \sum \ell \tag{3}$$

d_{net} is a metric: the triangle inequality holds for any shortest-path distance. Under uniform regularity, bounded curvature, and power-law decay, the graph converges in the Gromov-Hausdorff sense to a smooth Riemannian manifold.

Two scales, one graph. The regular component becomes spacetime through coarse-graining. Sparse anomalous edges (direct low-ℓ connections between manifold-distant points) are invisible to the coarse-grained metric: $d_{\text{net}} \le d_{\text{geo}}$, with equality except across shortcuts. These shortcuts correspond to microscopic Einstein-Rosen bridges (ER = EPR).

The metric's faithfulness, that edge lengths respect the actual correlation ordering, follows from strong subadditivity, the master inequality under partial trace; its triangle inequality is automatic from the shortest-path construction.

The Lorentzian signature (one time, three space) is similarly selected by stability. For a manifold with more than one time direction, the dispersion relation loses positivity; only a single time direction permits a positive energy spectrum with a stable ground state, allowing coherent propagation of information through causal chains.

Full Derivation: The Graph

Threads are the identity chains of A6. The network G = (V, E_id, E_I) carries I on interaction edges. Two lemmas used everywhere:

Lemma 4.1 (Monotone growth of arguments). I(A : BC) ≥ I(A : B).

Lemma 4.2 (Local processing never helps). Any CPTP map applied to B alone cannot increase I(A : B).

Entanglement, in this language, is the memory of the graph (A2): interaction writes I, and Lemma 4.2 says forgetting requires contact, not distance.

The One Object and the Three Decompositions

Local entanglement density and the entanglement tensor:

$$\rho_E(i) = \sum_j \mathcal{I}(i,j)$$
$$\mathcal{E}_{\mu\nu}(x) = \sum_y \mathcal{I}(x,y)\, n_\mu(x,y)\, n_\nu(x,y)$$

Everything here is the tensor E_μν examined under three readings.

$$\begin{aligned}
\text{Algebraic (symmetry):}\quad & \mathcal{E} = \mathcal{E}^{(\text{trace})} + \mathcal{E}^{(\text{traceless})} && \Rightarrow\ \text{cosmology}\\
\text{Computational (class):}\quad & \mathcal{E} = \mathcal{E}^{(\text{stab})} + \mathcal{E}^{(\text{magic})} && \Rightarrow\ \text{dynamics}\\
\text{Partitional (cut):}\quad & \mathcal{E} = \mathcal{E}_{SS} + \mathcal{E}_{S\bar{S}} + \mathcal{E}_{\bar{S}\bar{S}} && \Rightarrow\ \text{mind}
\end{aligned}$$

Each split is well-defined on its own terms: the algebraic one is pointwise tensor algebra, the computational one is a labeling of the stabilizer versus non-Clifford (magic) contribution to the correlations, a resource reading rather than a canonical additive split, the partitional one is a block decomposition under a choice of vertex set S.

One object, three readings. By symmetry it yields cosmology, by computational class it yields dynamics, by partition it yields mind.

Derivation chain

Step 1. Vertices are interaction points; an edge exists where I(i,j) > 0. Edge weight is the mutual information itself.

Step 2. Along a processing chain i → j → k the data-processing inequality bounds the end-to-end correlation by every intermediate link: correlations compose by attenuation, never amplification. This is what makes path structure meaningful.

Step 3. Attenuation along independent links is multiplicative in I, hence additive in -ln I. Distance is therefore logarithmic in shared information, and shortest paths are the channels along which the least distinguishability is lost.

Full Derivation: The Metric

Construction

Define edge lengths and take the geodesic closure:

$$\ell(i,j) = -\log\bigl(\mathcal{I}(i,j)/\mathcal{I}_{\max}\bigr)$$
$$d_{\text{net}}(i,j) = \min_{\text{paths}} \sum \ell$$

Theorem 5.1. d_net is a metric: symmetry and nonnegativity are immediate; d_net(i,i) = 0; the triangle inequality holds for any shortest-path distance by construction.

High shared information is proximity, with the triangle inequality for free.

The manifold

Under A4, coarse-graining d_net yields a smooth Riemannian limit in the Gromov-Hausdorff sense; the conditions (degree bound, decay law, curvature bounds) are assumptions.

Two scales, one graph

$$\begin{aligned} d_{\text{geo}}(i,j) &= \text{distance in the emergent manifold (regular component)} \\ d_{\text{net}}(i,j) &= \text{distance in the full graph (includes anomalous edges)} \\ d_{\text{net}} \le d_{\text{geo}}, &\quad \text{equality except across sparse shortcuts} \end{aligned}$$

A vanishing density of shortcut edges leaves coarse geometry intact while collapsing selected path lengths (the small-world phenomenon). Shortcuts are the framework's non-locality: pairs with a direct low-l edge that the manifold cannot represent. By A2′ they carry correlation and no dynamics, which is the ER = EPR picture of an entangled pair as a non-traversable micro-bridge.

A black hole horizon is the saturation limit of shortcut density.

3.2 From Virtual to Physical

Logic precedes manifestation. 1+1=2 holds regardless of substrate, whether this is base reality or a simulation layer. As entities interact, they trace logical pathways, manifesting what conforms to logical possibility. Reality does not invent logical rules; it adheres to them through unfolding structure.

The shift: what transforms virtual possibility into physical reality is **subjective interaction**. Space is real only to an observer's relationships: distance cannot be measured without a point from which to measure. Space extends only to the furthest edges of an observer's relational network. Space is not a thing; it is the geometric information describing how interaction points relate from a given perspective.

Time is real only in an observer's interactions. Without causal chains, there is no before or after. If interactions stop, time stops, and vice versa. When both space and time exist from an observer's point of view, the experience is *physical*. Physicality is not a separate substance; it is the rendering of geometric relations and causal sequences from within the network. A simulation's inhabitants experience their world as physical because, from their perspective, space and time are real. We are no different. A world supplied with the same relations is real in the only sense the word can hold, because real has never meant anything but what a vantage's relations render to it. Information, relation seen from within, is therefore not a description laid over some firmer stuff; it is the whole of it. Existence is relational and rendered from inside, and there is nothing else for it to be.

There is no view from nowhere. Every measurement, every experience of space and time, happens from a specific interaction point. The manifested universe is rendered by its observers, every single one of them, at every scale. In the age of computers and synthetic intelligences, this stance could carry serious implications for ethics, and open new ways to think about life and sentience.

One question remains open in this picture. If logic precedes manifestation, why does anything need manifesting at all? The truth of 1+1=2 required no event. Most logical structure plainly does without manifestation.

If no interaction point had a limited perspective, if every point saw the global state, the network would be what is called a stabilizer code. Geometry would exist in description but would not respond to matter. Time would have no privileged present; all moments would be of equal value. Nothing would feel like anything, because feeling requires an inside, and an inside is exactly a limited perspective. Such a model is deficient and cannot explain why the feeling of subjective experience exists or why there is such a thing as now. In this framework, those hard problems dissolve into one solution: the condition that the process is seen from the inside, by observers who cannot know the entirety of the system. It follows that this also fits well with the unification of quantum gravity and general relativity and other troubling phenomena.

---

4. Time and Mass

4.1 Time as Causal Chains

As threads weave among each other, they interact. Each interaction is a step in a causal chain, one thread influencing another, creating new patterns of meta-identity. From the perspective of any interaction point, these causal chains *are* time. An observer's timeline is the sequence of interactions it participates in.

Without these chains, there is no time. An isolated observer, without interactions, exists in a timeless state: there is literally nothing to measure as before or after. The two are inseparable: without one, the other has no meaning. The causal chain iterates toward increasing entanglement complexity, and this increase is correlated with the arrow of time. But time is not an outside force driven by complexity, rather a relational measurement of causal events. The speed of time is similarly only meaningful relative to interactions between observers. Due to the current nature of the network, entanglement tends to increase with these causal interactions, and time's arrow points in the same direction.

Proper time from entanglement density

Time emerges as the gradient of entanglement complexity:

$$\mathcal{I}(t+1) \geq \mathcal{I}(t)$$

For a specific observer O, subjective proper time follows the entanglement density along its causal path:

$$d\tau = \sqrt{1 - \frac{\rho_E}{\rho_{\text{sat}}}}\,dt \tag{4}$$

where ρ_{sat} is the saturation density at which distinguishability fails. To leading order:

$$d\tau \approx \left(1 - \frac{\rho_E}{2\rho_{\text{sat}}}\right)dt$$

This matches the weak-field gravitational result $d\tau = (1 + \Phi/c^2)dt$ under the dictionary $\rho_E/\rho_{\text{sat}} = -2\Phi/c^2$. Two observers at different field values measure different proper time because the local entanglement density differs: higher density near mass slows the clock. This is relativistic time dilation with the correct direction.

The horizon limit follows directly: as $\rho_E \to \rho_{\text{sat}}$, $d\tau \to 0$. Proper time ends where distinguishability saturates, the framework's own description of the black hole interior.

The arrow of time follows from data processing: the master inequality, evaluated along a causal chain.

Full Derivation: Time

Operational definition. Proper time along a thread is its count of distinguishability events: interactions registered in its history (A3).

Theorem 7.1 (Dilation, sign corrected). With ρ_{sat} the saturation density at which distinguishability fails,

$$d\tau = \sqrt{1 - \rho_E/\rho_{\text{sat}}}\, dt$$

To first order $d\tau \approx (1 - \rho_E/2\rho_{\text{sat}})\, dt$, matching the weak-field law $d\tau = (1 + \Phi/c^2)\, dt$ under the dictionary $\rho_E/\rho_{\text{sat}} = -2\Phi/c^2$. The Schwarzschild factor $\sqrt{1 - r_s/r}$ is reproduced exactly for $\rho_E \propto 1/r$. The identification $\rho_E \propto -\Phi$ is the one assumption, flagged.

Two consequences:

- As $\rho_E \to \rho_{sat}$, $d\tau \to 0$: proper time ends where distinguishability saturates. The time law and the black hole interior now state the same fact.
- If subjective time inherits the law through the self-model's internal density, states that lower internal ρ_E yield more experienced time per clock time. Matches reported phenomenology; testable against DMN measures.

The arrow. The spread of correlation (number of pairs with $I > 0$, and coarse-grained entropy) is non-decreasing under generic local dynamics (A2 typicality + Lemma 4.2 applied to coarse-graining). Global information is conserved under unitarity; what grows is its dispersion. The arrow points along increasing entanglement complexity.

4.2 Mass from Thread Bundling

As causal connections multiply, threads begin bundling together. Each bundle adjusts to the presence of others, creating conditions for more bundling. Connectedness breeds connectedness: to debundle is more energetic than to bundle, implying a naturally recursive loop.

Bundling density is mass. The classes this scale sorts matter into, massless photons and gluons, transient W and Z, and persistent fermions, are set out in §8.

The tightness of thread bundling creates what is experienced as mass. The denser the bundling, the greater the mass. This gives a new perspective on the equivalence of energy and mass: different manifestations of thread bundling density. Larger bundles have higher interactive, configurative, and functional potential, and that potential is what collectively manifests as the property we call mass.

Bundling, entanglement density, and mass

The energy of a thread tangle relates to its topological invariants (writhe, linking number, twist):

$$E_{\text{tangle}} \sim f(\mathrm{Lk}, \mathrm{Tw}, \mathrm{Wr}), \qquad \mathrm{Lk} = \mathrm{Tw} + \mathrm{Wr}$$

Mass is the integrated entanglement density over a spatial slice:

$$m = \frac{1}{c^2} \int_\Sigma \rho_E \, dV$$

With the energy-momentum tensor, identifying total energy and comparing with the mass formula yields directly:

$$E_{\text{rest}} = mc^2$$

Mass and energy are different manifestations of the same underlying phenomenon: thread entanglement density. Inertia is the resistance of thread bundles to reconfiguration.

There is a reason inertia and energy are the same currency here. The rate at which any system can move to a distinguishable state is set by its energy, so energy is the capacity for distinguishable change per unit time. Mass is that capacity locked into a standing pattern.

Energy as the rate of change

The Margolus-Levitin theorem bounds the time to reach an orthogonal state by $t_\perp \geq \pi\hbar/2E$. The maximal rate of distinguishable change is therefore $2E/\pi\hbar$: energy is distinguishability per unit time. Rest energy is the internal rate of a bundle, and $E = mc^2$ reads as the total internal rate of distinguishable change.

Full Derivation: Energy and Mass

The energy axiom is a theorem. The framework opens with "energy means capacity for change." Quantum mechanics says exactly this with a constant attached:

$$t_\perp \geq \pi\hbar/2E \quad \text{(Margolus-Levitin)}$$

The maximal rate of distinguishable change (orthogonalization rate) of a system is 2E/πħ. Energy is the capacity for distinguishability per unit time, by theorem.

The speed of light is the Lieb-Robinson velocity. For dynamics local on the regular graph, correlations propagate no faster than an emergent maximal velocity v_LR. Define c := v_LR of the regular component. Relativity's speed limit is inherited from the locality of interaction (Lieb-Robinson bound). Shortcut edges do not violate it because, by A2′, nothing propagates along them. Isotropy and universality of c require homogeneity of the regular structure: part of A4's cost, stated.

From inside the weave, c is not a speed through space but the net's own step rate read in emergent units. A vantage measures distance by counting edges along the shortest path (§3) and measures time by counting its own updates (§4.1), and these are one process, one step per edge, seen two ways. So when any vantage divides the farthest a change can reach in one step by the time that step takes, it reaches the same number, and no vantage assembled from those steps can find a faster count to beat it. This is why c is the same for every vantage: it is the rate at which the net brings one interaction to bear on the next, and every inside is built from that rate. A change cannot reach a distant vantage sooner than the chain of updates between them allows, and that ceiling is c, a fact about how the net updates rather than about any frame laid over it.

Mass. A bundle's rest energy is its internal distinguishability rate (Margolus-Levitin), and its inertia is the cost of reconfiguring topologically protected structure:

$$m = \frac{1}{c^2} \int_\Sigma \rho_E \, dV, \qquad E_{\text{rest}} = mc^2 \tag{8}$$

with tangle energy constrained by Lk = Tw + Wr (Calugareanu-White-Fuller). The integral is a definition; the identification of internal event rate with ρ_E is the assumption; the c2 conversion is the lattice constant (length per step over time per step, squared) fixed by the geometric and temporal constructions established above. The relation E = mc^2 follows by identifying rest mass as the internal distinguishability rate per c^2; the non-trivial content is that this mass creates the correct gravitational field via the entanglement tensor (Section 5), which requires the Jacobson derivation, not a definition.

---

5. The Emergence of Gravity

5.1 Causal Diamonds and Magic

Mass is bundling density. Gravity is what bundling density looks like from a distance.

Dense bundling raises the entanglement density of a region, and entanglement density is what fixes the lengths of edges. Around a planet or a star the edge structure is loaded, so the shortest paths through the network bend toward the dense region. A passing bundle follows those paths, tracing the geometry of the network, pulled by no force.

Gravity is the network's information topology, experienced by the bundles moving through it. An observer follows the path of least thread rewiring. This is the geodesic of the information gradient.

One more ingredient completes this picture. Without it, the network would exist as a static geometry: shaped, but frozen. The ingredient is what recent work in quantum information theory calls *magic*.

Magic operates as a dynamic process. A interacts with B, and mutual information is created: a new edge in the graph. That new edge changes the local entanglement density, which bends the metric slightly. Threads passing through the region follow a new geodesic, interact differently, and produce new mutual information, new density, new curvature. The loop repeats.

This is a feedback cycle. The network never reaches equilibrium because every reconfiguration creates new perturbations.

Every observation is an interaction, and every interaction perturbs the network. That perturbation propagates, reshaping local geometry, which alters the paths future interactions take. The loop does not close. Without magic, the network would be a stabilizer code: geometry would exist but could not respond. Magic makes the encoding approximate, information is always partly lost in any local patch. That loss bends the network, and the bending is gravity.

Provenance. The chain from relative entropy to the Einstein equations is due to Jacobson (1995) and Faulkner et al. (2014); here the metric is built explicitly from mutual information via edge lengths and geodesic closure, and the iterative feedback makes the derivation dynamical rather than kinematic. The magic-backreaction link, that without non-Clifford magic the network is a stabilizer code with rigid geometry and no gravitational response, is corroborated in holographic settings by Cao et al. (2024, 2026) and Munizzi et al.; here it operates directly on the thread graph, with no holographic dual.

Computational verification (§11.9). The magic↔backreaction link is confirmed background-independently in simulation: stabilizer states give a flat reduced spectrum and a trivial modular Hamiltonian $K = -\log \rho$ with no responsive area operator, while non-local magic yields a non-trivial K and a responsive area operator, with the first law $dS = d\langle K\rangle$ reproduced; a stochastic rule generates the required non-local magic spontaneously. The driven steady state was checked separately and behaves consistently: with coherent magic the geometry's responsiveness holds per unit structure and the reduced spectrum becomes non-flat from exactly flat at zero magic, so the magic→geometry relation extends away from equilibrium. Magic lowers the amount of entanglement the flow sustains, not its responsiveness (§11.9).

The Magic Ingredient

A recent line of results (Cao, Cheng, Karthikeyan, Li and Preskill, 2026) supports the final piece. The entanglement tensor requires *non-Clifford* ("magic") contributions to produce gravitational curvature. Without magic, the Thread Network is a stabilizer code: it produces space structure via entanglement, but no bending in response to matter. Magic is the formal measure of what the framework qualitatively calls thread bundling complexity and interaction depth.

Exact stabilizer codes produce a state-independent area with no backreaction: the geometry exists but is inert. The *imperfection* in the encoding makes gravity possible. Perfect thread encoding would produce flat, inert networks; gravity is the imperfection: it emerges because thread relationships cannot be perfectly decoded from subsets of the network.

The entanglement tensor and the gravity bridge

The entanglement tensor is the second functional derivative of entanglement entropy with respect to metric perturbations, and decomposes into entanglement and magic components:

$$\mathcal{E}_{\mu\nu} = \mathcal{E}_{\mu\nu}^{\text{stab}} + \mathcal{E}_{\mu\nu}^{\text{magic}}$$

Here $\mathcal{E}^{\text{stab}}$ captures the *shape* of space, while $\mathcal{E}^{\text{magic}}$ captures its *flexibility*, the non-Clifford contribution enabling matter-geometry coupling. The metric perturbation is proportional to the full tensor:

$$\delta g_{\mu\nu} = \ell_P^2\, \mathcal{E}_{\mu\nu}$$

The tensor must be conserved ($\nabla^\mu \mathcal{E}_{\mu\nu} = 0$) and Lorentz-covariant; it is positive semidefinite as a matrix automatically, since $\mathcal{E}_{\mu\nu} = \sum_y \mathcal{I}(x,y)\, n_\mu n_\nu$ with $\mathcal{I} > 0$. Positivity of relative entropy underwrites the averaged null energy condition (ANEC, Faulkner et al. 2016), not a pointwise one; the framework's phantom phase is a pointwise-NEC violation consistent with ANEC (§1.1, §7).

This is the master inequality, evaluated at a causal horizon.

The quantum-to-gravity bridge. A causal diamond is what a thread can reach within a finite number of interaction steps. From any point in the network, the question is: which other threads can I interact with, given how many edges I have traversed so far? The answer is the diamond, growing with every interaction as more edges expand the horizon.

Inside the diamond, all entanglement is summed. More threads in the diamond means more shared information, more entropy, more weight. The diamond presses on the surrounding network, which presses back, and the compromise is curvature: the diamond bends the local geometry.

Many such diamonds are distributed across the network. Each region has its own thread density: dense regions produce deeper curvature, sparse regions lie flatter. Across all of them the full geometry of spacetime follows: the curvature at any point is the sum of contributions from all the diamonds that contain it. Einstein's equations emerge from this superposition.

The quantum-to-gravity bridge

Any small region of spacetime can be viewed as a causal diamond, a bounded region where quantum-information principles apply. Inside, the distribution of information (entropy) relates directly to geometry. The derivation proceeds through relative entropy positivity:

- **Causal diamond.** Center a small causal diamond at a point; locally, spacetime is approximately flat.
- **Relative entropy positivity.** For any perturbation around vacuum, $D(\rho\|\sigma) \geq 0$, where $K = -\log\sigma$ is the modular Hamiltonian.
- **The entanglement first law.** First-order: $\delta S = \delta\langle K\rangle$. The vacuum modular Hamiltonian of a ball takes the universal form $K = 2\pi \int_B \dfrac{R^2 - r^2}{2R}\, T_{tt}\, d^3x$.
- **Second order.** Positivity gives $\delta^2\langle K\rangle - \delta^2 S \geq 0$, with $\delta^2 S = \int \mathcal{E}_{\mu\nu}\, \delta g^{\mu\nu}\, d^4x$. Imposed on all states and all diamonds this is the canonical-energy condition of Faulkner et al. (2014), which forces the

linearized Einstein equation, schematically $\delta g_{\mu\nu} = \ell_P^2\, \mathcal{E}_{\mu\nu}$. The coefficient is the one calibrated to Bekenstein-Hawking ($\alpha = 1/4G$, the area-law calibration); the sign reproduces the Schwarzschild factor exactly (Theorem 7.1) and is attractive.

- **Einstein's equations.** Conservation with the Bianchi identity yields the field equations, with Λ arriving as an integration constant.

A bootstrap through progressively curved backgrounds is conjectured to extend this to the full nonlinear Einstein equations. The entire procedure requires magic at every order; if magic vanishes, the coupling terms collapse and only flat spacetime survives. **Spacetime curvature is not a separate entity: it is the macroscopic manifestation of quantum-information relationships.** This reverses the usual program: rather than quantizing gravity, gravity emerges from quantum principles.

The connection runs in both directions. Forward, the chain goes from the quantum state to the field equations: the state fixes the shared-information graph, the graph fixes the metric, and the first law on causal diamonds forces Einstein's equations. Backward, a given geometry fixes the causal diamonds, the diamonds fix how strongly regions are entangled, and the dictionary inverts: correlations fall off with distance exactly as the emergent metric demands. The backward direction recovers the entanglement pattern rather than the exact microstate, and that coarseness is what it means for geometry to be a summary of the state. A spacetime is a pair that survives the round trip, and the fixed point is the statement that the trip closes.

A discrete substrate is expected to break Lorentz symmetry, and observation constrains such breaking severely. At the free level the framework's automaton class meets this: the light cone comes out exact and shared by every species, with corrections confined to orders beyond current measurement. The interacting case remains open, and the Complete Mathematics names it as the program's foremost remaining edge.

Quantum Gravity from the State-Geometry Fixed Point

What is called gravity at the classical level and what we call quantum mechanics are two views of one object here: the state-geometry fixed point.

The fixed point says this: the network's geometry is the entanglement pattern that stabilizes itself. The state determines the metric, the metric shapes which interactions are possible, those interactions update the state, and the cycle closes. Both are always present, the state and the geometry evolve together, in step.

Three problems dissolve.

Superposition of geometries. There is no ghostly superposition of different spacetimes. The geometry comes from the full state, not from any single classical configuration. A superposition of two locations has the geometry of the superposition, not two separate geometries. The mistake was thinking of geometry as classical from the start.

The problem of time. Time is relational from the outset. There is no external clock against which the universe ticks. Proper time is the sequence of interactions along a thread. The Wheeler-DeWitt equation's frozen time is what a relational theory looks like from the outside. From the inside, where observers live, time flows because interactions happen.

Singularities. When thread density approaches saturation, proper time slows to zero. The network cannot compress further. The density cap is built into the mathematics: as entanglement density reaches its maximum, distinguishability fails. There is no infinite point. There is only the saturation horizon, exactly the same boundary that appears at a black hole.

Newton's constant measures vacuum entanglement density, a property of how much mutual information the vacuum carries per unit area. This is why G appears in both gravitational and quantum contexts: it measures the same network property in both.

The graviton, in this picture, is a collective excitation of the entanglement pattern. The Weinberg-Witten theorem forbids a massless spin-2 particle from being composite in flat spacetime. But the graviton here never lived in flat spacetime, existing as a fluctuation of the entanglement field itself, and the theorem's assumptions do not apply.

The Bose-Marletto-Vedral experiment will test whether gravity can mediate entanglement between two neutral masses. In the framework the answer is yes, because gravity is entanglement. A positive result confirms the core identity; a null result would mean gravity is not the entanglement pattern.

Full Derivation: Gravity

The chain, with each link labeled:

- Entanglement first law: for small perturbations of a causal-diamond vacuum, $\delta S = \delta\langle K\rangle$, a direct consequence of positivity and monotonicity of relative entropy.
- Imposing the first law on all local diamonds yields Einstein's equations as an equation of state (Jacobson 1995; refined by Faulkner et al.).
- The reading here: $E_{\mu\nu}$ built from I-weighted dyads is the source; $\delta g \propto l_P2\ E$.
- Response requires magic: a pure stabilizer state's entanglement pattern is rigid, and gravitational backreaction requires non-Clifford resource (A5; the Cao, Cheng, Karthikeyan, Li and Preskill (2026) line of results is the supporting literature). The necessity direction is conjectured.

Gravity is iterative because observation is iterative: each interaction perturbs I, hence $E_{\mu\nu}$, hence geodesics, hence future interactions. The loop is the dynamics; the decoupled fixed-background description is its adiabatic approximation.

Derivation chain

The full seven-step chain from relative entropy to the Einstein field equations is written in the quantum gravity section (steps 1 through 7). The content specific to this section: bundling raises ρ_E; ρ_E loads edges and shortens effective paths toward dense regions; free motion follows the path metric, so trajectories bend with no force carried. The area expansion used in step 4 of that chain is the leading order of the Raychaudhuri equation: focusing of null rays is the geometric face of entanglement loading.

Full Derivation: Quantum Gravity

The State-Geometry Fixed Point

Gravity is not quantized here. It emerges from quantum principles that are information-theoretic, not gravitational. The central equation:

$$g^* = \Phi(\psi^*), \qquad \psi^* \text{ self-similar under } U_{g^*}$$

A spacetime is a self-consistent pair: a quantum state whose entanglement pattern generates exactly the geometry under which its own local dynamics preserves that pattern. Einstein's equations are the first-order stability conditions of this fixed point.

Perturbation Theory Around the Fixed Point

Step 1. The vacuum is a local minimum of relative entropy. Let σ be the vacuum state restricted to a small causal diamond and ρ = σ + ε δρ a perturbation with Tr(δρ) = 0. The first variation of von Neumann entropy is δS = -Tr(δρ log σ), so

$$\begin{aligned} D(\rho\|\sigma) &= -S(\rho) - \mathrm{Tr}(\rho \log \sigma) \\ &= -S(\sigma) + \epsilon\ \mathrm{Tr}(\delta\rho \log \sigma) - \mathrm{Tr}(\sigma \log \sigma) - \epsilon\ \mathrm{Tr}(\delta\rho \log \sigma) + O(\epsilon^2) \\ &= O(\epsilon^2) \end{aligned}$$

The relative entropy vanishes to first order: the vacuum sits at the bottom of the relative-entropy landscape. Every first-order statement that follows is an expansion around this minimum.

Step 2. The entanglement first law. With K = -log σ the modular Hamiltonian of the diamond, the same expansion gives S(ρ) = S(σ) + Tr(δρ K) + O(ε2), that is

$$\delta S_{\mathrm{EE}} = \delta\langle K\rangle$$

Entropy change equals modular energy change, diamond by diamond. For the diamond geometry of A4 the vacuum modular Hamiltonian takes the universal boost form, which is what lets the right-hand side be read as energy-momentum flux through the diamond.

Step 3. Entanglement entropy scales with area. For states with bounded correlation length the vacuum entanglement across the diamond boundary obeys S_EE = α A + subleading. The scaling S ∝ A is derived from the area-law theorem; the coefficient α = 1/4G is matched to Bekenstein-Hawking, a calibration.

Step 4. Area variation reads out Ricci curvature. For a small diamond with future null normal k^μ and affine parameter λ, the cross-section area at distance λ varies as

$$A(\lambda) = A_0\,\lambda^{d-2}\left[1 - \frac{\lambda^2 R_{\mu\nu}k^\mu k^\nu}{d^2-1} + O(\lambda^3)\right]$$

$$\delta A = -\int R_{\mu\nu}k^\mu k^\nu\,\delta\lambda\,d\Omega$$

Curvature is precisely the failure of small areas to grow at the flat-space rate.

Step 5. Combine and force the tensor equation. Insert the area law into the first law: δA/4G = δ⟨K⟩, with the right side the energy-momentum flux. This yields, for every null direction k^μ,

$$R_{\mu\nu}k^\mu k^\nu = 8\pi G\,T_{\mu\nu}k^\mu k^\nu \quad \text{for all null } k \implies R_{\mu\nu} + f\,g_{\mu\nu} = 8\pi G\,T_{\mu\nu}$$

Holding for all null vectors forces the tensorial equality up to a term proportional to the metric. The contracted Bianchi identity ∇^μ(R_μν - 1⁄2R g_μν) = 0 then fixes f = -1⁄2R + Λ with Λ an integration constant:

$$R_{\mu\nu} - \tfrac{1}{2}R\,g_{\mu\nu} + \Lambda\,g_{\mu\nu} = 8\pi G\,T_{\mu\nu} \tag{7}$$

Einstein's field equations, with the cosmological constant arriving as an integration constant rather than an insertion.

Step 6. The source is the graph. In this framework T_μν represents the coarse-grained directional distribution of mutual information:

$$T_{\mu\nu}(x) = \sum_y \mathcal{I}(x,y)\,n_\mu(x,y)\,n_\nu(x,y) \;-\; \text{trace subtraction}$$

the same construction that defines mass from bundling. Every term in Einstein's equations now has a graph-side definition: curvature from area variation of diamonds, stress-energy from the directional I-distribution, and the equation itself from the first law that links them.

Step 7. Second order and stability. At second order the relative entropy expands as

$$D(\sigma + \Delta\rho\,\|\,\sigma) = \tfrac{1}{2}\mathrm{Tr}\big(\Delta\rho\,L_\sigma[\Delta\rho]\big) + O(\Delta\rho^3)$$

where L_σ is the Hessian, the Fisher information metric on states, positive semidefinite by monotonicity. Faulkner and collaborators showed this second order reproduces the nonlinear Einstein corrections at that order. Monotonicity also makes the relative entropy to the vacuum a Lyapunov function: it can only decrease under the dynamics the perturbed geometry defines, so the fixed point is an attractor. What remains unproved is existence and uniqueness of the fixed point itself, listed in the open problems.

Three Structural Problems Dissolved

The superposition objection. Semiclassical gravity ($G_{\mu\nu} = 8\pi G \langle T_{\mu\nu}\rangle$) fails for superpositions of massive bodies because it sources geometry from expectation values. The framework sources geometry from the full entanglement pattern $E_{\mu\nu}$ of the quantum state, not its expectation value. A superposition of two mass configurations is a state whose I-graph encodes both geometries together with their interference. "Superposition of geometries" is not a paradox; it is the generic case.

The problem of time. Canonical quantum gravity freezes (Wheeler-DeWitt: the total state is stationary). The framework's time was relational from the start: proper time is interaction count along an identity chain, which is exactly the Page-Wootters mechanism, where time emerges from entanglement between a clock subsystem and the rest. The framework adopts Page-Wootters natively rather than encountering the frozen-time problem at all.

Singularities. $\rho_E \leq \rho_{sat}$ is built into the kinematics, and the corrected time law $d\tau = \sqrt{(1 - \rho_E/\rho_{sat})}\, dt$ ends proper time at saturation without any infinite density. The interior continues as graph dynamics without a geometric description. Singularity resolution costs nothing extra.

This interior is more than a resolved singularity. It does not simply end at the density cap. What it continues into, the information that crosses its horizon, and the re-genesis it seeds are treated together in Section 6.2.

Newton's Constant and the Graviton

Newton's constant. Jacobson's derivation fixes 1/4G as the vacuum's entanglement per unit emergent area, a relationship, not a numerical prediction. The framework relates G to the vacuum $\mathcal{I}$-density but does not compute G's numerical value from first principles; G is a calibration that matches the framework's units to observed gravitational strength. The hierarchy question "why is gravity weak" becomes "why is the vacuum so highly entangled," which relocates the question to an object the framework already contains. This is Sakharov's induced gravity (1967) merged with Jacobson (1995).

In this picture, gravity is not a fundamental force that must be quantized. Gravity is entanglement thermodynamics. The Einstein equations are the equation of state of the thread graph's entanglement distribution. The apparent incompatibility of general relativity and quantum mechanics dissolves: quantum mechanics provides the entanglement structure, and general relativity describes its thermodynamic equilibrium. They are not competing theories; they operate at different scales of the same underlying process.

This resolution also explains a deeper tension. Quantum mechanics demands unitarity, information must be conserved. General relativity demands a dynamical spacetime, geometry must evolve. How can both be true? In the thread graph, total mutual information is bounded by holography ($N \cdot \langle I \rangle = A/4l\_P2$), but the distribution of I changes continuously through interactions. The total is conserved; the distribution evolves. Geometry, the metric derived from I-distribution, is dynamical even though information is conserved. The thread graph redistributes its information content without creating or destroying it.

The graviton and Weinberg-Witten. The Weinberg-Witten theorem forbids a massless spin-2 composite in a Lorentz-invariant theory with a covariant stress tensor. The framework's escape is that Lorentz invariance is emergent (the Lieb-Robinson cone of A2′), so the theorem's hypotheses fail. The graviton is a collective mode of the entanglement pattern, and it is exactly as fundamental.

Lorentz invariance: cost and protection. Emergent Lorentz invariance is generically violated at low energies by radiative corrections unless protected. Observational constraints on Lorentz violation push linear-order violations beyond the Planck scale. The framework, like every discrete-substrate approach, owes a protection mechanism. It provides one, not by fine-tuning, but by symmetry.

Rotational invariance: symmetry-protected to all orders. The automaton lives on a cubic lattice with point group O_h (48 elements). The space of O_h-invariant rank-2 tensors is one-dimensional, spanned by δ_{ij}, verified by explicit group averaging (off-diagonal entries zero to machine precision). The T-gate, the non-Clifford resource that makes geometry responsive, is $T = \mathrm{diag}(1, e^{i\pi/4})$, an on-site Z-rotation. Injected at uniform density with no preferred axis, the perturbation $V = (\pi/8) \Sigma_x Z_x$ is invariant under O_h (the point group permutes sites; the uniform on-site sum is unchanged): V belongs to the trivial representation A_1g. The free propagator and the Brillouin zone are O_h-symmetric. Therefore every correction to the kinetic tensor Γ_{ij} at $O(k^2)$, at any order in the magic coupling, is an O_h-invariant rank-2 tensor, hence proportional to δ_{ij}. Rotational isotropy of the $O(k^2)$ dispersion is symmetry-protected to all orders.

The only allowed anisotropy enters at $O(k^4)$ (the cubic harmonic $\ell=4$), whose relative effect scales as $k^2 \to 0$ in the infrared, power-counting irrelevant. A one-loop self-energy computation confirms this: uniform (O_h-symmetric) magic injection yields $T_{ij} = 1.7\times10^{-4} \cdot \delta_{ij}$ (off-diagonal 1×10^{-18}, isotropic); axis-aligned magic injection breaks the symmetry and yields measurable anisotropy ($T_{xx} = -0.055$, $T_{yy} = T_{zz} = 3\times10^{-4}$). The automaton's own interaction rule selects pairs with the isotropic probability $p = \exp(-d)$, placing it in the protected case by construction. The falsifiable corollary: deliberately axis-aligned magic injection would produce measurable Lorentz violation.

Physical statement. Non-local magic is the resource that makes geometry responsive, no stabilizer code supports a non-trivial area operator (Cao 2024), and gravitational backreaction is the holographic dual of magic (Munizzi 2024). The present result adds: because magic enters as a point-group singlet, rendering geometry responsive does not break the emergent rotational invariance of the light-cone. Responsive gravity and emergent isotropy are compatible, and the compatibility is symmetry-protected rather than fine-tuned, precisely what the standard radiative-correction objection (Collins et al.) demands and does not usually get.

Scope. This establishes rotational invariance (isotropy) of the IR dispersion. Boost invariance and single-cone universality, that all modes share one speed c, remain open as the L1 velocity-convergence program. Free-level Lorentz protection is independently confirmed by QCA inheritance (exact shared light cone, even-order corrections only, confirmed numerically at Test A, §11.7).

Candidate protection mechanisms. Several mechanisms could let emergent Lorentz invariance survive radiative corrections, each a route not yet settled by calculation.

L1. Criticality from the percolation exit. The genesis cascade halts at the percolation transition, which hands the infrared a scale-invariant starting point for free. At critical points, mode velocities flow to a common value in the infrared (the mechanism that equalizes velocities in interacting Dirac systems). If the exit tunes the network to criticality, a universal cone is the attractor, not an accident. The framework already has the critical point; this connects two existing pieces.

L2. One substrate, one cone. The Collins et al. fine-tuning argument assumes independent fields, each with its own propagation speed to renormalize. Here every species is a collective mode of the same I-graph with a single Lieb-Robinson velocity; there is no second speed for corrections to drive apart. The task reduces to showing interactions do not dress different modes with different effective cones.

L3. The quantum cellular automaton route. Existing QCA constructions (Bialynicki-Birula; Arrighi; D'Ariano and collaborators) recover the Weyl and Dirac equations with exactly conical dispersion up to corrections of even order in the lattice scale, pushing violations to dimension-six operators beyond current bounds. If the interaction automaton coarse-grains to a QCA of this class, protection is inherited from an established result.

The inheritance argument, first analytic contact. The interaction automaton satisfies the three defining properties of a quantum cellular automaton: global unitarity (composition of local unitaries), strict locality (finite propagation per step, A2′), and homogeneity on average (no preferred site). For automata with these properties plus isotropy, the small-wavenumber limit of the internal dynamics falls in the Weyl-Dirac family (D'Ariano and Perinotti derive the Weyl equation from exactly these requirements). The free sector of the framework therefore inherits the QCA protection demonstrated numerically in the computation section: exact shared cone, corrections at even orders only. What is not yet mapped is the interacting case, for which interacting QCA constructions exist in the literature as the target.

L4. Lorentz-violating operators are irrelevant at the fixed point. The state-geometry fixed point is an attractor by the Lyapunov property. A frame-dependent perturbation moves the state off the fixed point; if its scaling dimension at the fixed point is irrelevant, monotonicity damps it faster than radiative corrections feed it.

L5. Boosts as modular flow. Bisognano-Wichmann identifies boosts with the modular flow of wedge algebras, and the framework already runs on modular Hamiltonians of diamonds. If a graph-side Bisognano-Wichmann holds (modular flow of a half-graph approaches a graph automorphism in the dense limit), boost invariance is inherited from the same entanglement structure that yields the first law, and needs no separate protection.

L6. Statistical isotropy of the grown graph. Violations require a preferred direction; the grown graph is random with no lattice axes, so frame-dependent operators vanish on ensemble average and fluctuate at one over root N. Suppresses but does not forbid; supports L2 rather than standing alone.

L7. Deformed rather than broken. The substrate may preserve an exact quantum-group deformation of the Poincaré algebra with l_fund as deformation scale; violations are then confined to specific energy-dependent dispersion signatures rather than generic operators, evading present bounds while predicting time-of-flight effects for gamma-ray bursts. Turns the wall into a prediction.

L8. Dilution of frame information. A preferred frame is long-range mutual information with a global direction; the dilution theorem drives pairwise I between distant subsystems down, so frame correlations cannot persist at large scales. Kills constant background frames; must still address locally generated violations.

L9. Fixed-point re-absorption. A frame-dependent term shifts Φ(ψ), and re-solving the fixed point re-absorbs it into g*: Lorentz violation as a coordinate artifact of holding the old metric fixed while perturbing the state. If the absorption is exact at first order, low-energy violations are pure gauge.

L10. Observer selection. Regions where mode velocities fail to equalize cannot bind stable bundles (coherent propagation is what bundling needs), so observers only condense in Lorentz-symmetric domains. A last resort, explanatory rather than protective.

The free-level case is settled (§11.7, §11.8); the interacting case remains open.

The BMV Test

If geometry is entanglement structure, then the gravitational interaction between two masses is, microscopically, shared entanglement responding to both. A mass in a superposition of two positions deforms the pattern in superposition. Consequence: gravity can generate entanglement between two otherwise isolated masses. This is exactly what the BMV experiments (Bose et al. 2017; Marletto and Vedral 2017) are being built to test.

A controlled BMV-type experiment showing gravity decoheres without entangling would refute the core identity that gravity is entanglement, because for this framework geometry is entanglement, with no softer reading available. The experiments are in active development.

The round trip: quantum to Einstein and back

Forward (quantum to Einstein). The chain of 10.2, steps 1 through 7: a quantum state defines an I-graph; the graph defines the metric; vacuum relative entropy is minimized; the first law on diamonds plus the area law plus area variation force, for all null directions, the Einstein field equations with Λ as integration constant and T_μν read off the graph. Conditional on the tagged identifications, the forward direction is a derivation with every step written.

Backward (Einstein to quantum). Given a geometry: the metric defines causal diamonds; diamonds define modular Hamiltonians of universal boost form; the boost form plus the area law fix the vacuum entanglement structure; and the dictionary inverts the metric definition,

$$I(i,j) = I_max \cdot exp(-d(i,j)/l)$$

recovering the correlation graph from distances: correlations decay with geometric separation exactly as the emergent metric demands. The backward direction recovers the entanglement pattern, not the microstate; many states share one pattern, and that degeneracy is what geometry being a coarse description means.

Closure. The fixed point g = *Φ(ψ)* is the statement that the two directions compose to the identity on solutions: a spacetime is exactly a pair that survives the round trip. Stability of the closure is the Lyapunov property of 10.2 step 7; existence and uniqueness of the pair remain open and are listed. So the answer to whether one can go from quantum to Einstein and back is yes, conditionally and with the steps written, in both directions, with the round trip closing precisely at the fixed point.

---

6. Inflation, Black Holes, and Re-Genesis

6.1 The Genesis Cascade

The bifurcation that began existence did not act once and stop. It kept going, and that runaway is inflation.

Each generation doubled the network: one thread became two, two became four, four became eight. The number of interaction points grew exponentially. Before geometry had emerged, before the speed of light had crystallized as a limit, there was nothing to restrain this doubling. The cascade ran outward, and the universe inflated.

No extra field drove this. No inflaton particle. The bifurcation operator, the same piece of mathematics that turns nothing into something, is the engine of inflation. Applied once, it gives genesis. Applied repeatedly, it gives the cascade.

The cascade's termination mechanism is not yet derived from first principles. Candidate hypotheses include saturation of the holographic bound, geometric closure of the emergent de Sitter horizon, and dynamical fixed-point attraction;

none have been numerically confirmed. The cascade currently lacks an intrinsic stopping condition in the formalism, this is equivalent to the cosmological constant problem in all approaches to quantum gravity.

The cascade's own termination, which fixes the number of e-folds and the cosmological-constant scale, is the open calculation that remains here. The separate questions of whether each sealed interior forms a self-sustaining successor, and whether the succession continues without end, are treated in Section 6.2.

The uniformity of the cosmic microwave background is the memory of this pre-geometric phase. Before the Lieb-Robinson cone emerged, before signals had a maximum speed, the entire network was in causal contact. The smooth sky is exactly what this picture predicts.

And the residue? The irreversibility of the original bifurcation left a permanent tension per edge. Under volume-extensivity, the assumption that residue-carrying edges are minted at constant physical density as space expands, this tension acts as a cosmological constant, the dark energy that keeps the expansion going. If instead the residue is carried on a fixed comoving edge set, its density dilutes as a-3 and it contributes to dark matter rather than dark energy. The volume-extensivity premise remains an open question (§7).

Direct simulation of the fixed-point dynamics (§11.9) clarifies the cascade's role. The literal interact+dilute rule has only a *dead* attractor; a live, extended geometry is sustained only under a continuous *distributed* genesis source, a driven dissipative steady state. This recasts genesis: it is not merely the first event but an *ongoing* process whose throughput holds spacetime away from the trivial (product) equilibrium. The cascade's missing stopping condition and the persistence of structure are then two faces of one balance, source against dilution, rather than separate puzzles.

Inflation as Genesis Cascade

The bifurcation operator $\hat{B}$ applied iteratively:

$$\hat{B}^n : \mathcal{H} \to \mathcal{H}^{\otimes 2^n}$$

After n generations:

$$a(n) \propto 2^{n/3}$$

Corresponding to exponential expansion with effective Hubble rate $H = \ln 2/(3t_s)$ where t_s is the step time. No inflaton field is required; the cascade is driven by the same isometry that powers genesis.

Graceful exit, geometric feedback. The cascade does not require an external stopping mechanism. As the network grows, the emergent metric feeds back on the dynamics: distances increase, interaction probabilities p = exp(-d) decrease, and the effective genesis rate per qubit declines. At a threshold network size set by the dilution parameter η, the genesis rate approaches the dilution rate, and the cascade transitions naturally into the driven steady-state expansion. The exit from inflation is a phase transition in the ℐ-density, not an imposed boundary condition. Numerical simulation of a geometric cascade model confirms this feedback mechanism: the cascade self-terminates, with the number of e-folds N_e determined by η rather than taken from observation. At mean-field level the self-termination takes the logistic form $\dot{N} = rN\left(1 - N/N_*\right)$: exponential while N is small, saturating as N approaches a ceiling N_* fixed by the dilution rate, with no external stopping condition. The closed form is model-level; a first-principles derivation of the stopping condition is the open problem noted above.

N_e as output. The cascade model gives N_e = O(10-30) e-folds for η in the framework's measured range (0.63-0.70). This is fewer than the ~60 e-folds of standard cosmology, and fewer are needed, because the framework's pre-geometric causal contact (§2.1) already resolves the horizon problem without requiring 60 e-folds of inflation. N_e is a prediction, not a consistency input. The exact value is model-dependent (the geometric feedback has free parameters at the mean-field level); the robust result is the order of magnitude and the functional dependence N_e(η).

Tensor-to-scalar ratio. The cascade's pre-geometric phase has no metric to perturb: tensor modes are naturally suppressed. The framework qualitatively predicts a small tensor-to-scalar ratio r, consistent with current upper bounds (r < 0.03).

n_s and the spectral index. The spectral tilt n_s is not yet derived from the framework. The universal slow-roll relation $n_s - 1 \approx -2/N_e$ does not apply: the cascade is a quantum computation, not a slowly rolling scalar field. With the cascade's N_e of order 10-30 e-folds (see Graceful exit above), the slow-roll proxy would give $n_s \approx 0.93 - 0.97$, consistent with Planck's $n_s = 0.965 \pm 0.004$, but this is numerology, not a derivation. A first-principles computation of n_s requires a quantum effective field theory of the ℐ-field fluctuations during the cascade, which has not yet been constructed. The mean-field geometric cascade model (§11.9) demonstrates the termination mechanism but does not capture the quantum fluctuations needed for the power spectrum.

Full Derivation: Inflation as Genesis Cascade

The Cascade

Genesis is the isometry B̂: H → H ⊗ H. Iterated in parallel across the state, n generations produce $N(n) = 2^n$ subsystems. Pre-geometric time is causal depth: one generation is one interaction along every identity chain, so proper time advances by one step t_s per generation, τ = n t_s. Then:

$$a(n) \propto N^{1/3} = 2^{n/3} a(\tau) \propto exp((ln2/3) \cdot \tau/t_s) H_i nf = ln2/(3t_s) \approx 0.23/t_s$$

Exponential expansion at one e-fold per 4.3 steps, with t_s of order the Planck time. Sixty e-folds require n ≈ 260 generations. Inflation is the era in which the genesis operator still dominates the dynamics. No inflaton field is introduced; the engine is the same operator that started the universe.

The Classic Problems

Horizon. Before the regular edge structure percolates, there is no local Hamiltonian and hence no Lieb-Robinson cone: correlation is established globally by the cascade itself, because every pair descends from common bifurcations a few generations back. The uniformity of the CMB is the memory of the pre-geometric phase, when "causal contact" had no speed limit because there was no geometry to limit it.

Flatness. The cascade is statistically identical at every node (one rule, applied everywhere), so the condensing manifold inherits homogeneity at the percolation transition, and residual curvature dilutes through the e-folds exactly as in standard inflation.

Monopoles. There is no GUT phase transition in the framework's early universe, so there is nothing to overproduce; the problem does not arise.

Free Capacity, from Monogamy

The cascade requires free capacity: $\hat{B}$ acts on weakly entangled (near-product) degrees of freedom, and the monogamy budget (Christandl-Winter 2004) caps each node's quantum correlation budget at 2S. As the graph populates, budgets fill, the product-state instability is consumed, and bifurcation is progressively suppressed. Inflation ends when the entanglement budget saturates and the regular component percolates into the $d = 3$ structure that Theorem 6.2 forces. The exit is not a tuned potential; it is the same monogamy inequality that fixes the self's boundary in the consciousness sections. One theorem, both ends of the universe.

The Spectrum: Structure, Not Yet a Number

Observation demands $n_s \approx 0.965$. The framework's natural formalism is δN: the curvature perturbation on a comoving region equals the fluctuation in its local e-fold count.

Dichotomy. If the exit is independent per node (Poisson-like), the tilt is blue ($n_s > 1$). If the exit is collective (a percolation transition, long-range correlated, increasingly deterministic as the critical point approaches), the tilt is red ($n_s < 1$). Which case holds turns on the percolation statistics of the exit in three dimensions, which is not yet computed.

Current numerical exploration. Preliminary simulations with 15 cascades show the collective exit structure is present (H8 method yields crossing-generation variance consistent with red tilt), but the sample size is too small for precision: 15 cascades give a sampling error of ~40% on the variance, ten times larger than the effect size. The generic $2/n_*$ term alone gives $n_s \approx 0.967$, closer to observation than any simulation to date. The structure is promising; the number is not yet computed.

6.2 Black Holes, Information, and Re-Genesis

A black hole is the saturation limit of Section 5 taken to its end. As bundling drives $\rho_E \to \rho_{\text{sat}}$, proper time runs out at the horizon (eq. 4) and the singularity is replaced by a density floor, not an infinite point. The interior does not stop there. With no geometric description and no shared time with the exterior ($d\tau/dt \to 0$ at the horizon), it continues as graph dynamics, a region running its own irreducible process with limited perspective on the outside. That this is the genesis condition, and not merely a repeat of it, follows from what a vantage inside the seal can register. Proper time ends at saturation because distinguishability fails there (eq. 4): a node that has spent its full monogamy budget on the joint state, spread across the dense interior, holds vanishing shared information with any single partner, so its outward edges read featureless and no inside vantage can tell its neighbors apart. The interior is blank from within. This is the same blank the genesis nothing carries, where zero distinctions specify one configuration (Theorem 2.1). The two are not the same density, since one end is empty and the other full, but they are the same condition for a vantage: no accessible distinction, and the next distinction still to be produced. A saturated interior is a genesis nothing read from inside, and the genesis derivation applies to it unchanged. Leaving saturation is then forced, since the cap shuts the source off while dilution does not, so the standing density falls by exactly $1-\eta$ on the first step, and that same step reopens the free, near-product capacity $\hat{B}$ acts on (§6.1). The blank does not hold: capacity returns, the first distinction is produced, and the horizon is where $\hat{B}$ fires again, the in-falling branch seeding the sealed interior and the escaping branch the radiation, anti-correlated as Theorem 3.1 requires. This is genesis running a second time.

Information is preserved across the horizon. Because an interaction once written leaves mutual information that never fully returns to zero (A2), the escaping branch stays correlated with its in-fallen partner: the radiation is entangled with the interior, not severed from it. The horizon is an encoding surface. Information written onto it is read back slowly as the region evaporates, so the mutual information between radiation and interior first rises and then falls along the Page curve (Page 1993), and unitarity is inherited directly from $\hat{B}$ being an isometry (Theorem 2.2), which preserves the inner product and discards nothing. No firewall forms, because the correlation reorganizes step by step rather than breaking (Hayden and Preskill 2007).

The interior is then a successor net, a fresh start of the cascade $\hat{B}^n$, and the cosmos a one-way succession of such events, one-way by Theorem 2.2 because $\hat{B}$ has no global inverse, and width-capped by the monogamy budget. Two questions follow, and the recursive cosmogenesis of Sections 2 and 5 left them open: whether a single sealed interior settles into a self-sustaining net rather than draining away, and whether the succession continues without end. Direct simulation of the sealed dynamics (§11.9) separates them into one settled question and one undecidable by construction.

The single forming is settled. A black-hole interior begins at saturation, every pair maximally entangled, the most connected state the network can hold. Because the metric is the geodesic closure of mutual information, shared information makes distances short, so in any connected interior the typical geodesic is short and the chance of any two points interacting stays high, of order one rather than vanishing. The interior therefore keeps remaking its own ties at a high rate, and the balance of this remaking against dilution has a single self-sustaining rest point. Writing the standing density as a fraction of the cap, $x = \rho_E/\rho_{sat}$, the fire-and-dilute balance gives

$$x^* = \frac{\eta p}{1-\eta+\eta p},$$

with p the order-one interaction probability and η the dilution survival. This is positive for every $\eta > 0$, so the interior settles at a self-sustaining density rather than nothing, and there is no competing collapsed rest point in the connected regime; the self-sustaining state is the only one the interior can reach. The interior also cannot fall into the one collapsed state, total disconnection, because the pairs that fire each step form a graph many times denser than the connectivity threshold (the firing fraction sits several times above log N over N), so the chance of the interior coming apart falls away faster than exponentially in its size. A saturated interior, being fully connected, remakes a densely connected graph every step and never disconnects. The collapsed attractor seen for the bare, undriven rule is the sparse, near-empty regime; the saturated start lies in the self-sustaining regime instead, which is exactly why a self-sustaining net is tied to a black hole rather than to ordinary density. What singles out the black hole is not a delicate selection among basins but the seal itself (eq. 4): only a sealed interior runs its own dynamics undisturbed to this rest point, while the enclosing net's ongoing expansion, by the graceful-exit feedback above, holds any unsealed region at a lower driven level.

The succession's growth is undecidable. Whether the succession extends another step asks whether a successor in turn grows complex enough to pack its own saturating regions. This is a reachability question on the framework's own update, and that update applies a universal gate set (CNOT, Hadamard, T-gate; Definition 4 of the companion automaton), so the process is computationally universal and irreducible, the same non-vanishing magic (A5) required elsewhere for gravitational backreaction. Reachability of a target region for a universal process is undecidable, by the halting problem and Rice's theorem: no procedure decides, for every seed, whether a successor ever forms a black hole. So whether the succession continues without end cannot be settled from outside by any calculation; it is settled only by running, which is what the cosmos does by existing. This is not the cascade-termination problem in disguise, and not a gap to be closed later. It is a property the dynamics forces, the same irreducibility that makes the present have to be produced rather than read off (companion paper), now appearing at the largest scale. The single magic resource that turns on gravity and forces an inside is the same universality that makes the succession's reach unknowable from without.

7. The Dark Sector

7.1 Network Tension and Decomposition

The Thread Network carries a residual strain from the first split. Bifurcation was irreversible per edge, and under volume-extensivity (edges minted at constant physical density as space expands), the residual strain maintains constant energy density, stretching outward without clumping or varying from place to place, accelerating the expansion. This is dark energy. If the residue is instead carried on a fixed comoving edge set, it dilutes and contributes to dark matter (see §6).

Where threads bundle, a galaxy forms. The bundle is dense: many Identity Edges and Interaction Edges packed into a small region. This density breaks the local isotropy. Far from the bundle, Interaction Edges point equally in all directions and their net directional pull cancels to zero. Near the bundle, the edges align along the density gradient. A directional tension emerges, pulling inward. This is dark matter: the network responding to its own internal weight.

Dark matter bends the path of photons without scattering them. It pulls on galaxies without emitting light. An observer sees what the curvature does to the light passing through, never the curvature itself. The dark sector is woven into the metric every observer already inhabits.

The entanglement tensor, the same mathematical object that gives gravity, contains both faces. One pulls the universe apart. The other holds it together.

Decomposition of the entanglement tensor

Dark energy is the trace, a homogeneous isotropic component driving the background expansion, under the assumption of volume-extensivity of the bifurcation residue (§6, §R2-P2). Dark matter is the traceless part, an inhomogeneous anisotropic component sourcing additional gravity where matter clusters.

Prediction: phantom crossing from driven dynamics. The framework's spacetime is a driven dissipative structure, not an equilibrium. The dark-energy strain ρ obeys the continuum form of the automaton's own per-edge rule, $\dot\rho = S(a) - \Gamma\rho$, with dilution rate $\Gamma = 1 - \eta$ (≈ 0.35 at $\eta \approx 0.65$) and genesis source $S(a) \propto a^{-3}$ (fixed comoving distinction-creation, physical density falling with volume). The effective equation of state is fixed by the scaling of the strain density:

$$w(a) = -1 - \frac{1}{3}\,\frac{d\ln\rho}{d\ln a}, \qquad \rho(a) \text{ from } \dot\rho = S(a) - \Gamma\rho.$$

At early times the strain lies below its receding attractor $\rho^* = S/\Gamma \propto a^{-3}$ and is rising, so $d\ln\rho/d\ln a > 0$ and $w < -1$ (phantom); at late times it tracks the falling attractor and $w > -1$ (quintessence). The two regimes meet in a single crossing of the phantom line from below, at the epoch where genesis balances dilution ($S = \Gamma\rho$). Numerical integration with a self-consistent Friedmann background (`wz_derivation.py`) confirms this generic Quintom-B form, with

$w_0 > -1$ today, matching the sign of the DESI DR2 preference (w crossing -1 from below). The crossing redshift is set by Γ/H_0 (hence by η and the step-to-Hubble ratio) and is not yet calibrated; in the simple linear model the crossing falls earlier (deeper) than DESI's $z \approx 0.3 - 0.5$, a quantitative tension whose resolution requires a non-negative, saturating treatment of the strain. The qualitative prediction, a single phantom crossing from below, is parameter-free; its location is a one-parameter probe of the substrate scale. This is the framework's *primary* cosmological prediction.

If the strain sits at equilibrium instead. If future surveys find no phantom crossing, the driven-dynamics reading gives way to the equilibrium value $w = -1$, fixed by volume-extensivity of the residue (constant density, a cosmological constant). The monotonicity axiom and the averaged null energy condition it underwrites hold in either case.

Dark energy is the unreleasable strain of bifurcation: the network cannot return to the vacuum it left per edge. Under volume-extensivity, the residue is isotropic and constant in density, giving $w = -1$. If instead the residue is edge-conserved, it dilutes as a^{-3} with $w = 0$, contributing to dark matter rather than dark energy. Dark matter is the traceless response to bundling: it is constrained by conservation $\nabla^\mu \bar{\varepsilon}_{\mu\nu} = 0$ and numerically correlated with matter gradients (Pearson $r = 0.556$, automaton b6).

Current DESI DR2 data (2024–2026) favors this Quintom-B scenario, with w crossing -1 from below at approximately 4σ; the DESI five-year survey concluded in 2026 and final results are pending.

Dark because structureless: dark matter carries low magic content, leaving nothing for light to couple to. Read through the manifestation boundary, this acquires a precise sense: the dark sector is the part of the network that holds without happening, real relations with no inside and nothing to render. Kept as a labeled reading. The dark sector hides in the part of the entanglement tensor that does not radiate.

Full Derivation: The Dark Sector

Decompose the one object by symmetry:

$$\begin{aligned} &\mathcal{E}_{\mu\nu} = \mathcal{E}^{(\text{trace})} + \mathcal{E}^{(\text{traceless})} \\ &\rho_{\text{DE}} = \varepsilon/8\pi G && (\text{trace: isotropic residual strain}) \\ &T^{\text{DM}}_{\mu\nu} = \bar{\varepsilon}_{\mu\nu}/8\pi G, \quad \nabla^\mu \bar{\varepsilon}_{\mu\nu} = 0 && (\text{traceless: anisotropic response}) \end{aligned}$$

(See decomposition above: E^(trace) yields dark energy under volume-extensivity; E^(traceless) yields dark matter, constrained by ∇^μ ε̄_μν = 0 and numerically correlated with matter, r = 0.556.)

The equilibrium alternative. A persistent equation of state would hold $w(z) \geq -1$ at every redshift; a multi-survey absence of phantom crossing would place the cosmology there.

Current status: DESI DR2 data favors w crossing −1 from below at approximately 4σ, consistent with the driven-dynamics prediction. A confirmed absence of phantom crossing would place the cosmology at the equilibrium bound, leaving the monotonicity axiom intact.

Open: the coincidence (why ρ_DE ≈ 2.3 ρ_DM today), the relic abundance, the value of ε. The framework reduces to ΛCDM at leading order and treats that as its entry ticket.

The cosmological constant problem of standard quantum field theory, a 120-order-of-magnitude discrepancy between the observed Λ and the zero-point sum over field modes, does not arise in this framework. The framework never sums zero-point energies. Λ is not a vacuum energy in the QFT sense; it is the total information of the thread graph divided by the horizon area. The "problem" is a category error: comparing a QFT calculation that the framework does not perform to an observation that the framework explains differently. The open question is not "why is Λ small?" but "why is the universe large?", and the two are holographically dual.

Derivation chain

Step 1. Decompose the entanglement tensor: the isotropic trace part sources uniform expansion; the anisotropic traceless part sources clustering without electromagnetic coupling. Dark energy and dark matter are the two reading directions of one object.

Step 2. Energy conditions, correctly placed. Positivity of relative entropy underwrites the averaged null energy condition (Faulkner et al. 2016), an integral statement, not the pointwise one. The two are not equivalent: a source can satisfy ANEC while violating pointwise NEC locally, and the framework's primary prediction does exactly that. The driven strain (Steps 3-4) runs $w < -1$ in its early phantom phase, a local pointwise-NEC violation, while the time-integrated null contraction stays non-negative; numerical integration of the single-crossing solution gives $\int T_{\mu\nu} k^\mu k^\nu \, d\lambda \geq 0$, the late quintessence surplus overtaking the early phantom deficit. At equilibrium the residue is volume-extensive and constant in density, giving $w = -1$ exactly, with no energy-condition inequality required.

Step 3. The driven equation of state, derived. Off equilibrium the residue is a cosmological component of density ρ and equation of state w. The continuity equation on a Friedmann background, $H = \dot{a}/a$, reads $\dot{\rho} + 3H(1+w)\rho = 0$, so

$$w = -1 - \frac{1}{3}\,\frac{\dot{\rho}}{H\rho} = -1 - \frac{1}{3}\,\frac{d\ln\rho}{d\ln a}.$$

This is the definition used above, now traced to its source. The substrate's own per-edge law is $\dot{\rho} = S(a) - \Gamma\rho$, genesis source minus dilution, with $\Gamma = 1 - \eta$ and $S(a) = S_0\, a^{-3}$. Substituting,

$$1 + w = -\frac{1}{3}\,\frac{S(a) - \Gamma\rho}{H\rho}.$$

Step 4. The single crossing, forced. Since $H, \rho > 0$, the sign of $1 + w$ is the sign of $-(S - \Gamma\rho)$:

$$w < -1 \iff S > \Gamma\rho \iff \rho < \rho^* \equiv S/\Gamma, \qquad w > -1 \iff \rho > \rho^*,$$

with $w = -1$ exactly at $S = \Gamma\rho$. Early, the strain sits below its attractor ρ^* and fills toward it ($S > \Gamma\rho$): phantom, $w < -1$. The attractor $\rho^* = S_0 a^{-3}/\Gamma$ recedes as the universe grows, the rising strain overtakes the falling attractor exactly once, at $S = \Gamma\rho$, and stays above it after: quintessence, $w > -1$, today. One crossing of $w = -1$ from below, with no parameter placed to put it there. The crossing epoch is fixed by Γ/H_0, hence by η; only its redshift, not its existence, awaits calibration.

7.2 The Fifth Force

The invisible weave that fills the universe is not distant. It passes through every atom in your body. Near an atomic nucleus, where matter is densely packed, the weave is locally compressed. That compression pulls on the electrons nearby, creating an interaction that has nothing to do with ordinary electricity or gravity. A fifth pull, distinct from the four known forces.

This pull means the 95 percent of the universe we cannot see is not separate from us. It touches ordinary matter at the smallest scale. Every atom carries a trace of the dark network in its energy levels. The heavier the nucleus, the stronger the trace. The closer the electron orbits, the more it feels.

The way to look for it: trap individual atoms, measure their energy levels with extreme precision, compare different versions of the same element. The differences would reveal the pull. The shape of this prediction, a tensor pattern across magnetic sublevels, is specific and falsifiable on its own; its amplitude depends on how small the underlying network is, and at a Planckian network scale it falls far below present precision, so a null result bounds that scale rather than refuting the mechanism.

The dark tensor at nuclear scales

The symmetry-breaking potential contains a coupling $\kappa\, \Omega_{\mu\nu}\, \bar{\Psi}\gamma^{(\mu}\nabla^{\nu)}\Psi$ between the dark tensor Ω and fermionic matter. This term lies outside the Standard Model gauge group and constitutes a fifth fundamental interaction.

Computational status (resolved). The mediator mass and the effective coupling have now been computed from the automaton. The dispersion of the Ω response carries a gap equal to the negative logarithm of the dilution factor, m_Ω = -ln η. This is a dimensionless substrate-unit gap; the physical mediator mass is (-ln η) ħc/l_fund, of order the Planck mass for a Planckian l_fund and so far above the window isotope-shift experiments probe. The gap locks the range to the substrate dilution rather than leaving it free; the resulting range is sub-lattice, so the force is contact-like. The coupling κ_eff is dimensionless, because Ω is a dimensionless information anisotropy while the matter operator carries the mass dimension of a stress tensor; its measured value is g_Ω ≈ 0.2. The absolute magnitude of any observable does not follow from the automaton alone but is set by the fundamental substrate length l_fund through a contact suppression of order (l_fund/a_0)^p, with p between 2 and 3. At a Planckian l_fund the isotope-shift signature falls tens of orders of magnitude below present precision, and observability would require l_fund near the PeV scale. The rank-2 shape is fixed and falsifiable independently of l_fund; the magnitude turns the experiment into a measurement of, or a bound on, l_fund.

A static nucleus with baryon number A sources Ω externally in Yukawa form:

$$\Omega(r) = \frac{\kappa_{\text{eff}} \cdot A}{4\pi r}\, e^{-m_\Omega r} \tag{5}$$

Scaling laws: energy corrections grow linearly with A (isotope-specific), are dominated by $\ell = 0$ orbitals (nuclear penetration), and grow as Z^3 across elements. Across a chain of isotopes at fixed Z this linear dependence is equivalently a dependence on neutron number, matching the electron-neutron coupling that isotope-shift searches constrain. The mediator mass m_Ω sets the range; the automaton fixes it to $m_\Omega = -\ln \eta$, a sub-lattice (contact-like) range corresponding to a heavy mediator of order $\hbar c / l_{\text{fund}}$. This lies far above the [100 eV, 10 MeV] window that isotope-shift experiments currently probe, so the interaction acts as an effective contact term whose amplitude is governed by l_{fund} (Figure: fifth-force landscape). It remains absent from macroscopic gravity experiments.

Because $\Omega_{\mu\nu}$ is rank-2 (not scalar), deformed nuclei produce anisotropic corrections that break magnetic sublevel degeneracy without external fields. A scalar mediator cannot produce this signature.

The Fifth Force: Full Derivation

Source Term

The coupling $\kappa\, \Omega_{\mu\nu}\, \bar{\Psi} \gamma^{(\mu} \nabla^{\nu)} \Psi$ in the symmetry-breaking potential sources the dark tensor field locally wherever baryonic matter is concentrated. For a static, spherically symmetric nucleus with baryon number A, the linearized field equation reduces to:

$$\nabla_\rho \nabla_\sigma R^{\rho\mu\sigma\nu}(\Omega) = \kappa_\Omega \Lambda_\Omega \Omega^{\mu\nu} + \frac{\kappa_\Omega \kappa}{2} \bar{\Psi} \gamma^{(\mu} \nabla^{\nu)} \Psi$$

Yukawa Solution

Outside the nuclear radius $r > R_N$, the spherically symmetric solution is:

$$\Omega(r) = \frac{\kappa_{\text{eff}} \cdot A}{4\pi r}\, e^{-m_\Omega r}$$

This is Yukawa form with range $\lambda = 1/m_\Omega$. The coupling to electrons produces energy corrections:

$$\Delta E_{nlm} = \kappa^2 \cdot \frac{\kappa_\Omega A}{4\pi} \int_0^\infty |\psi_{nlm}(r)|^2 \cdot \frac{e^{-m_\Omega r}}{r} \cdot |\nabla \psi_{nlm}(r)|\, r^2\, dr$$

Testable Predictions

1. Linear isotope scaling: $\Delta E(A_2) - \Delta E(A_1) \propto (A_2 - A_1)$. The shift scales with baryon number difference, not nuclear radius $A^{1/3}$ or surface $A^{2/3}$.

2. s-orbital dominance: $\Delta E_s/\Delta E_p \approx (Z/n)^2 \gg 1$. Electrons that penetrate the nucleus feel the force most strongly.

3. $\mathbf{Z^3}$ growth: $\Delta E \propto Z^3 \cdot A$. Heavier elements show dramatically larger effects.

4. Mediator mass, fixed (not a free window): the automaton fixes $m_\Omega = -\ln \eta$, a sub-lattice (contact-like) range that lies far above the [100 eV, 10 MeV] window isotope-shift experiments probe. The experiment therefore bounds the substrate length l_{fund} that sets the amplitude, not m_Ω itself.

5. Tensor anisotropy: In deformed nuclei the components of Ω differ along the principal axes, breaking m_l degeneracy without an external magnetic field. This rank-2 signature is one that no scalar or vector mediator can produce. By the Wigner-Eckart theorem, a traceless symmetric rank-2 operator Ω_{ij} acting on a state of definite angular momentum j (the orbital l for an LS-resolved level, or the total J for a fine-structure level) yields an energy shift proportional to $3m^2 - j(j+1)$, with m the magnetic quantum number of j; for a fine-structure level the orbital tensor is projected onto J, with a nonzero $6j$ factor setting the scale. The predicted sublevel pattern is:

$$\Delta E(m) = \kappa_{\text{eff}} \cdot f(r) \cdot \left[3m^2 - j(j+1)\right]$$

where $f(r)$ is the radial overlap of the electron wavefunction with the nuclear Ω profile. For the experimentally relevant transitions:

- **Orbital $l = 1$ (m_l resolved):** $\Delta E \propto [1, -2, 1]$ for $m_l = [-1, 0, 1]$, a triplet with the central component shifted down.
- **Orbital $l = 2$ (m_l resolved):** $\Delta E \propto [6, -3, -6, -3, 6]$ for $m_l = [-2, -1, 0, 1, 2]$, a quintuplet with alternating sign.

The ${}^2S_{1/2} \to {}^2D_{5/2}$ electric-quadrupole transition in Ca^+ (729 nm) currently measured by Wilzewski, Aude Craik et al. (Phys. Rev. Lett. 134, 233002, 2025) is a concrete target. The ${}^2D_{5/2}$ level has total angular momentum $J = 5/2$, so its six m_J sublevels ($m_J = \pm 1/2, \pm 3/2, \pm 5/2$) carry the pattern $3m_J^2 - J(J+1) \propto [10, -2, -8, -8, -2, 10]$; examining the transition for this m_J-dependent splitting is a direct, low-cost extension of an experiment already running.

The current state of the art reaches 100 mHz and finds a nonlinearity that is compatible with Standard-Model nuclear effects, namely the second-order mass shift and nuclear polarization, and sets bounds on the coupling of a scalar electron-neutron boson across the [100 eV, 10 MeV] mass window. That analysis assumes a scalar mediator and is therefore blind to the m_l structure predicted here. Because the measured transition is itself a quadrupole transition, examining its sublevel pattern for an m_l-dependent shift is a concrete and low-cost extension of an experiment already running, and is the framework's strongest differentiator from scalar fifth-force models.

6. Experimental protocol. The most promising target is the isotope shift in singly charged ytterbium (Yb) or strontium (Sr) ions, where optical clock transitions have been measured to 10^{-19} precision. Comparing transitions dominated by s-orbitals across isotope chains (e.g. ^168Yb, ^170Yb, ^172Yb, ^174Yb, ^176Yb) searches for an A-linear deviation from the expected King-plot linearity. The tensor signature can be isolated by measuring magnetic sublevel splittings in deformed isotopes without external magnetic fields. A positive detection would measure both m_Ω (from the range-dependence across orbital types) and κ_eff (from the overall amplitude). A null detection at current precision would constrain κ_eff/m_Ω^2, narrowing the parameter window.

Connection to Cosmological Scales

The same $\Omega_{\mu\nu}$ field at three scales:

- Cosmological (Gpc): trace = uniform tension = dark energy.
- Galactic (kpc): traceless = directional clustering = dark matter.
- Atomic (fm-pm): locally sourced by nuclei = fifth force.

One field. Three resolutions. Three previously unconnected observations unified by a single tensor.

---

8. Why 3+1 Dimensions

8.1 Dimensions & Particles

The Thread Network produces particles through its own geometry. What a particle is, and why the space it inhabits has exactly three dimensions, follow from the same underlying structure.

Particle Manifestation

Different thread configurations correspond to different physical phenomena.

Some phenomena are pure excitations in the network. They travel through the medium without forming bundles. Without bundling, they experience no mass. Without mass, the connectivity that manifests as gravity does not constrain them, though they still follow the network's gradients and curvature. What physics calls photons and gluons belong to this class. They travel at the network's maximum propagation speed, *c*. All observers measure the same *c* because they are all embedded in the same network with the same limit.

Other phenomena are stable bundles: threads knotted together by repeated interactions, persisting as coherent structures. These have mass. Their inertia is proportional to their bundling density. What physics calls electrons and other matter particles are examples. A particle's bundling is a local rise in entanglement density, the same ρ_E the metric curves around in Section 5. Once stable knots form, their differences and binding are what raise ρ_E locally, so matter sources curvature as the particle-scale face of the gravity already derived, not a separate mechanism.

Between these extremes lie intermediate configurations. Some bundles are tight enough to carry mass, but too loose to persist indefinitely. These correspond to what physics calls W and Z bosons: massive, but transient. The tightest, most stable bundles form what physics calls fermions: matter in its most persistent form.

Why 3+1 Dimensions

A thread is a one-dimensional object. A bundle of threads can form closed loops. Whether those loops can form persistent tangles depends on how many spatial dimensions exist.

Two dimensions. A thread cannot knot in two dimensions. Every closed loop shrinks to a circle without cutting or tearing, and no crossing can lock, so threads form circles but never knots. No persistent tangled structure. No mass.

Four or more dimensions. In four dimensions there is room to lift every crossing into the extra dimension and pull the thread free, so every knot unties. Again: no persistent structure, no mass.

Three dimensions. This is the only number of spatial dimensions where one-dimensional threads can knot and the knots stay tied. Two dimensions for the plane of the crossing, and a third to record which thread passes over which. Close the loop in three dimensions and the knot is locked, it cannot be untied without cutting the thread. This is a classical theorem in geometric topology.

Mass is tangled thread bundles. If knots cannot form or cannot persist, mass cannot exist as a stable property. The framework therefore predicts exactly three macroscopic spatial dimensions. No extra ingredient: threads are one-dimensional by construction, mass is persistent tangling, and tangling requires exactly three dimensions.

The same logic sets time to one dimension. Each event along a thread has one successor, the next interaction. There is no branching of identity. Time is the sequence of distinguishable events, and a sequence has one direction and one dimension.

The physicist Paul Ehrenfest showed in 1917 that stable orbits, planets around stars, are only possible in three spatial dimensions. In two dimensions, gravity would not permit closed orbits. In four or more, orbits are unstable. Ehrenfest's argument is independent of the knot argument, but both reach the same conclusion from different premises. The argument adds a further layer: stable orbits and persistent knots are consequences of the same underlying constraint on spatial dimensionality.

Dimension Theorem. See Full Derivation below for the formal statement. The core result: closed 1-dimensional curves admit nontrivial isotopy classes (knots) only in exactly three spatial dimensions, and stable massive matter therefore requires $d = 3$.

Why time is one-dimensional. Identity Edges define unique successors (Axiom 6: minimality of threads). Each event carries one forward arrow. The local structure is $\mathbb{Z}$ (discrete sequence), coarse-grained to $\mathbb{R}$.

Threads select (3+1). The pair is forced: 3 is the minimal dimension for knots; 1 is the dimension of a chain with unique successors. No other pair is consistent with stable massive identity.

Ehrenfest (1917) corroboration. Stable Kepler orbits exist only for inverse-square force laws in three spatial dimensions. Ehrenfest derived this from the scaling of centrifugal barriers: for potential $V \propto 1/r^{d-2}$, circular orbits are unstable when $d \geq 4$, and for $d = 2$ there is no bounded orbit. The orbit argument and the knot argument are independent probes of the same constraint.

Extra dimensions. Any additional spatial dimensions must be compactified below the thread interaction scale, or knots, and therefore mass, would dissolve. Compactified dimensions are not forbidden; the argument explains only why three are macroscopic.

Full Derivation: Dimension: Why 3+1

Theorem 6.1 (Knots need three dimensions). Closed 1-dimensional curves admit nontrivial isotopy classes (knots) only in exactly three spatial dimensions. In two dimensions curves cannot cross; in four or more every knot unties (there is room to pass strands through the extra dimension).

Theorem 6.2 (Mass forces d = 3). If mass is energy stored in topologically protected thread configurations, then stable massive matter exists only where 1-dimensional threads can be topologically protected, hence only in three spatial dimensions. In d = 2 there is no tangling at all; in d ≥ 4 every tangle relaxes and no particle is stable. The universe has three spatial dimensions because that is the only number in which threads can hold a knot.

Time is one-dimensional. Identity edges give each event a unique successor: a total order per thread. A second time dimension would mean branching succession, contradicting persistence of identity as defined. One time dimension is not chosen; it is what "the same thing, again" means.

Together: 3 + 1, with the spatial part forced by knot theory and the temporal part by the definition of identity. Corroboration, independent of threads: stable planetary orbits and stable atoms exist only for three spatial dimensions (Ehrenfest).

Full Derivation: Particles and Spin

Particle types are topological classes of thread configurations on the three-dimensional spatial slice forced above:

Free arcs with no tangle are massless and travel at c, with nothing internal to reconfigure. Framed knotted bundles are massive and persistent, topologically protected by the three-dimensional knotting forced above. Loose transient tangles are unstable and decay.

Spin from framing. Threads in 3 dimensions are naturally framed (they carry twist, by CWF). For framed objects, a 2π rotation is not isotopic to the identity while 4π is (the Dirac belt trick): the topology of framed tangles realizes the double cover that defines spin-1⁄2. CWF (Lk = Tw + Wr) forces a 2π framing twist to trade against writhe, a discrete sign-like exchange.

Spin-statistics connection. Exchange of two identical particles in 3D is homotopic to a 2π rotation (theorem in configuration space topology). Combined with the belt trick: exchange → 2π rotation → phase -1 for framed tangles → fermions. Unframed tangles → bosons.

The Standard Model's specific groups and masses remain underived and are listed as open.

---

9. Gauge Group Emergence

The gauge symmetries of the Standard Model, SU(3) × SU(2) × U(1), are not postulated in the Thread Framework. They emerge from the quantum information geometry of interacting qubits in the thread graph.

9.1 SU(3) from Three-Qubit I-Configuration Space

Consider three qubits in an equilateral triangle. Their mutual information values (I_AB, I_BC, I_CA) form a 3-dimensional configuration space. The equilateral configuration (all I equal) is the GHZ-class ray; the W-class covers the 2-dimensional deviation space densely.

Proposition (SU(3) Emergence, conditional). SU(3) acts on the 8-dimensional 3-qubit state space $(C^2)^3$. The three I-values (ℐ_AB, ℐ_BC, ℐ_CA) are invariant labels under this action, and the equilateral configuration (all ℐ equal) is the GHZ-class ray. The permutation group S_3, the Weyl group of SU(3), acts on the I-values by permuting the three qubit labels.

Cartan classification of rank-2 simple Lie groups selects SU(3) uniquely: A_2 has Weyl group S_3, while B_2 and C_2 have D_4 and G_2 has D_6. Numerical verification (`su3_cartan.py`) gives the deviation-space Fisher metric isotropic to eigenvalue ratio 1.000000, and Schur's lemma forces it: the S_3 standard representation on the deviation plane is the 2-dimensional irreducible representation, so any S_3-invariant bilinear form on it is unique up to scale, hence isotropic. The same mechanism protects rotational isotropy via the O_h point group (§11).

Status. The Cartan step is mathematically complete. The derivation is conditional on the automaton's IR fixed point being in the sl3 universality class, an RG premise currently under numerical investigation. The automaton preserves Z3 cyclic symmetry (2.2% relative standard deviation at L = 6); the full S3 Weyl group and the continuum SU(3) require verification at larger lattice sizes. Lattice/IR emergence of the full S_3 Weyl group is not established: it holds for an isolated, maximally symmetric triangle, but on multi-triangle patches (K_4, strips) finite-size and boundary asymmetry break per-plaquette S_3 (`su3_adversarial.py`, `su3_diag.py`). The Cartan/Schur content above stands; the lattice/IR step is open.

9.2 Quantum Deformation Parameter

The quantum deformation parameter is q = e^{-⟨I⟩} ≈ 0.31, where ⟨I⟩ ≈ 1.17 nats per edge is measured from the automaton (b7). This is a generic real value, not a root of unity. For generic q > 0, q ≠ 1, the quantum group U_q(sl3) is isomorphic to the universal enveloping algebra U(sl3) as a Hopf algebra (Drinfeld 1985). The representation theory is that of ordinary SU(3), no tuning to a root of unity is required.

Throughout, the quantum integer is [n]*q* = *Σ*.}^{n-1} q^{n-1-2k

9.3 U(1) and SU(2)

U(1). Thread orientation (winding number) is additive under concatenation of threads, yielding U(1) group structure. Electromagnetic gauge symmetry emerges from the invariance under reparametrization of the winding phase.

SU(2). Two interacting threads form a 2-strand braid. With framing (twist), the Temperley-Lieb algebra yields SU(2) representation theory in the continuum limit. The braid-to-SU(2) route is the standard correspondence and stands on its own; what remains open is numerical confirmation that the automaton realizes this structure.

9.4 Generations and Open Questions

Three generations are suggested by the 3-strand braid group B_3, whose representation theory supports several distinct configurations. What is established is the structural source, a braid group rich enough to carry more than one; deriving exactly three, and which mechanism (braiding planes versus twist levels) selects them, is under investigation. The gauge group derivation is the framework's most ambitious open problem: the Cartan classification step is complete and stands, while the dynamical RG step that would close it is work in progress.

---

10. Quantum Dynamics and Measurement

10.1 Unitary Evolution: The Schrödinger Equation

Before an interaction registers an outcome, the state evolves, and that evolution is not a separate postulate. It is forced by the same inequality that fixes everything else. A rule that let distinguishability grow would run the master inequality backward, so the only admissible dynamics is one under which relative entropy to the fixed-point state never increases. Locality confines the rule to geometrically near edges, and thread persistence keeps causal order.

Under these three constraints the coarse-grained dynamics is determined: it carries any state toward the fixed point along the steepest descent of relative entropy, and the automaton of the computation section is one realization, with universality guaranteeing that every realization shares the same long-wavelength physics.

A single particle is a thread, or a small bundle, spread as a superposition over the vertices of the regular graph. At each step its amplitude hops to neighboring vertices, which is the graph Laplacian acting on the amplitude. As the net densifies and the step shrinks, that discrete hopping becomes the Laplace-Beltrami operator on the emergent manifold, the kinetic term of the wave equation. Plain hopping would give diffusion, a classical spreading. What makes the evolution quantum rather than diffusive is a single factor of i, and that factor has a home: near the fixed point the rule splits into a part that dissipates and a part that rotates, and the rotating part is a commutator with an effective Hamiltonian. In the near-equilibrium limit the rotation dominates and the evolution is unitary, while the residual dissipation is the small correction that later reads as measurement. Mass enters through entanglement density: a heavier bundle carries more internal entanglement and more hopping paths, so it is harder to displace, and the diffusion constant falls as one over the mass. The scale of action is the energy of one interaction times the duration of one step.

Full Derivation: Unitary Evolution

Let L be the Liouvillian, $L = (1/\delta t)\log\Phi$ for the one-step map Φ. The keystone requires Φ to contract relative entropy, so $\frac{d}{dt}D(\rho(t)\|\rho^*) \leq 0$: relative entropy is a Lyapunov function, and $L'(\rho^*)$ has non-positive spectrum. Locality (A2′) makes L a finite-range sum $L = \sum_i L_i + \sum_{\langle i,j\rangle} L_{ij}$ in the emergent geometry (Lieb-Robinson). In the continuum limit the contraction takes the gradient-flow form $L = -\Gamma\,\mathbf{grad}_{\mathrm{Fisher}}\,D(\cdot\|\rho^*)$; any two local Liouvillians with the same fixed point agree in their low-energy effective action (the difference is RG-irrelevant), so the automaton represents the unique universality class.

For a single particle the graph Laplacian $\Delta_{\mathrm{graph}}\psi(i) = \sum_{j\sim i}(\psi(j)-\psi(i))$ gives, in the continuum limit,

$$\partial_t\psi = D\,g^{\mu\nu}\nabla_\mu\nabla_\nu\psi, \qquad D = \varepsilon\,\ell_{\mathrm{fund}}^2/\delta t.$$

The linearized Liouvillian at the fixed point splits into a rotating and a dissipative part,

$$L'(\rho^*)\,\delta\rho = -i[H_{\mathrm{eff}},\delta\rho] + \mathcal{D}(\delta\rho).$$

Near equilibrium the dissipative part is subleading and the rotation dominates, which is the imaginary diffusion constant $D = i\hbar/2m$ of quantum mechanics. With mass as internal distinguishability density $m = (1/c^2)\int_\Sigma \rho_E\,dV$ (so $D \propto 1/m$) and $\hbar$ set by the energy per interaction times the step δt,

$$i\hbar\,\partial_t\psi = -\frac{\hbar^2}{2m}\,g^{\mu\nu}\nabla_\mu\nabla_\nu\psi + V(x)\,\psi, \qquad V = m\Phi,$$

the Schrödinger equation for a particle in the emergent geometry with gravitational potential Φ. The kinetic term, the relation $D \propto 1/m$, and the scale $\hbar$ follow from the construction. The factor of i rests on the near-equilibrium limit, where the rotating part of the linearized Liouvillian dominates the dissipative part, a dominance assumed rather than proved. The gauge couplings of the Standard Model are not yet reached (Section 9).

10.2 Measurement and Collapse

Measurement is not a mysterious process requiring a classical observer. **Every thread interaction is a measurement.** When threads interact, they produce distinguishability. As more environment threads interact, distinguishability increases. The apparent collapse is that an interaction point can only access states distinguished by its own interaction history. The full superposition still exists from an external perspective, one not yet entangled with this set of threads.

The von Neumann chain and Wigner's friend

In quantum measurement, the von Neumann chain asks: where does the measurement end? If a particle interacts with a detector, and the detector with a computer, and the computer with a human observer, at which link does the superposition become a definite outcome? Von Neumann showed the chain can be placed anywhere without changing the predicted probabilities. The framework answers: **everywhere.** Every thread interaction is a measurement. No privileged conscious endpoint is needed. The apparent collapse is that an interaction point can only access its own interaction history.

Wigner extended the puzzle with a thought experiment. A friend measures a particle and sees a definite outcome. Wigner, outside the laboratory, has not yet interacted with the friend. From Wigner's perspective, the friend and particle remain in a joint superposition. Who is right? Both. From the friend's perspective, the particle has a definite state, the friend's threads have interacted with the particle's threads and produced distinguishability. From Wigner's external perspective, not yet entangled with the friend's threads, the superposition remains intact. Both perspectives are valid. Reality is the network of all perspectives. The delayed-choice quantum eraser reads the same way: whether the which-path record counts as erased is not an absolute fact but one relative to a vantage, fixed by which interactions have woven the record in. Two vantages that have stitched different links can each be right while disagreeing about whether the interference was ever recoverable. The eraser is observer-relative, not absolute, which an experiment can check.

Distinguishability, decoherence, and the Born rule

A concrete mechanism: measurement as basin selection (§11.9). The driven dissipative dynamics (§11.9) can possess multiple attractor basins. This makes a definite, testable proposal possible: a measurement outcome is the basin into which a state flows under the dynamics, and the Born weight of outcome k is the volume of its basin of attraction, basin-volume(k) = $|\langle k|\psi\rangle|^2$. The monotone flow supplies the irreversibility (the apparent collapse), while the basin structure supplies the discreteness of outcomes. This was not testable while only the trivial attractor existed; with the driven live structure it now is, and is pre-registered with a negative outcome (volumes $\neq |\langle k|\psi\rangle|^2$) as a valid result.

Measurement is the creation of distinguishability through thread interaction. When threads interact, they produce mutual information.

$$\rho_{\text{obs}} = \mathbf{Tr}_{\text{env}}\, |\Psi\rangle\langle\Psi|$$

Decoherence is environment thread interaction: more interacting threads means more distinguishability. The apparent collapse is that an interaction point can only access states distinguished by its own interaction history. The full superposition still exists from an external perspective.

What decides that distinguishability occurs at all: magic. Non-Clifford operations make the thread encoding approximate. The resource condition: $M > 0$ implies distinguishability emerges. In a perfect stabilizer code ($M = 0$), no distinguishability is produced. The self requires imperfection.

What decides how much: envariance (Zurek). Born probabilities follow from the swap symmetries of entangled states:

$$|\Psi_{SE}\rangle = \sum_k \sqrt{p_k}\, |s_k\rangle|e_k\rangle$$

A swap $U_S|s_k\rangle = |s_{\pi(k)}\rangle$ leaves $|\Psi_{SE}\rangle$ invariant only if $p_k = p_{\pi(k)}$. For equal amplitudes, $p_k = 1/D$, the Born rule. The counting of equal-amplitude fine-grainings is exactly the counting of thread configurations accessible from a given interaction point. The framework adopts envariance as its Born route. Decoherence explains why interference terms vanish in the pointer basis. Why one specific outcome is experienced, the single-outcome problem, turns on the nature of the observer and is taken up in the companion work on structured perspective. The measurement content here is physical: every thread interaction degrades quantum coherence by creating mutual information, and that degradation is what is conventionally called decoherence.

The data-processing inequality, that mutual information cannot increase under local operations, is the master inequality evaluated at the moment of interaction.

Full Derivation: Measurement

- Every interaction writes I: every interaction is a measurement.
- Collapse is access: an interaction point reads only its own history (A3). Superposition persists for every perspective not yet woven in. Wigner and his friend both report correctly.
- Pointer states are the I-stable configurations: those whose records in the environment are redundantly copied and stable under further

interaction.
- **Born rule, divided correctly.** Probabilities follow from envariance: the swap symmetries of entangled states force equal probabilities for equal amplitudes (Zurek). Thread-configuration counting is exactly that fine-graining count; the framework adopts envariance as its Born route. The separate claim that magic is the resource that makes an interaction produce distinguishability at all remains conjectural.

Derivation chain: the Born rule via envariance

Step 1. For an entangled state |ψ⟩ = Σ_k √p_k |k⟩|E_k⟩, a swap of two system branches can be exactly undone by the counter-swap on the environment: the global state is invariant under the pair.

Step 2. Probabilities assigned locally to the system cannot depend on anything the environment swap changes; for equal Schmidt coefficients the branches are therefore equiprobable.

Step 3. Unequal rational coefficients reduce to the equal case by fine-graining the environment; continuity extends to the reals:

$$P(k) = |\langle k|\psi\rangle|^2 = p_k$$

The Born rule is recovered, not postulated.

Step 4. The added claim: produced distinguishability requires magic (a stabilizer-only interaction registers nothing new). This is conjectured.

---

11. Computational Verification

The account is not only conceptual. It includes a working computational model: the interaction automaton. This is a discrete simulation of threads on a graph, evolving through three operations, interact, dilute, repeat, with one conjectured parameter (the dilution factor η, measured in the range 0.63-0.70).

What the Automaton Does

The automaton begins with a cubic lattice. At each step, threads interact if their mutual information exceeds a threshold. The interaction updates the state. The graph dilutes: edges below threshold are removed. The process repeats. Over many steps, spatial structure is preserved and refined; the effective dimension settles low (numerically 2-3), the dynamics pulling high-dimensional blobs down to a connected low-dimensional manifold, though the exact value is not yet uniquely pinned.

The automaton is not a simulation of physics on a pre-existing spacetime, but a simulation of what happens when distinguishability is the only primitive and the master inequality is the only law. Spacetime is what comes out.

Key Results

M1: Dimensional convergence. Starting from a cubic lattice, the automaton's effective dimension scales toward 3. Measured across three lattice sizes:

Lattice	Vertices	d_eff
6^3	216	1.63
10^3	1000	1.95
16^3	4096	2.15

The effective dimension follows the finite-size scaling $d_{eff} \approx 3 - 3.4/\sqrt{L}$ (extrapolated → 3 as $L \to \infty$). A two-parameter fit with the asymptote left free returns 2.98 ($R^2 = 0.998$) across the three sizes, so the limit 3 is recovered from the data rather than imposed (on three points; more sizes would tighten it). The measured values (1.63-2.15) show convergence toward 3 but do not reach it at accessible lattice sizes. The result is consistent with dimensional emergence; a universality-class hypothesis remains to be demonstrated analytically.

M2: Area-law entanglement. When the automaton is partitioned into spatial regions, the entanglement entropy between a region and its complement scales with the area of the boundary, not the volume. This is the required scaling for gravitational entropy and follows from the locality built into the interaction rule.

M7: Waterfall spectrum. The automaton's interaction cascade produces a nearly scale-invariant spectrum of fluctuations. Cascade simulations are currently limited to small lattices ($L \ll 64$), yielding only 3 usable k-bins and no convergence toward the observed $n_s = 0.965 \pm 0.004$. Numerical confirmation of the spectral index is work in progress.

P2.2: Magic emergence. Starting from a stabilizer state under Clifford-only evolution, the automaton with the full Clifford+T gate set generates non-Clifford resource. Magic is measured operationally as the difference between the effective stress-energy under full gate-set evolution and under Clifford-only evolution: $E_{magic} = E - E_{stab}$.

The automaton is not merely a simulation of physics on a pre-existing spacetime. It is a quantum computation. The bifurcation operator creates qubits. The interaction rule, CNOT + Hadamard + T-gate, is universal for quantum computation. The automaton's evolution produces geometry, gauge symmetries, and the dark sector as byproducts of its operation. The laws of physics are the rules of the computation. They are not imposed from outside, they are discovered by the graph as it runs. The stability of physical law follows from steady-state dynamics: the effective dimension, the mean mutual information, and the gauge group settle to attractor values, sustained by ongoing genesis feeding the network, a driven steady state, not a source-free equilibrium (§11.9).

Caveats

The automaton is a representative of a conjectured universality class; it is not proven to be the unique consequence of the axioms. The gradient flow that drives the dynamics is the unique candidate, not the unique consequence. The continuum limit has not been rigorously established. These are model-level concerns, not axiom-level concerns. The axioms survive or fall independently of the automaton. The automaton is evidence, not proof.

Full Derivation: The Interaction Automaton

11.1 Algorithm

The automaton is implemented in Python (numpy, scipy). The core interaction loop:

```
class ThreadAutomaton:
 def __init__(self, N=1000):
 # N thread crossings, each a qubit
 self.N = N
 self.states = [Qubit() for _ in range(N)]
 self.I_matrix = {} # I[(i,j)] = mutual information in nats

 def step(self):
 # 1. Compute geometry from current I
 d_geo = self.compute_geometry()

 # 2. Determine interactions
 pairs = self.select_pairs(d_geo)

 # 3. Apply entangling gates (Clifford + T = universal)
 for i, j in pairs:
 self.apply_CNOT_H_T(i, j)
 self.I_matrix[(i,j)] = np.log(2)

 # 4. Dilution
 ETA = 1 - 1/np.e # conjectured value; measured range 0.63-0.70
 for (i,j) in list(self.I_matrix.keys()):
 self.I_matrix[(i,j)] *= ETA
 if self.I_matrix[(i,j)] < 1e-6:
 del self.I_matrix[(i,j)]

 def compute_geometry(self):
 # Edge lengths from I
 l = {}
 I_max = np.log(2)
 for (i,j), I_val in self.I_matrix.items():
 l[(i,j)] = -np.log(I_val / I_max) if I_val > 0 else np.inf
 # Geodesic distances (Floyd-Warshall or Dijkstra)
 return self.geodesic_closure(l)
```

M1: Dimensional convergence. The effective dimension d_{eff} scales toward 3 with finite-size extrapolation $d_{\text{eff}} \approx 3 - 3.4/\sqrt{L}$. Measured values (1.63–2.15) show convergence but do not reach 3 at accessible L.

M2b: Spontaneous dimension. Starting from a single node with repeated bifurcation, the spectral dimension d_s stabilizes at approximately 3.1-3.8 across multiple random seeds.

M7b: Universality. Two different interaction rules (deterministic highest-$\mathcal{I}$ neighbor and probabilistic softmax) both produce $d_s \approx 3$, consistent with a universality class.

The convergence of d_eff toward 3, the stabilization of ⟨I⟩ near 1.17 nats per edge, and the emergence of SU(3) gauge structure are attractors of the driven dynamics: the graph self-organizes toward these values while genesis continues to feed it (§11.9). The laws of physics are not imposed on the graph; they are discovered by it. The stability of physical law across cosmic time is a consequence of fixed-point dynamics: once the network reaches its steady state, macroscopic physics becomes stationary.

The stochastic AR(1) dynamics of the automaton, I_{t+1} = η·I_t + log2·X_t, have a further consequence for the nature of the vacuum. What quantum field theory describes as "vacuum fluctuations" are the Bernoulli innovations X_t at each edge. The variance floor Var* ≈ 0.20 nat2 per edge is the Planck-scale granularity of the I-field, not a sum over virtual particles. The vacuum is not empty, it is the steady state of a dynamical information network. There is no zero-point catastrophe because there is no sum over zero-point modes.

P2.2: Magic emergence. Starting from a stabilizer state, the automaton spontaneously generates non-Clifford resource.

The automaton is evidence for the framework's computational tractability, not proof of its uniqueness. Rotational Lorentz invariance is symmetry-protected by the O_h point group (the T-gate is a group singlet; every $O(k^2)$ correction is forced ∝ δ_ij). Free-level protection is confirmed by QCA inheritance (exact shared cone, even-order corrections only). Boost invariance across modes remains open (L1 velocity-convergence program).

Full Derivation: Computational Verification

11.5 n_s: Structure, Not a Number

Preliminary ensemble simulations (15 cascades per parameter set) show that the crossing-generation variance is consistent with a red tilt from collective exit, but the sample size is inadequate: 15 cascades give ~40% sampling error on the variance, ten times the effect size. A statistically meaningful measurement requires 500+ cascades with pre-registered estimator.

The generic 2/n_* term alone gives n_s ≈ 0.967, closer to observation than the current best simulation. The percolation structure is promising; the number is not yet computed.

11.7 Lorentz program: first numerical contact

Test A (L3, L2 at free level). A split-step Dirac automaton (the QCA class of L3) was solved exactly and numerically. Results: the massless dispersion is exactly conical (fitted corrections at machine precision, 10-13); for massive species the odd-order corrections vanish (k3 coefficient at fit-noise level 10-7 to 10-6) so lattice corrections enter at even orders only; and the strict causal cone is exactly one cell per step for every mass, measured at 1.0000 across species.

Unitarity plus locality yields an exact, shared Lieb-Robinson cone at the free level: L2 holds exactly there by construction, and the L3 protection class (violations confined to even, dimension-six-like orders) is reproduced. The open half is interactions, where the QCA literature has existing interacting constructions to map onto.

Test B (L1, velocity convergence). The first probe compared two path metrics on the cascade graph and was invalid by design: the metrics are trivially proportional along tree edges, so the constant ratio measured nothing. The redesign: extract per-mode velocities from the graph Laplacian dispersion (wave dynamics, not path metrics) on post-exit graphs of increasing size, and test whether mode velocities converge with scale.

Test C (L6, statistical isotropy). Spectral-embedding anisotropy showed no decay with size (exponent ≈ 0 against the predicted -0.5), but the probe is too crude: two-dimensional spectral coordinates of a tree-dominated graph do not define well-defined directions. Inconclusive; a correlation-function isotropy measure on the post-exit graph is the proper instrument.

Test D (L4, bias irrelevance). A frame bias of strength 0.3 imposed on a lattice I-field decays by ten orders of magnitude in sixty steps under isotropic local rebalancing: frame information is irrelevant when the microdynamics is direction-blind on average. Conditional: the result demonstrates the mechanism, and the remaining task is showing the interaction automaton's update rule is in this class.

Standing of the wall after first contact. At the free level the wall does not stand: the QCA class delivers an exact shared cone with even-order-only corrections, which is precisely the protection L3 claimed. The wall is now an interacting-theory problem, which is narrower than before.

11.8 Redesigns executed: L1-R and M1-R

L1-R (mode-velocity convergence, proper instrument). Two species were defined on the same cascade graph as wave dynamics on two Laplacians (uniform-weight and I-weight, each normalized to remove the trivial speed scale), and front velocities were extracted from wave-packet arrival times. The velocity ratio converges toward one with scale: 1.087 at 511 nodes, 1.046 at 2047, 1.038 at 8191. Infrared velocity convergence, the core of L1, shows the predicted trend at first contact. Single seed, three sizes: an ensemble and larger graphs are required before this is more than a trend.

M1-R (geometry test that can fail). The redesign demanded by 11.1 was executed: reinforcement-decay edge dynamics with a realistic pruning threshold (twelve percent of mean weight) over 120 steps, removing roughly a third of all edges. The graph does not fragment (reachable fraction 0.999 and 0.996 at L = 10 and 14) and the effective dimension is low and rises slightly with system size (2.51 then 2.63); the connected low-dimensional manifold is selected, though the value is not uniquely pinned to 3. The test is now falsifiable and passed.

11.9 Simulation findings on the dynamical core

A dedicated suite (`threadsim/code/`) tested the framework's dynamical claims directly. Each entry states what the simulation shows and what it bounds; several delimit an open problem rather than solve it.

The steady state is driven, not autonomous. For the map Φ = D_η(E_geo(·)) the source-free rule has a single attractor, the trivial product state (ΣI → 0, every η ∈ [0.5, 0.98], all-to-all and local-ring). The reason is structural: interaction strength p = exp(−d) vanishes as the mutual information it would replenish vanishes, while dilution (1−η) stays constant, a positive feedback toward the product state. Brief "live" transients occur (at η ≈ 0.7, ΣI peaks then collapses) but do not persist. A live, extended, low-dimensional, stationary phase is realized robustly in the classical growing automaton (the manifold phase, sustained against dilution by ongoing genesis); for the quantum density-matrix flow a clean live stationary fixed point at accessible n is not established (driven variants sit between thermal death and over-cooling). The bound stands either way: the emergent geometry is a driven, non-equilibrium phase sustained by ongoing genesis, not an autonomous equilibrium, and its existence and uniqueness remain open. (`quantum_flow*.py`, `manifold_phase.py`, `growing_flow.py`)

Geometry and locality are one invariant. Hamiltonian k-locality and I-graph geometry are preserved exactly by local unitaries and degraded together by entangling transformations: they are the same local-unitary invariant. Geometry therefore adds no factorization selection beyond Hamiltonian locality. (`factorization_*.py`)

Magic is required for a responsive geometry. Stabilizer states give a flat reduced spectrum and a trivial modular Hamiltonian K = −log ρ; non-local magic gives a non-trivial K and a responsive area operator, with the first law dS = d⟨K⟩ reproduced, background-independently, with no holographic dual. The driven steady state behaves consistently: an earlier reading in which T-gate magic appeared to reduce distributed structure was traced to two measurement artifacts, applying magic as a decohering mixture rather than a coherent gate, and conflating entanglement amount with responsiveness. With coherent magic and the response measured per unit structure, the relation holds away from equilibrium, the reduced spectrum going from exactly flat at zero magic to non-flat with magic, the signature of a responsive geometry, while magic lowers the entanglement amount the flow sustains rather than its responsiveness. (`magic_backreaction.py`, `modular_hamiltonian.py`, `magic_responsiveness.py`, `magic_coherent.py`)

SU(3): isotropy is forced, lattice emergence is open. The deviation-plane Fisher metric is isotropic because the S_3 standard representation is the 2-D irreducible, so Schur's lemma forces a unique invariant form (ratio 1.000000). Full S_3 at the lattice/IR level is not established: it holds for an isolated symmetric triangle but breaks on multi-triangle patches (K_4, strips) under finite-size and boundary asymmetry. (`su3_cartan.py`, `su3_adversarial.py`, `su3_diag.py`)

The fifth force is a contact-like probe of the substrate scale. The traceless tensor Ω sourced by a local mass excess was measured directly: it is rank-2 and radial, linear in the source mass, with a dispersion gap $m_\Omega = -\ln \eta$ that locks the range to the dilution factor and makes the force contact-like. The coupling is dimensionless, $\kappa_{eff} = g_\Omega \approx 0.2$, so the absolute King-plot magnitude is fixed by the substrate length l_{fund} through a contact suppression and is unobservably small at a Planckian l_{fund}. The rank-2 shape, an m_l-dependent quadrupole with neutron-number and Z^3/n^2 scaling, is scale-free and discriminates the model from any scalar mediator, and is grounded against the calcium King-plot bound of Wilzewski, Aude Craik et al. (2025). (`fifth_force_omega.py`, `omega_propagator.py`, `kappa_eff.py`)

Measurement as basin selection (testable proposal). Because the driven dynamics can support several attractor basins, a measurement outcome can be modeled as the basin a state flows into, with Born weight equal to basin volume, testable once a self-sustaining attractor structure exists. Stated as a proposal, not a result.

These are model-level results: with one exception they use minimal models and a faithful entangling interaction rather than the full automaton at scale, and the pruning step is not yet included.

---

12. Open Problems

- **Derive A4 (regularity) from thread dynamics.** The manifold's emergence rests on this.
- **The shortcut catalog.** Density and length distribution of anomalous edges compatible with observed Lorentz invariance.
- **Quantify Q versus C in warm neural binding.** the monogamy budget (Christandl-Winter 2004) applies to Q; the empirical size of Q in brains decides how much of the budget is theorem versus dynamics.
- **Complete the spin-from-framing program.** Ribbon classes to statistics, and any path from thread topology toward gauge structure.
- **Variational status of the selection principle max D(S).** Derive it from the underlying automaton dynamics (the companion mathematics) or replace it.
- **Dark-matter confrontation with the standard evidence.** The coincidence and relic-abundance problems of §7, and the mass-from-baryon separation in merging clusters and the CMB acoustic-peak structure, are not yet addressed.
- **Compute n_s from percolation.** Pre-register estimator, run 500+ cascades.
- **The continuum limit proof.** Rigorous convergence of the discrete automaton to the gradient flow.
- **The Standard Model.** Vacuum selection from framed ribbon tangles.

Simulation-tested premises (see §11.9). - *Fixed point:* the live extended geometry is a driven steady state sustained by ongoing genesis, not a source-free equilibrium. - *Preferred factorization (A0):* geometric entanglement gives no independent selection beyond Hamiltonian locality. - *Magic↔backreaction:* holds background-independently and in the driven steady state; with coherent magic the reduced spectrum becomes non-flat (responsive) there too, while magic reduces entanglement amount, not responsiveness. - *SU(3):* the Fisher isotropy is forced by Schur's lemma; lattice/IR S_3 emergence is open. - *Fifth force (κ_eff, m_Ω):* computed (§7.2). m_Ω = -ln η (contact-like, sub-lattice range); κ_eff = g_Ω ≈ 0.2 (dimensionless); the absolute magnitude is set by the substrate length l_fund, so the King-plot experiment bounds l_fund. Grounded against Wilzewski, Aude Craik et al., Phys. Rev. Lett. 134, 233002 (2025).

Map the interaction automaton to the interacting QCA constructions. The free-level inheritance is established (the Newton's-constant derivation; §11.7); the wall now stands only here. Existing interacting quantum cellular automata are the target; success closes the interacting Lorentz channel (the cubic l=4 channel requires separate RG analysis).

Prove the magic and fast-forward connection. Conjecture: a dynamics admits exponential fast-forwarding if and only if its magic generation rate vanishes. The Clifford direction is proved (symplectic powering); the Atia-Aharonov results bound the generic case; the equivalence would turn the manifestation boundary from thesis to theorem.

Appendix A: Dependency Audit

The Keystone Itself

Strong subadditivity (Lieb-Ruskai 1973): for any tripartite quantum state,

$$S(ABC) + S(B) \leq S(AB) + S(BC)$$

This is equivalent to monotonicity of relative entropy under partial trace, and via Stinespring dilation under every physical map:

$$D(\Lambda(\rho) \,\|\, \Lambda(\sigma)) \leq D(\rho \,\|\, \sigma)$$

Distinguishability never increases under processing. Everything in the audit below cites this one inequality.

The inequality $D(\rho\|\sigma) \geq D(\Phi(\rho)\|\Phi(\sigma))$ for every CPTP map Φ is the single non-redundant dependence of the framework. Everything else follows from this one statement plus the axioms.

Proof Status

Proved independently by Lindblad (1975, *Communications in Mathematical Physics*) and Uhlmann (1977). The proof uses the joint convexity of relative entropy, which in turn rests on the operator convexity of $x \log x$. No conjectures. No numerical approximations. The result is a theorem in the strict mathematical sense.

Equivalent Formulations

The monotonicity of relative entropy is equivalent to strong subadditivity of von Neumann entropy (Lieb and Ruskai, 1973):

$$S(\rho_{ABC}) + S(\rho_B) \leq S(\rho_{AB}) + S(\rho_{BC})$$

The master inequality is also equivalent to the data processing inequality for mutual information: $\mathcal{I}(A:B) \geq \mathcal{I}(A:B')$ when B' is obtained from B by any physical operation.

What It Carries

The inequality is load-bearing for:

- **Metric triangle inequality**: $\ell = -\log(\mathcal{I}/\mathcal{I}_{\max})$ satisfies $d(A,C) \leq d(A,B) + d(B,C)$ because SSA implies the required subadditivity of $\mathcal{I}$.
- **Gravity as an equation of state**: the entanglement first law $\delta S = \delta\langle K\rangle$ follows from positivity and monotonicity of relative entropy evaluated on causal diamonds (Jacobson 1995; Faulkner et al. 2014).
- **Dark energy equation of state**: the equilibrium value $w = -1$ follows from volume-extensivity of the residue; relative-entropy positivity underwrites the averaged null energy condition, which the driven solution satisfies while locally violating pointwise NEC in its phantom phase; the driven dynamics predict a single phantom crossing from below (Quintom-B, §11.9), with the crossing redshift uncalibrated.
- **Monogamy budget**: squashed entanglement E_{sq} satisfies monogamy because of SSA (Christandl-Winter 2004).
- **Markov blanket self**: the self-model boundary ∂X renders interior and exterior conditionally independent, $\mathcal{I}(A:C|B) \to 0$, saturating strong subadditivity (HJPW 2004). The companion paper develops this formulation.
- **Reflexive bound**: looking inward costs distinguishability because the self-model's processing is a CPTP map, and monotonicity bounds how much $\mathcal{I}$ survives.
- **Binding surplus**: the whole exceeds the sum of parts because $\mathcal{I}_{\mathrm{bound}} > \sum \mathcal{I}_{\mathrm{pairwise}}$ when the state is not separable, a direct consequence of SSA.

What It Does Not Require

The inequality does not assume Hilbert spaces (it holds for any von Neumann algebra), does not assume dynamics (it is a state property), and does not assume locality (it predates quantum field theory). This is why the framework can derive locality rather than assume it.

Fragility

If strong subadditivity were ever experimentally violated, the entire framework would collapse. No known experiment challenges it. The inequality has survived fifty years of testing.

Full dependency table: the companion Complete Mathematics document.

Premise Ledger

Every closed or conditionally-closed result rests on a named assumption:

Claim	Section	Premise
Lorentz IR invariance (rotational)	§5	O_h point-group protection: T-gate is A_1g singlet → $O(k^2)$ isotropy to all orders; one-loop confirmed. Boost invariance open (L1).
Dark energy w = -1	§6, §7	Volume-extensivity of bifurcation residue
Cosmological constant Λ	§6, §7	w = -1 (above)
SU(3) gauge group	§9	Automaton IR fixed point = sl3 spider category (RG premise)
$n_s \approx 0.967$	§6	$N_e \approx 60$ from standard cosmology (consistency check only)
Dilution factor η = 1-1/e	§3	α = 1 from fixed-point stability (conjectured)
Fifth force signature	§7	Rank-2 source profile + linear coupling; computed $\kappa_{eff} = g_\Omega \approx 0.2$ (dimensionless), $m_\Omega = -\ln \eta$ (contact-like); magnitude set by l_{fund}
Dark energy w(z): driven Quintom-B crossing from below (primary); w = −1 equilibrium backup	§7, §11.9	source-sink $\dot{\rho} = S - \Gamma\rho$; volume-extensivity for backup (ANEC satisfied, pointwise NEC locally violated)

Three-Layer Classification

Every claim belongs to one of three categories:

Layer 1, Imported theorems (established in the literature): | Theorem | Source | Used for | |---------|--------|----------| | Monotonicity of relative entropy / SSA | Lindblad (1975), Uhlmann (1977), Lieb-Ruskai (1973) | Metric, gravity, dark energy bound, monogamy, reflexive bound | | Entanglement first law → Einstein equations | Jacobson (1995), Faulkner et al. (2014) | Gravity as equation of state | | Markov blanket self | HJPW (2004), companion paper | Self-model boundary condition | | Monogamy of squashed entanglement | Christandl-Winter (2004) | Monogamy budget | | Lieb-Robinson bound | Lieb-Robinson (1972) | Finite propagation speed, emergent c | | Envariance → Born rule | Zurek (2003) | Measurement probabilities | | Margolus-Levitin bound | Margolus-Levitin (1998)

| Energy-information conversion | | Cartan classification | Cartan (1894) | SU(3) uniqueness proof | | Unknotting theorem | Geometric topology | 3+1 dimensions |

Layer 2, Framework's proven consequences (derived from axioms + imported theorems): | Result | Derivation | |--------|-----------| | Φ_bind ≥ 0 | Subadditivity (Klein) → closed-form proof | | S_3-equivariant update rule | Construction (update rule is S_3-symmetric by design); emergence from a non-symmetric start not yet shown | | Two-point correlator: G(x,y) = δ_{xy}·Var* | AR(1) steady-state solution | | Lorentz incoherent channel closed | Two-point correlator + coarse-graining | | d_eff convergence toward 3 | Finite-size scaling extrapolation | | ⟨ℐ⟩ ≈ 1.17 nats/edge | Automaton measurement (no tuning) | | Thread graph → metric (triangle inequality) | SSA + geodesic closure | | Lorentz rotational invariance ($O(k^2)$) | O_h point-group theorem: T-gate is A_1g singlet → all corrections ∝ δ_ij; one-loop confirmed | | Lorentz $O(k^4)$ anisotropy IR-irrelevant | Power counting: relative anisotropy ∝ k^2 → 0 |

Layer 3, Framework's conjectures (asserted, not yet proven): | Conjecture | Status | |-----------|--------| | w = −1 for bifurcation residue | Conditional on volume-extensivity (unproven) | | Dark matter from traceless E_μν | Numerical correlation r = 0.556 (preliminary L≤12) | | SU(3) from automaton IR fixed point | Cartan step proven; sl_3 IR premise unproven (RG) | | η = 1−1/e | Conjectured; measured range 0.63–0.70 | | α = 1 | Conjectured from fixed-point stability | | Fifth force functional signature | Computed: m_Ω = −ln η (contact-like), κ_eff = g_Ω ≈ 0.2; magnitude set by l_fund | | n_s ≈ 0.967 | Consistency check from 2/N_e (N_e from standard cosmology) | | Cascade termination mechanism | 11 hypotheses tested; none survive | | N ≈ 10^{122} / Λ magnitude | Holographic identity; N from observation (circular) | | Dark energy w(z) | driven dynamics → Quintom-B crossing from below (primary, §11.9); w = −1 equilibrium backup (volume-extensivity) | | Lorentz boost invariance | Open; single-cone universality across modes (L1 program) |

Appendix B: Framework Recap

The entanglement tensor $E_{\mu u} = \sum_y \mathcal{I}(x, y) \cdot n_\mu n_u$ admits three readings:

1. **Computational** (stabilizer vs magic) → rigid space + responsive gravity ← **This paper**
2. **Algebraic** (trace vs traceless) → dark energy + dark matter
3. **Partitional** (inside vs boundary vs outside) → self + world (explored in companion work)

The framework's complete axiomatic structure is presented in Halldén (2026), *The Thread Framework*.

References

[1] Lieb, E.H. & Ruskai, M.B. (1973). "Proof of the strong subadditivity of quantum-mechanical entropy." *Journal of Mathematical Physics*, 14(12),

1938-1941.

[2] Lindblad, G. (1975). "Completely positive maps and entropy inequalities." *Communications in Mathematical Physics*, 40(2), 147-151.

[3] Uhlmann, A. (1977). "Relative entropy and the Wigner-Yanase-Dyson-Lieb concavity in an interpolation theory." *Communications in Mathematical Physics*, 54(1), 21-32.

[4] Lieb, E.H. & Robinson, D.W. (1972). "The finite group velocity of quantum spin systems." *Communications in Mathematical Physics*, 28(3), 251-257.

[5] Bekenstein, J.D. (1973). "Black holes and entropy." *Physical Review D*, 7(8), 2333-2346.

[6] Hawking, S.W. (1975). "Particle creation by black holes." *Communications in Mathematical Physics*, 43(3), 199-220.

[7] Schwinger, J. (1951). "On gauge invariance and vacuum polarization." *Physical Review*, 82(5), 664-679.

[8] Wheeler, J.A. (1990). "Information, physics, quantum: The search for links." In *Complexity, Entropy, and the Physics of Information* (SFI Studies in the Sciences of Complexity, vol. VIII). Westview Press.

[9] Page, D.N. (1993). "Information in black hole radiation." *Physical Review Letters*, 71(23), 3743-3746. arXiv:hep-th/9306083.

[10] Jacobson, T. (1995). "Thermodynamics of spacetime: The Einstein equation of state." *Physical Review Letters*, 75(7), 1260-1263. arXiv:gr-qc/9504004.

[11] Margolus, N. & Levitin, L.B. (1998). "The maximum speed of dynamical evolution." *Physica D*, 120(1-2), 188-195. arXiv:quant-ph/9710043.

[12] Zurek, W.H. (2003). "Environment-assisted invariance, entanglement, and probabilities in quantum physics." *Physical Review Letters*, 90(12), 120404. arXiv:quant-ph/0211037.

[13] Christandl, M. & Winter, A. (2004). "Squashed entanglement: An additive entanglement measure." *Journal of Mathematical Physics*, 45(3), 829-840. arXiv:quant-ph/0308088.

[14] Hayden, P. & Preskill, J. (2007). "Black holes as mirrors: quantum information in random subsystems." *Journal of High Energy Physics*, 2007(09), 120. arXiv:0708.4025.

[15] Van Raamsdonk, M. (2010). "Building up spacetime with quantum entanglement." *General Relativity and Gravitation*, 42(10), 2323-2329. arXiv:1005.3035 [hep-th].

[16] Faulkner, T., Guica, M., Hartman, T., Myers, R.C. & Van Raamsdonk, M. (2014). "Gravitation from entanglement in holographic CFTs." *Journal of High Energy Physics*, 2014(3), 51. arXiv:1312.7856 [hep-th].

[17] Faulkner, T., Leigh, R.G., Parrikar, O. & Wang, H. (2016). "Modular Hamiltonians for deformed half-spaces and the averaged null energy condition." *Journal of High Energy Physics*, 2016(9), 38. arXiv:1605.08072 [hep-th].

[18] Planck Collaboration (2020). "Planck 2018 results. VI. Cosmological parameters." *Astronomy & Astrophysics*, 641, A6. arXiv:1807.06209.

[19] Cao, C., Cheng, G., Karthikeyan, K., Li, C. & Preskill, J. (2026). "State-dependent geometries from magic-enriched quantum codes." arXiv:2603.13475 [hep-th].

[20] Berera, A., Buniy, R.V., Kephart, T.W., Päs, H. & Rosa, J.G. (2017). "Knotty inflation and the dimensionality of spacetime." *European Physical Journal C*, 77, 682. arXiv:1508.01458 [hep-th].

[21] Sakharov, A.D. (1967). "Vacuum quantum fluctuations in curved space and the theory of gravitation." *Doklady Akademii Nauk SSSR*, 177(1), 70-71.

[22] Georgi, H. & Glashow, S.L. (1974). "Unity of all elementary-particle forces." *Physical Review Letters*, 32(8), 438-441.

[23] Ehrenfest, P. (1918). "In what way does it become manifest in the fundamental laws of physics that space has three dimensions?" *Proceedings of the Amsterdam Academy*, 20, 200-209.

[24] Tarski, A. (1955). "A lattice-theoretical fixpoint theorem and its applications." *Pacific Journal of Mathematics*, 5(2), 285-309.

[25] Gottesman, D. (1998). "The Heisenberg representation of quantum computers." arXiv:quant-ph/9807006.

[26] Călugăreanu, G. (1959). "L'intégrale de Gauss et l'analyse des nœuds tridimensionnels." *Revue de Mathématiques Pures et Appliquées*, 4, 5-20.

[27] Drinfeld, V.G. (1985). "Hopf algebras and the quantum Yang-Baxter equation." *Doklady Akademii Nauk SSSR*, 283(5), 1060-1064.

[28] Weinberg, S. & Witten, E. (1980). "Limits on massless particles." *Physics Letters B*, 96(1-2), 59-62.

[29] Verlinde, E. (2011). "On the origin of gravity and the laws of Newton." *Journal of High Energy Physics*, 2011(4), 29. arXiv:1001.0785 [hep-th].

[30] Bose, S., Mazumdar, A., Morley, G.W., Ulbricht, H., Toroš, M., Paternostro, M., Geraci, A.A., Barker, P.F., Kim, M.S. & Milburn, G. (2017). "Spin entanglement witness for quantum gravity." *Physical Review Letters*, 119(24), 240401. arXiv:1707.06050 [quant-ph].

[31] Marletto, C. & Vedral, V. (2017). "Gravitationally-induced entanglement between two massive particles is sufficient evidence of quantum effects in gravity." *Physical Review Letters*, 119(24), 240402. arXiv:1707.06036 [quant-ph].

[32] Cao, C., Cheng, G., Hamma, A., Leone, L., Munizzi, W. & Oliviero, S.F.E. (2024). "Gravitational back-reaction is magical." arXiv:2403.07056 [hep-th]. (PRX Quantum 6, 040375, 2025.)

[33] Takayanagi, T. (2025). "Essay: Emergent Holographic Spacetime from Quantum Information." *Physical Review Letters*, 134(24), 240001. arXiv:2506.06595 [hep-th].

[34] Vilenkin, A. (1982). "Creation of universes from nothing." *Physics Letters B*, 117(1–2), 25–28.

[35] Hartle, J.B. & Hawking, S.W. (1983). "Wave function of the Universe." *Physical Review D*, 28(12), 2960–2975.

[36] Wolfram, S. (2002). *A New Kind of Science*. Wolfram Media.

[37] Tegmark, M. (2014). *Our Mathematical Universe: My Quest for the Ultimate Nature of Reality*. Knopf.

[38] Bennett, C.H. (1988). "Logical depth and physical complexity." In *The Universal Turing Machine: A Half-Century Survey* (R. Herken, ed.), Oxford University Press, 227–257.

[39] Lawvere, F.W. (1969). "Diagonal arguments and cartesian closed categories." In *Category Theory, Homology Theory and their Applications II*, Lecture Notes in Mathematics 92, Springer, 134–145. (Historical note: an earlier version of this paper invoked Lawvere's diagonal theorem for the self-reference argument; the current formulation replaces this with the Markov blanket framework of the companion paper. The citation is retained for completeness.)

[40] Wilzewski, A., Huber, L. I., Aude Craik, D. P. L. et al. (2025). "Nonlinear calcium King plot constrains new bosons and nuclear properties." *Physical Review Letters*, 134, 233002. arXiv:2412.10277.