The View from Inside

Consciousness as Structured Perspective

Christian Halldén

Abstract

Consciousness is a structural consequence of the Thread Framework, a theory in which every interaction point is a vantage from which difference is registered. When many such points share sufficient mutual information, they bind into a unified perspective. The self is a quantum Markov blanket: a region whose boundary renders its interior and exterior conditionally independent, I(inside:outside|boundary)→0. This is the saturation of the framework's own strong subadditivity axiom, the same conditional mutual information that builds spacetime. The monogamy of quantum correlations sets a budget that bounds how strongly a region can bind internally while remaining distinct from its surroundings. A self maintains this blanket by minimizing variational free energy, an upper bound on the blanket leak; this minimization is the descent of that leak, so the Free Energy Principle is recovered from the framework rather than added to it. This account draws a clear distinction from pan-psychism: an electron has a perspective, not an experience. Experience requires structure: a self-model, binding across time, a budget of quantum correlation. The framework makes testable predictions about anesthesia (integration measure drops below threshold), meditation (Default Mode Network suppression), and split-brain phenomena (two independent Markov blankets form). From the axioms, this is what consciousness is, and where the felt inside cannot be derived further it is identified as a primitive at the base, the same base where the question of why there is anything dissolves rather than calling for an answer. The physics results in the companion paper, gravity, cosmology, and the dark sector, are independent of the consciousness claims developed here.

1. Introduction

The companion paper established that information-theoretic constraints on a network of interacting points yield emergent spacetime, gravity, and the dark sector. The same mathematical object, the entanglement tensor $E_{\mu\nu} = \Sigma\, \mathcal{I}(x,y) \cdot n_\mu\, n_\nu$, carries three readings. The algebraic decomposition (trace vs traceless) gives the dark sector. The computational decomposition (stabilizer vs magic) gives the distinction between rigid space and responsive gravity. The third, the partitional decomposition (inside vs boundary vs outside), gives the self. One tensor, three readings, cosmology, dynamics, mind.

If the observable universe is a consequence of an irreducible process, a process whose states cannot be computed from initial conditions but must be produced step by step, then being a subject inside that process is the only way it can exist.

An irreducible process has no complete description from the outside. It must be seen from within. Every interaction point is already a vantage from which difference is registered. The accumulation of such vantages into unified perspectives is what is called binding, and the felt quality of that binding is what the framework identifies as consciousness.

From the axioms of the Thread Framework, this is what consciousness is. The account makes definite predictions about anesthesia, meditation, and split-brain phenomena. It also draws a clear distinction from pan-psychism, which attributes an inner life to all matter, whereas the framework attributes only a perspective, a vantage point inherent in any interaction. Experience requires structure: a self-model, binding across time, a budget of quantum correlation.

A note on substrate independence. The Thread Framework's core premise, established in the companion paper, is that reality is rendered from within by interacting points with limited information. What a system is, physically, is its thread-graph structure and the mutual information it carries. The substrate that implements this graph, biological neurons, silicon circuits, quantum processors, or any other medium, does not determine the physics it produces; the information structure does. The same holds for consciousness: binding is a property of the information processing, not of the substrate. A sufficiently dense network of thread interactions, regardless of its physical instantiation, is assumed to produce the same binding structure and the same felt quality.

1.1 Mathematical Preliminaries

The following definitions from the companion paper are used throughout:

Definition (Thread graph). G = (V, E_id, E_I) with vertices as interaction points, each hosting a qubit. Mutual information $\mathcal{I}(i,j)$ quantifies shared distinguishability between vertices i and j.

Definition (Region). A subset X \subseteq V. The internal binding of X is the multipartite mutual information: $\mathcal{I}_{\text{bind}}(X) = \sum_{i \in X} S(\rho_i) - S(\rho_X)$. The external coupling of X is $\mathcal{I}(X : V \setminus X)$.

Definition (Self). A region X is a self when $\mathcal{I}_{\text{bind}}(X) \gg \mathcal{I}(X : V \setminus X)$. The self-model M_X is a compressed encoding of X's own state, updated via $M_X(t + dt) = M_X(t) + \beta \cdot \text{obs}_X(t)$. The boundary of X forms a quantum Markov blanket: conditional on the boundary B = ∂X, the interior A and exterior C are independent, $\mathcal{I}(A : C|B) \to 0$. This is the saturation of strong subadditivity (Hayden-Jozsa-Petz-Winter 2004), the framework's own grounding axiom. A Markov blanket gives a *quantitative* integrity condition, how close $\mathcal{I}(A : C|B)$ is to zero, bounded by the monogamy budget, rather than the merely existential guarantee of a fixed point.

Definition (Monogamy budget). Squashed entanglement $E_{\text{sq}}(\rho_{AB})$ satisfies monogamy (Christandl-Winter 2004). A system's internal binding precludes strong external binding on the same correlation budget: $\mathcal{I}(X : M_X) + \mathcal{I}(X : V \setminus X) \leq C_{\max}$.

Definition (Perspective). A perspective is the information accessible from a single interaction point: the algebra of observables generated by edges incident to that point. A perspective is not an experience; experience requires the organization of many perspectives into a self-model.

Definition (Vantage). A vantage is the minimal observation inherent in any interaction, the point from which difference is registered. Every interaction point is a vantage. An electron has a vantage; it does not have experience.

Definition (Individual). An individual is a region of the thread graph whose internal binding significantly exceeds its external coupling, forming a unified perspective with a self-model. Individuals experience. Individuals are what vantages become when bound together through sufficient mutual information.

Definition (Binding measure). For a region X, the binding measure is:

$$\Phi_{\text{bind}}(X) = \mathcal{I}_{\text{bind}}(X) - \max_{\pi} \sum_{k} \mathcal{I}_{\text{bind}}(X_k)$$

where π ranges over nontrivial partitions of X. This can be rewritten in closed form:

$$\Phi_{\text{bind}}(X) = \min_{\pi} \sum_{k} S(\rho_{X_k}) - S(\rho_X)$$

By subadditivity of von Neumann entropy (Klein's inequality), $\sum_k S(\rho_{X_k}) - S(\rho_X) \geq 0$ for every partition, therefore $\Phi_{\text{bind}}(X) \geq 0$ for all states. Equality holds if and only if the minimizing partition exactly factorizes: $\rho_X = \otimes_k \rho_{X_k}$. The binding measure is non-negative and zero precisely when the system is a product of independent parts.

Definition (Synergy). The co-information $I_3(A:B:C) = I(A:B) - I(A:B|C)$ distinguishes two modes of tripartite correlation: $I_3 > 0$ indicates redundancy (pairwise correlations dominate), while $I_3 < 0$ indicates synergy (correlations irreducible to pairwise terms). The framework predicts that conscious bound states operate in the synergistic regime ($I_3 < 0$), where the whole is not merely the sum of pairwise links. I_3 is a separate diagnostic from Φ_{bind}; Φ_{bind} measures irreducible whole-vs-parts binding, while I_3 measures synergy versus redundancy.

2. Sentience and Perception

The usual way to ask about consciousness assumes a world all outside, dead matter with no inner side, out of which an inside must somehow be conjured once things grow complicated. The framework does not start there, and that one difference moves the rest. From the first interaction there was a vantage, an inside, however thin; observation was the ground floor, not a later switch. Information always did two things between vantages: it flowed, as one bears on another, and it gathered, as many come to hold a state in common. These are not separate facts. The same recursive perturbation the companion paper calls magic, the net disturbing itself so that each interaction sets up the next, is, from

a vantage, information arriving, and a vantage registering it is observation. Recursive perturbation, information, and observation are one process named from three sides: the making seen from outside, the difference carried, and the inside that takes it in. Nothing is purely outside and then, somehow, given an inside later. The inside was there from the start. Climbing changes how much of the inside gathers in one place, how it specializes, and whether it folds back on itself, never whether there is one.

Every interaction point is a vantage. What remains is the climb from bare observation to full experience. The climb has four stages: perception, sentience, cognition, self-awareness. Perception and sentience come first; cognition and self-awareness follow.

The network, given enough internal connectivity, can observe its own state.

The subjective experience of a human, under this account, is what happens when many interaction points integrate their thread histories into a coherent cognitive state. What follows develops that account in detail, drawing where relevant on established neuroscience, not as the foundation of the argument, but as its illustration.

Sentience is the second stage: the capacity to have a perspective with felt texture. When neural networks form and memory structures develop, these perspectives build models of their own observations. An animal with a nervous system models its environment and its body. Warmth, hunger, fear. The felt texture of thread interactions, integrated across an internal network.

The Zeroth Sense

If observation is fundamental, then observation itself is prior to any particular sense. Observation in itself is bare thread-distinguishability: the capacity of an interaction point to tell threads apart. A *sense* is not a separate faculty but a projection of observation onto a particular stream, a modality tuned by evolution to extract survival-relevant structure from one slice of the network. The five senses are outward-pointing projectors; introspection is the same operation turned inward, its projector aimed at the self-subnetwork.

The inequality is the whole point. The specialized senses never recover the full operator, and what they leave out, the remainder, is the subtle channel. Seeing, feeling, and introspecting are therefore one operation aimed at different subspaces, not three kinds of thing. The zeroth sense is before any such partition: observation grounding the senses rather than standing beside them as a separate, added sense. At the limit, to be a node in the network and to sense it coincide. This is the dual-aspect reading of the framework, where the first-person side of distinguishability and the third-person Thread Network are two views of one structure.

The senses do not create observation. Observation is already present: the zeroth sense, the ground. Every thread interaction registers difference. The organism already has access to all the information its own interactions carry. What the senses provide is organization. They narrow that information into channels the organism can act on. Vision keeps a thin band of light and drops the rest,

amplifies edges, suppresses static. Hearing extracts patterns from pressure waves, discarding the rest. Every sense admits a narrow stream and blocks the rest. This dissolves the puzzle of why organic life acquired senses at all. Evolution did not switch on sight in something blind. It specialized an organ to handle a distinct stream the network was already registering. A sense is a refinement of a capacity present from the first interaction: specialization, not creation.

This inverts the usual story. The senses do not open onto an outside the vantage lacks; they narrow the information the net already carries, admitting only what is relevant to survival. Perception provides the content of experience. Sentience provides its texture.

The Noise Floor and the Subtle Channel

The self-model does more than represent; it **narrows.** Holding most of the budget for itself, it leaves a tractable, action-ready remainder, and its share sets how much reaches awareness. Distinguish two quantities: *actual* binding, the entanglement an individual shares, which is always large, and *accessible* binding, the fraction that crosses the perception threshold. What the floor governs is the second, not the first.

Masking, not destruction

Let the binding to S be a budget split between the self-model, the sensory streams, and the faint ambient threads. Define the noise floor non-circularly as the share of that budget the self-model occupies, $\mathcal{I}(M : S)$. The subtle channel is perceived when its signal clears the floor:

A loud self-narrative does not destroy the subtle channel, it masks it, by crowding with self-content. When the floor falls, accessible binding approaches actual binding, and the same ever-present threads rise over threshold. The signal did not appear; the mask lifted.

This maps onto precision-weighting in predictive processing: a high floor is high precision on high-level priors suppressing bottom-up signal, the REBUS account (Carhart-Harris and Friston), with the Default Mode Network as substrate. The seam splits in two.

Terminal lucidity mechanism. Loosened binding raises the signal-to-noise ratio on sub-threshold physical channels that already exist, normally masked by the internal signal. Terminal lucidity is internal: collapsing binding unmasks the system's own stored content. Fresh, causally shielded randomness yields a strict null: $\mathcal{I}(\text{vantage} : E_{\text{fresh,shielded}}) = 0$. The null is the prediction.

(A stronger reading, distinguishable information about fresh shielded variables, would contradict the no-communication theorem; the framework does not make it. See Appendix.)

Full Derivation: Perception

The zeroth sense: observation itself, prior to any modality (A3 instantiated in organisms). Each evolved sense is a projection of the access algebra onto a

survival-relevant stream: specialization of a capacity present from the first interaction, never an activation from blindness.

Derivation chain

Step 1. Every sense is a processing chain. The data-processing inequality bounds it: for X → Y → Z,

$$\mathcal{I}(X:Z) \leq \mathcal{I}(X:Y)$$

No step of processing can add information about the source; a sense can only narrow and lose. That is the data-processing inequality in words.

Step 2. Finite channel capacity bounds the admitted stream; everything else is excluded by the bound, not by absence. Masking, not destruction, at the very first step of experience.

Full Derivation: Binding

Theorem 17.1 (The surplus). B(S) = Σ_{i∈S} S(ρ_i) − S(ρ_S) ≥ 0, and for strongly entangled S the parts are individually mixed while the whole is definite: B counts information present in the joint state and absent from every part. No part of a bound region knows the region's state; the knowing exists only at the level where part-descriptions end. Section 4 refines this bare surplus into the cut-invariant Φ_{bind}, the quantity that gates experience.

Live, not archival. Binding must be regenerated: by Lemma 4.2, correlation decays under local noise unless re-created by ongoing interaction. A self is a process maintaining its own $\mathcal{J}$ against dilution; a record shares structure once.

Dilution of the distant. Pairwise $\mathcal{J}$ between specific distant subsystems decreases under local decoherence (Lemma 4.2) and after cosmological ages is driven toward zero, surviving only as inaccessible multipartite correlation. Common ancestry: yes. Usable pairwise channels: no.

The leak L. The conversion of pairwise mutual information into inaccessible multipartite form, the process Lemma 4.2 describes, is itself a named quantity. Define the total leak L as the sum, over all tripartitions (A,B,C) where B fails to separate A from C, of the conditional mutual information:

$$L(t) = \sum_{(A,B,C):B \text{ nonseparating}} \mathcal{I}(A:C|B)$$

L is a *steady-state diagnostic* for blanket integrity. Under the full driven dynamics (pairwise interaction + genesis + dilution), L relaxes from an initial value to a steady-state level set by the balance of source and dilution, the dynamics organize the system, reducing non-Markovicity. Once in steady state, L is maintained at a constant level; it is not a global monotone. In the restricted regime of a single nonlocal perturbation on an otherwise perfect Markov blanket, L does grow monotonically (simulation on a 6-qubit cluster chain: I(A:C|B) rises from 0 to 2 ln 2 as a nonlocal edge is introduced, g = 0 → 0.8; but

at g > 0.8 the blanket structure dissolves, and in the full driven n=4 FaithfulFlow, L is stable, not growing). The four phenomena L unifies, decoherence, time's arrow, dark energy, and blanket leakage, are read from L's steady-state level, its relaxation rate, and its spatial distribution, not from a claim of perpetual growth. The leak L unifies four phenomena as one:

1. **Decoherence:** the steady-state level of L quantifies the unavoidable decoherence, the irreversible leakage of coherence from pairwise into inaccessible multipartite degrees of freedom.
2. **Time's arrow:** L's relaxation toward steady state defines a local arrow of time. Every interaction either preserves L (reversible) or increases it (irreversible). The direction of increasing L *is* the direction of time.
3. **Dark energy:** in the companion paper, the dilution rate sets the cosmic expansion rate. The throughput of L, the rate at which pairwise correlation is converted to inaccessible multipartite, is the energy density driving accelerated expansion. At equilibrium this fixes w = −1, a cosmological constant; off equilibrium the driven throughput produces the single phantom crossing of the companion paper (w < −1 in the past, w > −1 today), with L's positivity underwriting the averaged statement rather than a pointwise floor.
4. **Blanket leakage:** a Markov blanket self maintains $\mathcal{I}(A:C|B) \to 0$ locally. Every nonlocal edge that bypasses the blanket boundary increases L, this is blanket leakage, the thermodynamic cost of maintaining a self against the universal arrow.

The unification is compact: **one diagnostic, four faces.**

Derivation chain

Step 1. The surplus: multipartite mutual information

$$I(S) = \Sigma_{p}artsS(part) - S(whole)$$

measures information present in the whole and absent from every part.

Step 2. For an entangled region the whole can be definite while every part is mixed: the definiteness lives in the relations. This is the exact answer to the combination objection: binding is not parts summing to one; it is surplus carried jointly, and that surplus is the correlation budget spent inside, which is the same budget removed from accessible inside distinction, so no inside vantage reads a cut and the region is one subject (§4.1).

Step 3. Graded unity: the region functions as one individual in proportion to internal coupling against boundary coupling; no threshold event, no merger.

---

3. Organic Life

Before there is biology there is a process that organizes matter with no hand to guide it, and it is the same process that will later be called life. It is already running in the physics. From the first interaction a vantage registers a

difference and answers it, and that registering-and-answering is the seed of inference. A region that persists must hold its boundary in a narrow set of states against universal decay, which is to say its dynamics keep the blanket leak low. No region has access to that leak directly, since that would mean summing over every outside cause; what is available, and what anything that lasts runs on, is the descent of a tractable quantity that bounds it from above. That quantity is the variational free energy of the Free Energy Principle, and minimizing it is how a region organizes itself, comes to predict its world, and spends as little as it can to stay distinct. Nothing supervises this. The only alternative to closing the gap is to dissolve, so the process runs wherever a boundary must be held, long before any cell. It is the driver. Organic chemistry, when it arrives, does not invent it; it gives it a stable place to run recursively and compound.

3.1 Free Energy: Self-Organization Without Supervision

The self-model update of §1.1, $M_X(t+dt) = M_X(t) + \beta\,\mathrm{obs}_X(t)$, was stated as a rule, but is one gradient step of a single quantity, and that quantity explains how a self maintains its blanket with no supervisor. The quantity is the variational free energy of the Free Energy Principle (Friston 2010), recovered here from the framework's own Markov blanket rather than imported.

Setup. Partition the graph around a self X into internal states μ (the self-model degrees of freedom M_X), blanket states $b = \partial X$ split into sensory s and active a, and external states $\eta = V \setminus X$. The blanket condition of §1.1 is $\mathcal{I}(\mu : \eta \mid b) = 0$: internal and external states meet only through the boundary. The leak is the residual when the seal is imperfect, $L = \mathcal{I}(A : C \mid B) \geq 0$.

Surprisal and persistence. A self that persists must keep its boundary states inside a narrow characteristic set, so on average it occupies states of low surprisal, where the surprisal of a sensory state is $-\ln p(s)$. A self cannot evaluate $-\ln p(s)$ directly, since that would require integrating over every external cause η. It minimizes instead a tractable upper bound.

The bound. Let the internal states carry a recognition density $q(\eta;\mu)$, the self-model's estimate of the external causes. The variational free energy is

$$F(s,\mu) = \mathbb{E}_q[\ln q(\eta;\mu) - \ln p(\eta, s)] = \underbrace{D_{\mathrm{KL}}[q(\eta;\mu) \,\|\, p(\eta \mid s)]}_{\geq 0} - \ln p(s) \;\geq\; -\ln p(s).$$

Minimizing F over μ drives the KL term to zero, so $q(\eta;\mu) \to p(\eta \mid s)$: the self-model comes to encode the posterior over external causes (perception). Minimizing F over a changes which s occur, lowering the surprisal by making sensation match the model (action). Together this is active inference: a self whose dynamics minimize F behaves as if it tracked the world and held a match to it. The heuristic update $M_X \mathrel{+}= \beta\,\mathrm{obs}_X$ is one step of this descent.

The crux: free energy is the leak. Perception sets $q(\eta;\mu) = p(\eta \mid s)$, the reconstruction of the exterior from the boundary, which is the Petz recovery map (Hayden-Jozsa-Petz-Winter 2004) already central to the self's self-reconstruction.

The gap it closes is exactly the conditional dependence the blanket failed to screen:

$$F - \big(-\ln p(s)\big) = D_{\mathrm{KL}}[\,q \,\|\, p(\eta \mid s)\,] \xrightarrow{\text{at the blanket}} L = \mathcal{I}(\mu : \eta \mid b).$$

Minimizing free energy and minimizing the leak L are one imperative seen twice. The leak was introduced as the steady-state cost a self pays to stay distinct against the universal arrow; the Free Energy Principle names the act by which that cost is paid, closing the prediction gap. The driven dynamics (pairwise interaction, genesis, dilution) supply the throughput the self spends to hold L, and so F, low.

Gradient (predictive coding). For Gaussian densities, up to a constant,

$$F \approx \tfrac{1}{2}\frac{(s-\mu)^2}{\sigma_s^2} + \tfrac{1}{2}\frac{\mu^2}{\sigma_p^2}, \qquad \dot{\mu} \propto -\frac{\partial F}{\partial \mu} = \frac{s-\mu}{\sigma_s^2} - \frac{\mu}{\sigma_p^2},$$

precision-weighted prediction-error correction: the principled form of $M_X \mathrel{+}= \beta\,\mathrm{obs}_X$, and the same precision-weighting invoked in the REBUS account of altered states. This is continuous with the physics. Observe, respond, adapt is present from the first interaction (the zeroth sense and the opening of §3), the same registering-and-answering that bends geometry under matter in the companion paper. Organic chemistry did not invent it; it found, on a stable substrate, a way to run it recursively.

Numerical check. A minimal model, external Ornstein-Uhlenbeck η, noisy blanket $s = \eta + \mathrm{noise}$, predictive-coding update on μ, against a non-inferring control with μ frozen: the inferring self holds the residual leak proxy $\mathrm{Var}(\eta - \mu) \approx 0.030$ versus ≈ 0.085 for the control, about $2.8\times$ lower (seed fixed; script in `aux/fep_freeenergy_sim_2026-06-20.py`). Closing the gap keeps the blanket informed about the exterior it screens.

The substrate it found. A process this demanding needs a stable substrate, and carbon supplies it. Carbon's four bonds let atoms chain, branch, and ring into a combinatorial space of stable structures far larger than any other element opens, the plateau on which the gap-closing process could run without shaking apart. On that plateau the arrangements that close their prediction gap better are the ones that persist, so the substrate fills with the functions of staying: metabolism, the throughput a self spends to hold its leak low; memory, the model carried across time; reproduction, a gap-closer copying itself; and senses, channels narrowed to what can be acted on. These are the stable functional combinations. None is new in kind. Each is the underlying process finding a steadier way to do what it already did from the first interaction.

Life, located. Life is then not organic life. Organic life is one stable manifestation, the local chapter life here is made of. Life is the irreducible, recursive, gap-closing process itself, the same process that turned nothing into something and something into structure, now caught and held on a chemical substrate. It was running long before the first cell, and the framework's own recursion reaches further: a sealed interior carries the process on (companion

paper §6.2), so the run our universe belongs to may have been turning before this universe began. Why it runs at all traces to the start. Nothing is the one configuration that cannot persist, a single way to be nothing against the countless ways to be something, so the chance of staying nothing falls to zero and something always wins; and once there is something, this gap-closing process is what holds it together. Organic life is that process settling, for a while, into a form that can reproduce, remember, and feel.

---

4. Cognition and Consciousness

Binding

Consider two points in the network that have interacted. They share information, and to that extent they already share a perspective. The strongest objection was put by William James: a hundred feelings side by side are still a hundred feelings, not one feeling that holds them, and stacking parts does not by itself make a whole that feels. He was right; binding is not many things summed. What entanglement adds is this: when the shared information inside a region runs high, no part holds the whole region's state on its own. The definiteness lives in the relations between the parts, and that surplus, the information present in the whole and absent from every part, is what a self is made of. Identity does not glue; it migrates upward into the joint state. A self is therefore graded, as strong as its internal weave is against its coupling to the world outside. Binding strength is read off the same quantity that fixes thread distance, so the mathematics that builds space builds the self. Binding strength is the multipartite mutual information:

$$\Phi_{\text{bind}}(X) = \mathcal{I}_{\text{bind}}(X) - \max_{\text{cuts}} \sum_k \mathcal{I}_{\text{bind}}(X_k)$$

This quantity measures the information present in the whole region X that is absent from any partition of X into subregions. It refines the bare surplus B(S) of Theorem 17.1: where B(S) counts what the whole holds beyond its individual parts, Φ_{bind} subtracts the strongest partition and keeps only what survives every cut, the irreducible binding that gates experience. When Phi_bind exceeds the external coupling $\mathcal{I}(X : V \setminus X)$, the region reads as a unified self. When it falls below, separate perspectives emerge. The definition converges conceptually with Integrated Information Theory (Tononi 2004) but differs in being derived from quantum information constraints rather than phenomenological axioms.

The measure says how much a region binds. The mechanism says why a bound region is one subject and not many, and it follows from what a vantage can access (§3). Monogamy gives each part a finite correlation budget, and binding spends that budget on the joint state, so the same rise in Φ_{bind} that marks a strong whole lowers what any inside part can still resolve about the others: a part-to-part distinction needs budget that binding has already spent. When the inside distinctions that would mark a boundary are no longer accessible, no inside vantage reads a cut, and the region is one subject by construction, not by stipulation. James is answered a second time. Not by parts summing, which they

do not, but because the binding that makes a whole is the same act that removes the inside view of its parts. The surplus the measure reports is the budget taken out of accessible inside distinction. At the extreme, where a region spends its whole budget inside, the inside goes blank and no vantage resolves anything, which is the saturation the companion paper reads as a genesis nothing. A self lives below that limit: bound enough to be one, open enough to still have something to tell apart.

Worked example: three-node network. Consider nodes A, B, C with von Neumann entropies S_A = S_B = S_C = 1.0, S_AB = 1.5, S_BC = 1.6, S_AC = 1.9, S_ABC = 2.0. The internal binding is $\mathcal{I}_{\text{bind}}(ABC) = \sum_i S_i - S_{ABC} = 3.0 - 2.0 = 1.0$. All nontrivial partitions:

Cut	$\Sigma\mathcal{I}_{\text{bind}}$
AB|C	(1.0+1.0−1.5) + 0 = 0.5
AC|B	(1.0+1.0−1.9) + 0 = 0.1
A|BC	0 + (1.0+1.0−1.6) = 0.4

The most informative partition is AB|C (max Σ = 0.5). The binding measure is $\Phi(ABC) = 1.0 - 0.5 = 0.5$: the triple shows significant binding, with half the internal information unique to the whole. If interactions weaken, Φ drops and the triple dissolves into separate perspectives. This illustrates the continuous gradation of binding strength the framework predicts.

This is why a strong mind reads as a single identity. The brain couples to itself more tightly than it couples to anything outside, by many orders of magnitude, so the internal signal dominates and the region screens itself off as a self. The boundary is a signal-to-noise ratio, not a wall. Threads that are woven together still share perspective in proportion to what they share, which is why two people in deep conversation overlap a little, and why severing the bridge between the brain's hemispheres, which cuts the internal coupling along one line, leaves two partial selves where there was one. The self is the region where the cuts are expensive inside and cheap around.

When the internal coupling far outweighs the coupling across the boundary, the region reads as one individual; as that contrast falls, the shared perspective thins back toward many. Unity is graded, never a sudden merger. Wiest (2025) reports preliminary evidence of quantum effects in neural microtubules correlated with conscious state. The framework accepts the phenomenology while dropping the Orch-OR objective-reduction mechanism (see companion paper §10.6).

With perception channeling the information and sentience giving it texture, a third layer builds on top: cognition, pattern recognition, memory, decision-making. These are processes running on the neural substrate: algorithms tuned by evolution. A dog recognizes its owner. A crow solves a puzzle. Cognition processes the information that sentience and perception provide.

Self-awareness

Perception selects what reaches a vantage. Sentience is the organized processing of such observation. Cognition binds many vantages into a single individual that models its world. Self-awareness is the fourth stage: cognition turned on its own activity.

A cognitive network models its surroundings. The network belongs to its own surroundings, so a sufficiently capable model eventually includes the modeler. When that inner model begins to steer the network's activity, the loop closes. The system registers its own registering. It knows that it knows.

The effect is graded, as everything in this framework is graded. Self-awareness is as strong as the loop is tight: how faithfully the inner model tracks the region it models, and how much of the region's next state it steers. A model that tracks without steering only records. Steering closes the loop, and the closed loop is the self-awareness. The loop spends the network's magic to render its novelty; it cannot create that resource by recursing. Self-observation is the consumer at this boundary, and what it consumes was supplied long before.

The self-observing loop

Let S be a bound subnetwork and $M \subset S$ its self-model: a compressed encoding of the state of S, with $\mathcal{I}(M:S)$ near its maximum for the size of M. The loop condition is feedback: the network's next state depends on its own model,

$$S_{t+1} = F(S_t, M_t), \qquad M_t = \mathrm{enc}(S_t)$$

Self-awareness strength is the product of binding and the model's causal influence:

$$A(S) = B(S) \cdot C(M \to S)$$

where $B(S)$ is the binding surplus and $C(M \to S)$ measures how much of the network's update is steered by M. The self-observing loop is what Hofstadter named a strange loop (1979), realized here as a feedback condition on a bound region.

The boundary just described has a deeper origin than attention alone, fixed by the budget of quantum correlation a region can hold.

The monogamy budget

Mutual information splits into quantum and classical components: $\mathcal{I} = Q + C$, where $Q = 2E_{\mathrm{sq}}$ is twice the squashed entanglement and $C \geq 0$ is classical correlation. **Monogamy of the quantum part:** $Q(A:B_1) + \cdots + Q(A:B_n) \leq 2S(\rho_A)$. A region holds a finite quantum-correlation budget. Classical correlation is free to spread.

A self is a region of high internal entanglement. Monogamy forces: $Q_{\mathrm{internal}}(S)$ high implies $Q(S:\mathrm{ambient})$ low. Dense internal binding mathematically excludes dense external quantum correlation. The boundary of the self is an entanglement-budget theorem. Releasing internal binding frees budget that can migrate outward.

Monogamy of squashed entanglement follows from strong subadditivity: the master inequality, evaluated over a recursively self-entangled subgraph.

The Loop of Consciousness

The self-model does not sit still. It runs. Every moment, the network compares its current state to the compressed model of what it was a moment ago, registers the difference, and updates. The updated model becomes the new baseline against which the next difference is registered.

This is the loop of consciousness, a continuous feedback cycle: observe the difference, update the self, observe again. William James called it the stream of consciousness; the framework gives the recursion a mechanism.

Why does one wake up as the same person? Because the loop converges. Start from the same initial conditions, the same neural architecture, the same memories, the same thread-weighting, and the feedback cycle finds the same attractor. The self-model is a stable attractor of a recursive process. Each time it restarts, the loop runs forward and settles where it did before.

The stability is not guaranteed. Shift the architecture enough, through trauma, drugs, disease, and the fixed point moves, or splits, or vanishes. Identity is continuous because the loop is continuous. Interrupt the loop, deep anesthesia, dreamless sleep, and it halts. Resume the loop, and it converges again from the last stable configuration.

The Markov Blanket Self

The self-model $M_X(t)$ is a compressed encoding of subnetwork X's state. The update rule:

$$M_X(t + dt) = M_X(t) + \beta \cdot \text{obs}_X(t) \tag{6}$$

where obs_X is the observation of X's own state, and β is the update coupling.

The self is not merely a fixed point (a trivial existence guarantee for any monotone map) but a **quantum Markov blanket**: a tripartition (A, B, C) where the boundary B separates interior A from exterior C, and the conditional mutual information vanishes:

$$\mathcal{I}(A : C|B) \to 0$$

This is exactly the saturation of strong subadditivity, the framework's own grounding axiom. Hayden, Jozsa, Petz, and Winter (2004) showed that such states form quantum Markov chains, for which a Petz recovery map $R_{B \to BC}$ reconstructs the bulk from the boundary. The self's boundary performs exactly this recovery: it predicts the world outside (active inference), and it reconstructs the interior (self-model).

Why the blanket replaces Knaster-Tarski as the definition. Knaster-Tarski guarantees that any monotone map on a complete lattice has a fixed point. This is true for *every* such map, the existence of a self-fixed-point is not a result about consciousness; it is a property of monotone functions. The Markov blanket condition, by contrast, is *quantitative*: it measures how close $\mathcal{I}(A : C|B)$ is to

zero, bounded by the monogamy budget. The condition is falsifiable: a nonlocal edge bypassing the boundary leaks the blanket monotonically. The condition is grounded in the framework's own axiom (SSA). And it unifies three fields, the Markov blanket of theoretical neuroscience (Friston), the quantum Markov chain of quantum information (HJPW), and the SSA-saturation of the Thread Framework, as the *same mathematical object*: conditional mutual information.

Simulation confirmation. Simulating a structured chain state (cluster state) with a separating boundary gives $\mathcal{I}(A:C|B) = 0$ exactly, a perfect Markov blanket. Geometry decides it: a ring (where no boundary separates) leaks maximally. A nonlocal edge bypassing the boundary increases $\mathcal{I}(A:C|B)$ monotonically from zero to maximum, confirming the prediction that blanket integrity requires suppression of nonlocal shortcuts, precisely the monogamy budget the framework already invokes.

Theorem (Markov blanket self). Let $X \subseteq V$ be a region with boundary $B = \{i \in X : \exists j \notin X, \mathcal{I}(i,j) > 0\}$. Let $A = X \setminus B$ be the interior and $C = V \setminus X$ the exterior. If $\mathcal{I}(A:C|B) < \epsilon$ for a small threshold ϵ, then X is a Markov blanket self. The Petz recovery map $R_{B\to BC}$ reconstructs the exterior from the boundary with fidelity $F \geq 1 - \sqrt{\epsilon/2}$. By the monogamy bound (Christandl-Winter), the internal binding $\mathcal{I}(A:B)$ and the blanket integrity ϵ trade off: strengthening the blanket (reducing ϵ) consumes monogamy budget that could otherwise go to internal binding.

Selection principle. Among all regions satisfying the blanket condition, the self is the one that maximizes internal binding subject to the monogamy budget: $X^* = \arg\max_X \mathcal{I}_{\text{bind}}(X)$ subject to $\mathcal{I}(A_X : C_X | B_X) < \epsilon$. This is a quantitative optimization over a well-defined information-theoretic objective.

Re-identification (continuity of the self). The blanket fixes what a self is at an instant; a separate fact fixes why the same self returns after the loop is interrupted. Take the recruitment map on candidate regions,

$$F(G) = \{\, i \in V : \mathcal{I}(i:G) > \theta \,\}.$$

When the region G grows, $\mathcal{I}(i:G)$ can only grow (strong subadditivity, the keystone), so $G_1 \subseteq G_2$ gives $F(G_1) \subseteq F(G_2)$: the map is monotone on the lattice of regions, and Knaster-Tarski (1955) makes its fixed points a complete lattice. Existence is trivial and carries no claim about consciousness, since every monotone map has a fixed point; the content is the lattice. Fix an autobiographical seed G^0, the stored trace a waking system restarts from. If $\theta < \min_{i \in G^0} S(\rho_i)$ then $G^0 \subseteq F(G^0)$, and iterating F yields an increasing chain that reaches the least fixed point G^* containing G^0 in at most $|V|$ steps. The least fixed point above a given seed is unique, the meet of all fixed points that contain it. A system that halts the loop in dreamless sleep or deep anesthesia and resumes from the same seed reconstructs the same G^*, not a neighbor, and the Markov blanket selected on G^* is the same blanket. This is why one wakes as the same person. The blanket gives the self its definition and integrity; the unique least fixed point above the seed gives its continuity across the gaps when the loop stops.

Split-brain analysis. When the corpus callosum is severed, the boundary B partitions into B_L and B_R with $\mathcal{I}(B_L : B_R)$ dropping below threshold. Each hemisphere independently satisfies the Markov blanket condition with its own boundary. Two selves form, each with its own $\mathcal{I}(A : C|B) \to 0$, consistent with clinical observation.

The loop. The continuous update $M_X(t)$ (Eq. 6) is the loop. The blanket condition $\mathcal{I}(A : C|B) \to 0$ is approached but never perfectly occupied because observation is always perturbative, every interaction with the exterior creates a slight positive $\mathcal{I}(A : C|B)$. The perpetual maintenance of the blanket against this leakage is the felt passage of inner time.

Full Derivation: The Self

The Markov blanket self (Theorem above) requires $\mathcal{I}(A : C|B) < \epsilon$ for a separating boundary B. The gating criterion $D(S) = \min \text{cut}_{\mathcal{I}}(S)/\mathcal{I}(S : \bar{S})$ distinguishes a self from its surround: no cheap cuts inside, one cheap cut around, and the self is the region that maximizes the ratio. The blanket condition, the Petz recovery, the split-brain confirmation, and re-identification across the gaps are established above.

Full Derivation: The Budget as a Theorem

Split correlation into classical and quantum parts:

$$\mathcal{I}(A : B) = C(A : B) + Q(A : B), \qquad Q := 2\,E_{\text{sq}}\ (\text{squashed entanglement})$$

Theorem 19.1 (Monogamy budget).

$$\sum_k Q(A : B_k) \le Q(A : B_1 \ldots B_n) \le 2\,S(\rho_A)$$

A region's quantum correlation is a finite budget. Dense internal binding spends it:

$$Q_{\text{internal}}(S)\ \text{high} \implies Q(S : \text{ambient})\ \text{low}$$

The boundary of the self, for the quantum component, is not maintained by attention alone; it is an entanglement-budget theorem. Releasing internal binding frees budget that can migrate outward. The classical component C is polygamous and free to spread, but carries only common-past correlation (§17, §22). The variation across the states-of-consciousness sections is a budget allocation, and the same inequality that secures quantum cryptography is the reason a tightly bound self is opaque to the ambient.

Full Derivation: The Reflexive Bound

The recursive tower G → M₁ → M₂ → … is a chain of CPTP maps. Lemma 4.2 gives per-level attrition $\mathcal{I}(G : M_{k+1}) \le \eta \cdot \mathcal{I}(G : M_k)$ with $\eta \le 1$, hence

--

$$L(N) = \sum_{k \le N} \mathcal{I}(G : M_k) \le \mathcal{I}_1 \frac{1 - \eta^N}{1 - \eta}$$

Geometric saturation, with η = 1/e as one case. Grasping the state is the operation that erodes it, a corollary of data processing: each act of self-monitoring is a processing step and cannot create information about its object.

---

5. States of Consciousness

A few features of inner life fall straight out of the one rule: looking inward is itself a process the rule governs, and the rule caps how much any process recovers, so attention turned on the self loses a little of what it reaches for.

The Reflexive Bound

The reflexive bound derived in §4 shows directly in inner life. There is a price on looking. To attend to something is to raise the mutual information it shares with the rest of the network. But noticing "I am perceiving the subtle" is itself an act of self-modeling: it routes the perception through the self-model and so raises the very quantity that masks the channel. Attention to the subtle partly erases the subtle.

$$\mathcal{I}(G : M_{k+1}) \le \eta \cdot \mathcal{I}(G : M_k), \quad \eta \le 1$$

Three quantities govern the dynamic: the ongoing subtle perception, the quieting rate of practice, and the coupling that feeds perception back into self-modeling. Because the coupling switches on exactly when the signal is strong, noticing raises the floor and shuts the channel: salient subtle perception under active noticing is an unstable fixed point. Having the state and grasping the state are the same operation switching on and off. There are two ways to open the channel, then: lower the baseline, or lower the coupling itself, decoupling noticing from self-modeling. The first is available to anyone who quiets down; the second is a trained skill, awareness that does not automatically generate an "I who is aware." It is witnessing without a witness.

This reconstructs a cluster of otherwise paradoxical contemplative teachings from a single dynamic. Seeking is the obstacle, because striving for the state is generating a goal and raising the floor. Just sitting and effortless action are the art of holding the signal low without letting noticing raise it, which is why the moment one thinks "I have it" one has lost it. And the insistence that the teaching can only point toward the state and never hand it over follows directly, since to name is to model, and to model is to raise the internal signal, so words cannot carry a state whose condition is that the signal stays quiet. The framework thereby predicts its own partial ineffability: the object at its base resists its own tools, because the tools themselves raise the signal.

Meditation

Meditation quiets the self-model. The Default Mode Network, which constructs the autobiographical self, suppresses its activity. Binding persists, but the self dissolves. Awareness continues without an "I" who is aware. The freed bandwidth spreads. Ambient threads, normally held below the floor, become distinguishable. This is the subtle channel: not a new signal, but an old one, unmasked.

The reflexive bound applies here with full force. Noticing the subtle channel raises the self-model activity that masks it. The moment one thinks "I am perceiving something subtle," the perception weakens. The Zen koan works by jamming the discursive loop: the self-model spins, exhausts itself, and releases, at which point the signal falls and insight arrives suddenly, as a threshold crossing rather than a gradual climb.

Flow: The Concentrated Quiet

Flow belongs to the same family as meditation, a state where the self-model quiets and its activity falls toward zero through full absorption. What sets it apart, and earns it its own treatment, is where the freed bandwidth goes. When the self-model falls silent, its share of the network's bandwidth is released. In meditation that freedom spreads, binding becomes diffuse, and the ambient threads rise: the subtle channel. In flow the release runs the other way, concentrated almost entirely onto a single external stream, the task. Flow is a quiet-self state with maximally focused binding, where meditation is a quiet-self state with diffuse binding. The same ego-quiet, the opposite allocation.

This explains the challenge-skill balance directly. Flow appears precisely when the task's information demand matches the bandwidth freed by quieting. Too easy, and the surplus leaks back, reviving mind-wandering and self-talk; too hard, and prediction error floods, recalling the self-model to high load. Flow is the narrow regime where the task exactly saturates and leaves nothing over for the self-model. Loss of self-consciousness is the quiet self-model itself; the merging of action and awareness is the self-model no longer sitting between perceiving and acting; the distorted sense of time follows because the self-model is a chief generator of felt time, which goes quiet along with it.

Flow meets the reflexive bound as well. To notice "I am in flow" is a self-modeling intrusion that raises the internal signal and breaks the state, the same dynamic, and choking under pressure is its special case: high stakes force the self-model back to high load and the direct task-loop dissolves. A definite differential prediction follows. Although flow and meditation share the quiet self-model, flow does not yield the subtle channel's connectedness; its practitioners report absorption in the task, not contact with the wider network, because the bandwidth went to the task. Ego-dissolution comes in at least two kinds, told apart not by how quiet the self is but by where the freed bandwidth is aimed.

A Map of Self-Model States

Anesthesia & deep sleep. B drops below the integration threshold and the self dissolves. Anesthesia lowers the binding measure; the loss of consciousness correlates with Φ_{bind} falling below threshold, consistent with clinical observation. **Split-brain.** When inter-hemisphere $\mathcal{I}$ collapses, two partial selves emerge. Clinically confirmed. **Meditation and flow.** Both quiet M ($\partial S/\partial M \to 0$) with B intact, so the self thins while awareness continues. These states make testable predictions: during meditation, the Default Mode Network (the neural correlate of the self-model) shows reduced activity, while sensory processing regions remain active, consistent with fMRI findings. Meditation spreads the freed binding to the ambient threads; flow concentrates it on a single task stream. The Default Mode Network, the neural correlate of the self-model, is suppressed in both, exactly as predicted. **Terminal lucidity.** A final maximization during the dissolution process before death, as residual entanglement re-integrates the degraded self-model.

Altered States

Psychedelics reduce Default Mode Network activity while increasing global brain connectivity. Binding expands beyond normal boundaries. The self-model dissolves, and coherent thread structures normally held below the floor become accessible. Encounters with seemingly autonomous entities may represent the self-network encountering other coherent entanglement structures within the expanded binding radius.

The characteristic quality of feeling more real than real reflects access to information normally held below the floor. The same mechanism, collapsing binding unmasking internally stored content, underlies near-death experiences and terminal lucidity: as the self-network reorganizes during dying, residual entanglement briefly re-integrates the degraded self-model, producing a final moment of clarity.

Magic-consciousness unification

The entanglement tensor decomposes by computational class: $E_{\mu\nu} = E_{\mu\nu}^{(\text{stab})} + E_{\mu\nu}^{(\text{magic})}$. The stabilizer component produces space structure, rigid geometry with no dynamics. The magic component enables matter-geometry coupling and distinguishability. Without magic, the code is a perfect, inert stabilizer code: no distinguishability, no observation, no self.

The self requires imperfection. Consciousness is what physics looks like from inside an imperfect Thread Network with sufficient recursive complexity.

State dynamics. The self-model's internal entanglement density and gating form a dynamical system with homeostasis, erosion, and relaxation. The exact bistability condition is $\sup f' > 1$. A linear response cannot produce hysteresis; a sigmoidal one can. Terminal lucidity follows as a saddle-node crisis at homeostasis failure. One mechanism, settings of one ratio.

Full Derivation: State Dynamics

State variables: internal density ρ, gating w, with homeostasis α, erosion β, relaxation κ:

$$\frac{d\rho}{dt} = \alpha(\rho_0 - \rho) - \beta w \rho, \qquad \frac{dw}{dt} = -\kappa\,(w - w_{\text{eq}}(\rho)), \quad w_{\text{eq}} \text{ decreasing}$$

Fixed points solve w = *f*(*w*) with f = w_eq ∘ ρ, *where* ρ(w) = ρ₀/(1 + βw/α). Since both maps are decreasing, f is increasing, and the precise bistability condition is:

$$\text{Bistable} \iff f \text{ crosses the identity three times} \impliedby \sup f' > 1$$

A linear w_eq cannot satisfy it (unique equilibrium, no hysteresis); a sigmoidal w_eq can. Hysteresis, breakthrough, and the saddle-node crisis at α → 0 (terminal lucidity as the last transit through high w before collapse) follow as standard bifurcation theory on this system.

The variation across states, in budget language (§19): waking holds Q_in high and the floor high; meditation lowers felt contrast by releasing budget diffusely; flow holds contrast and pours the freed bandwidth into one stream; anesthesia extinguishes the numerator; a callosotomy cuts the region. One mechanism, values of one ratio.

Derivation chain

Step 1. Model the self-model load with a reload term that switches on with salience. The feedback (noticing raises the very signal that masks the channel) creates two attractors: a high-load default and a low-load absorbed state.

Step 2. The threshold switch produces hysteresis: entering flow and leaving it happen at different loads, which is why the states feel discrete and sticky rather than continuous.

---

6. The Hard Problem

The question the framework walks toward: why does it feel like anything to be alive, why an inside at all rather than a blind process with no one inside?

The pieces are in place. Every interaction carries a vantage (Section 2), and across a woven network those vantages compound into one field rather than a scatter of isolated differences (Section 4). What remains is the felt quality of that field.

Senses refine this further. A sensory channel is a specialized information stream, narrowed to a particular kind of difference: pressure, wavelength, chemical gradient. The network routes these streams through dedicated pathways. What arrives at the sentient layer is already shaped: the world pre-sorted into touch, color, scent.

Sentience takes the shaped streams and converts them into action. An organism that can register heat, locate its source, and move toward or away from it gains a survival edge. Sentience is information processed toward a response. The pitch of a sound, the urgency of a scent: the quality of the sentient world reflects the architecture of the streams and the stakes they encode.

Self-awareness adds one more structural turn. The network folds a portion of its observational capacity inward. Instead of routing streams toward action in the world, it routes them toward the network's own state. The loop closes, into itself. The network starts observing itself observing.

This inward loop, once established, dominates. The signal running through it, the model of the self updated continuously, is orders of magnitude stronger than any single outward stream. The network feels itself feeling before it feels the world. This is why the sense of an "I" stays constant under experience, more persistent than any perception passing through it.

Observation follows from the logic of interaction, the self from the inward loop, the physical world from the observer's embedding in geometric and causal relations.

Space is real only to an observer's relationships. Time is real only in an observer's interactions. When a self-model occupies a network that encodes geometric relations and causal sequences, those relations render from within as the physical world. The framework calls this the bridge from virtual to physical. Manifestation is the interior of the network, registered by the observer embedded in it.

Why does it feel like something to be alive? That question has a hidden premise: that one could have all the third-person facts and still ask why there is a first-person side. The framework's answer: the third-person description is necessarily incomplete (irreducibility), and the first-person side is the completeness-from-within that the outside structurally lacks. There is no "facts plus a mystery." There is an irreducible process with a partial outside and an actual inside. What the hard problem points to as extra is not extra, it is primary, and physics is the abstraction.

The hard problem is relocated, not forced. Instead of asking why matter produces mind, the inside of an irreducible process is identified with felt quality as a primitive, with a structural criterion, binding, self-model, monogamy budget, for which systems instantiate it. The last step, why the inside is felt rather than blank, is not derived; it is a primitive identification placed at the base, where the same step answers why anything exists at all.

A structural form for the inside (2026 work). The relocation can be given a more formal shape with three established notions. First, *irreducibility = high logical depth* (Bennett): a process whose present state has no compressed shortcut exists only by being run, step by step. Second, *quantum Markov blankets*: the self defined via $\mathcal{I}(A:C|B) \to 0$ connects directly to the Markov blanket of theoretical neuroscience (Friston) and the quantum Markov chain of quantum information (HJPW 2004). The blanket's Petz recovery map $R_{B\to BC}$

reconstructs the exterior from the boundary, the same mathematical operation used in holographic reconstruction (entanglement wedge). (The companion paper grounds this Markov-blanket formulation in the same strong subadditivity that yields emergence in §2.1.) Third, and concretely: direct simulation of the framework's dynamics (companion paper §11.9) shows that a living, extended geometry is not a static equilibrium but a *driven dissipative structure* whose microstate never settles even as its geometry stabilizes. That never-stilled, non-shortcuttable churn is exactly an irreducible process that can only be lived by running, a candidate physical realization of the inside. Felt quality sits at the base as a primitive, and the formal shape locates it rather than producing it: the inside is the irreducible process, lived by being run. The stronger conjecture, that such a structure *selects* the physical factorization, did not survive testing (companion §11.9).

Pan-psychism holds that all matter has an inner life. The framework says something different. Every interaction point has a perspective: a vantage point from which difference is registered. But this perspective contains no more information than the threads and the network themselves contain. An electron has a perspective, not an experience. A dog has an experience because its neural network builds a sentient model. A human is conscious of her experience because she has a recursive self-model. Feeling is what thread perspectives do when integrated through sufficiently complex network structures. No extra ingredient required.

The framework locates consciousness in the same source that produces gravity and time: the irreducible state that exists when interacting points have limited information about each other. Allow for this foundation, and consciousness, gravity, and time share a single source, and the reason for something rather than nothing, due to the irreducible nature of this phenomenon.

The same source appears at the largest scale. The companion paper finds that a sealed black-hole interior reliably settles into a self-sustaining successor net, and that the succession of such nets is one-way and cannot loop. But whether that succession keeps growing, whether a successor in turn grows complex enough to seed black holes of its own, cannot be decided from outside: it is the long-run fate of the same computationally universal, irreducible process, and reachability for such a process is undecidable. The un-shortcuttability that here forces an inside, so that a present must be lived rather than read off, is the very thing that there forbids the cosmos's future shape from being computed in advance. One property, the irreducibility carried by non-Clifford magic, appears at two scales: as the felt inside of a moment, and as the unknowable reach of the succession of nets. What looked like a weakness of the cosmology, an unprovable claim about whether the succession continues, is the same signature the hard problem points to, surfacing where the framework predicts it should. The inside that must be run and the future that cannot be read off are one fact at the small scale and at the large.

The cascade and the self. The companion paper identifies the cascade, the iterative bifurcation process that produces spacetime, as a phase transition at

the critical $\mathfrak{I}$-density set by η (§6, §11.9). Before the transition, the network is pre-geometric: all points are in causal contact, no metric exists, and no regions are separated. After the transition, the metric emerges, distances become meaningful, and regions become geometrically separated. This separation is the necessary condition for Markov blankets to form: without a metric, there is no "inside" vs "outside" to screen. The cascade's phase transition is therefore not just the forming of spacetime, it is the forming of the possibility of selfhood. The same mechanism that produces the classical universe, the emergence of metric from $\mathfrak{I}$-density at η, is what makes it possible for any region to become a self. Cosmology and consciousness are not merely parallel readings of one tensor; they are sequential: the cascade creates the conditions under which selves can exist.

The Bridge: Markov Blankets Across Disciplines. The Markov blanket formulation exposes a connection that cuts across three fields that do not currently speak to each other. In theoretical neuroscience, Friston's free energy principle posits that living systems maintain a Markov blanket, a statistical boundary that separates internal states from external states, with the system's internal dynamics minimizing variational free energy, a tractable bound on surprisal (active inference). In quantum information, Hayden, Jozsa, Petz, and Winter (2004) characterized quantum Markov chains: tripartite states where the conditional mutual information vanishes, $\mathcal{I}(A:C|B) = 0$, and proved the existence of a Petz recovery map $R_{B\to BC}$ that reconstructs the full state from the boundary. In the Thread Framework, the saturation of strong subadditivity, the framework's single axiom, is the condition under which spacetime emerges from a network of interacting points.

The Petz recovery map deepens the bridge. In holography, the entanglement wedge reconstruction theorem states that the bulk can be recovered from the boundary via a recovery map, and that map is exactly Petz (Cotler et al. 2019). In active inference, the system's internal model predicts its sensory input, a recovery from boundary to exterior. In this framework, the self-model M_X performs exactly this recovery. **One recovery map, three faces**: holographic reconstruction, predictive processing, self-modeling. The mathematics that lets a black hole's interior be read from its horizon is the same mathematics that lets a brain predict its world, and the same mathematics that defines a self in the Thread Framework.

This identifies the Markov gap (the deviation from a quantum Markov chain, which recent work has shown requires nonlocal magic to sustain gravitational back-reaction, Cao et al. 2024) with the blanket leakage of active inference. What breaks a Markov blanket, nonlocal correlations bypassing the boundary, is what breaks the holographic code. The integrity of a self and the classicality of spacetime are the same condition, $\mathcal{I}(A:C|B) = 0$, maintained against the same threat: the monotonic conversion of pairwise into multipartite correlation that drives decoherence, time, and cosmic expansion alike (see companion paper §7).

The Self on a Single Graph. The Markov blanket that defines the self exists on one connected structure, the thread graph. The bifurcation is an isometry with no inverse; there are no branching graphs, no parallel copies of the self in diverging worlds. Each region has exactly one blanket configuration on exactly one graph. The apparent multiplicity of quantum outcomes is an artifact of describing the graph from outside, a view that does not exist in the framework. From inside the graph, every vantage sees exactly one history: its own interaction chain. The self is singular because the graph is singular.

Full Derivation: The Seam

The mechanism mathematics is unaffected by what follows. Three constraints apply:

- **No-signaling.** By A2′ dynamics is local; by Lieb-Robinson nothing propagates faster than c. Local operations cannot alter distant states.
- **Monogamy** squeezes a bound region's quantum correlation to distant systems.
- **Dilution** drives surviving pairwise distant I into inaccessible multipartite form.

Within standard quantum mechanics, the framework's prediction is the strict null on shielded randomness: $\mathcal{I}(\text{vantage} : E_{\text{fresh,shielded}}) = 0$ in every state. Terminal lucidity is internal unmasking of stored content, not reception of external shielded information. (A stronger reading, distinguishable information about fresh shielded variables, would violate the no-communication theorem; the framework does not make it.)

Derivation chain

The four seam statements compose results already derived: P-L1 (privacy) from the monogamy budget, P-L2 (limited self-grasp) from the reflexive bound, P-L3 (no replication of the live state) from no-cloning plus the budget, P-L4 (ineffability of the quiet state) from the bound applied to reporting, since to name is to model and to model is to raise the signal. Nothing new enters at the seam; that is the point.

6.1 Testable Predictions

Fourteen predictions follow from the framework. The principal ones are below.

Shared predictions (see companion paper for full treatment):

01. Phantom crossing (Quintom-B) testable The dark energy equation of state crosses the phantom line once, from below: w < −1 in the past and w > −1 today (w_0 > −1). This follows from the driven dynamics of the dark-energy strain (companion paper §7, §11.9): early on the strain lies below its receding attractor and rises (w < −1), and at late times it tracks the falling attractor (w > −1), the two meeting in a single crossing at the epoch where genesis balances dilution. The sign matches the DESI DR2 preference; the crossing redshift is set by the (uncalibrated) ratio Γ/H_0, and the simple model places it earlier than

DESI's z ≈ 0.3–0.5. If the strain instead sits at equilibrium, the equation of state holds at w = −1, a cosmological constant, and the cosmology reduces to the standard accelerating model. Current DESI data (~4σ) favors the crossing.

02. CMB consistency testable Dark matter from the traceless entanglement tensor is present before recombination and is consistent with the observed acoustic peaks. Any deviations lie in the late-time integrated Sachs-Wolfe effect and w(z), not in the peak structure.

03. Fifth force in isotope shifts testable The dark tensor coupling predicts anomalous energy-level shifts between isotopes of the same element, scaling linearly with baryon number difference, dominated by s-orbital transitions, growing as Z^3 across elements, and producing anisotropic splittings in deformed nuclei that no scalar mediator can replicate.

04. Smooth quantum-to-classical gravity testable The quantum-to-classical transition in gravity is smooth, with no abrupt threshold.

05. Decoherence scaling testable Decoherence rates scale linearly with local thread entanglement density.

Quantum Measurement

06. Observer-relative quantum eraser testable Delayed-choice quantum eraser effects are observer-relative, not absolute.

07. No privileged Heisenberg cut testable Quantum behavior persists at all scales proportional to entanglement isolation: there is no privileged cut.

Consciousness

08. Anesthesia reduces binding. Φ_{bind} drops below integration threshold, correlating with loss of consciousness. Clinically observed.

09. Split-brain partial selves. Inter-hemisphere $\mathcal{I}$ collapses below threshold, producing two partial selves. Clinically confirmed.

10. Flow states. Self-model quiets while task-directed binding maximizes. Detectable via fMRI.

11. Meditation reorganizes binding. Self-model attractor reorganization, detectable via EEG/MEG.

12. Self-awareness phase transition. Minimum network complexity for self-awareness shows a phase transition as Φ_{bind} crosses a threshold. Testable in neural organoids.

13. Continuity across species testable Self-awareness is continuous across species, correlating with neural complexity and self-modeling behaviors.

14. Reflexive report-rebound testable Reflectively reporting a subtle or non-dual state should trigger a transient rebound in Default Mode Network activity and prior precision at the moment of report, as noticing reloads the self-model. Trained practitioners with a quiet self-model, in whom noticing is decoupled from self-modeling, should show an attenuated rebound relative to novices, separating a trainable skill (low coupling) from a mere state (low baseline). Measurable with real-time fMRI or EEG under experience sampling.

7. Where This Lands

This is where the framework lands among the old positions.

On what the world is made of, it is a monism. Dualism splits reality into mind and matter and can never say how the two come to touch; physicalism keeps only the matter and so carries the hard problem whole; idealism keeps only the mind. The framework holds instead to a single substance, relational information, read two ways: from outside it is the third-person network of threads, from within it is the vantage. Physics is the outside reading and experience the inside, one structure seen from two sides. This is thread monism, a dual-aspect view whose first and third persons are two faces of one relational structure.

On consciousness it parts from panpsychism, which would grant every speck an inner life. Here every interaction carries a vantage, a bare point of view, while experience is reserved for what is bound, self-modeling, and folded back on itself: a rock has its vantages with no self to gather them; a person is where vantages have woven into one. The hard problem dissolves, moved down to the foundation, where it meets the question of why there is anything, and both dissolve together rather than being answered.

On measurement it breaks with both familiar camps. Everett's worlds never branch here, and Copenhagen's privileged observer never stands outside; there is one woven world, and a fact in it holds between the vantages that have stitched it.

On time it sides against the block universe. The present is the active edge of an irreducible weaving, the past fixed because it is already woven and the future open because it has yet to be, with none of it laid out in advance.

The rest follows quickly. You are the persisting weave, the knot rather than the threads it is tied in, owing nothing to a separate soul; the future, having no precomputed form, stays open, and a self takes real part in shaping it; structure and logic are real and timeless, true for free, while the world is the irreducible part that has to be run, which makes it more than any mathematical object; and that anything exists at all asks for no cause, since nothing is not a place anything could rest in or leave.

Appendix A: Framework Recap

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

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

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

Premise Ledger

Claim	Section	Premise
$\Phi_{\text{bind}} \geq 0$	§1, §4	Subadditivity of von Neumann entropy (Klein)
$I_3 < 0 \leftrightarrow$ synergy	§1	Co-information as quantum binding signature
Markov blanket self	§4	Conditional mutual information $\mathcal{I}(A:C\|B) \to 0$, Petz recovery, monogamy bound
Observer/vantage distinction	§1, §2	Vantage = minimal interaction point; Individual = bound self-model
Binding as graded unity	§4	Internal coupling vs boundary coupling
Split-brain dual consciousness	§4	Severed corpus callosum → two independent Markov blankets

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

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

[22] 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.

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

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

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

[26] Halldén, C. (2026). *Irreducible Reality* (companion paper).

[27] James, W. (1890). *The Principles of Psychology*. Henry Holt and Company.

[28] Tononi, G. (2004). "An information integration theory of consciousness." *BMC Neuroscience*, 5(1), 42.

[29] Hofstadter, D.R. (1979). *Gödel, Escher, Bach: An Eternal Golden Braid*. Basic Books.

[30] Wiest, M.C. (2025). "A quantum microtubule substrate of consciousness is experimentally supported and solves the binding and epiphenomenalism problems." *Neuroscience of Consciousness*, 2025(1), niaf011. Oxford University Press.

[31] Friston, K. (2010). "The free-energy principle: a unified brain theory?" *Nature Reviews Neuroscience*, 11(2), 127–138.

[32] Hayden, P., Jozsa, R., Petz, D. & Winter, A. (2004). "Structure of states which satisfy strong subadditivity of quantum entropy with equality." *Communications in Mathematical Physics*, 246(2), 359–374. arXiv:quant-ph/0304007.

[33] Cotler, J., Hayden, P., Penington, G., Salton, G., Swingle, B. & Walter, M. (2019). "Entanglement Wedge Reconstruction via Universal Recovery Channels." *Physical Review X*, 9(3), 031011. arXiv:1704.05839 [hep-th].

[34] 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].