The Boundary is a Rate: Coherence–Rupture–Regeneration as a Temporal Physics of Emptiness Realisation

Alexander Sabine
Independent Researcher
alexander@activeinference.institute

June 2026 (revised July 2026)

Abstract

Sandved-Smith, Fields, Doctor, Laukkonen and Hohwy argue that any system which persists must minimise surprisal, and so gather evidence for its own generative model, but that one dimension of this self-evidencing is provably impossible: evidencing the boundary that would constitute the agent as separate from its environment. Because no finite system can measure the entanglement entropy across its own boundary from inside, the separability of agent from environment is permanently unevidenceable, and they propose that the realisation of emptiness is the embodied recognition of this impossibility. We take their result to be correct and do not contest it. We ask what survives it when the self/environment boundary is read not as a static holographic screen but as a *rate*: the moving locus, in the Coherence–Rupture–Regeneration (CRR) framework, at which accumulated Fisher information saturates the Cramér–Rao bound, $C \cdot \Omega = 1$. Embedding CRR's three equations in the quantum-reference-frame (QRF) formalism the original argument employs, we find that the rupture instant is itself screen-inducing, so CRR inherits the impossibility rather than escaping it; that the single point at which the framework fails the formalism's unitarity requirement is, on a measurement-theoretic reading, a signature rather than a defect; that the forbidden spatial correlation is robustly bounded by, though not identical to, an accumulated temporal quantity; and that the convergence of many finite systems produces a genuine, macroscopic collective now through a synchronisation transition we establish by finite-size scaling against validated benchmarks. The contribution is not a refutation but a relocation: emptiness moves from an epistemic limit on measurement to an ontological claim about a present with no content of its own, met by the selectively weighted past of every other finite system. All results are consistency demonstrations under a chosen formalism, not empirical findings.

Keywords: self-evidencing; emptiness; quantum reference frames; holographic screens; free energy principle; Fisher information; Cramér–Rao saturation; synchronisation.

1 Introduction

To persist is to keep confirming a model of what one is and what surrounds one. In the free-energy principle (FEP), a system that endures *self-evidences*: it minimises the surprisal of its sensory states and thereby accumulates evidence for its own generative model. Sandved-Smith and colleagues identify a single belief this process can never reach. To confirm that one is *separate*—bounded off as an entity distinct from one's environment—would require quantifying one's entanglement with everything one is not, and a finite system provably cannot measure the entanglement entropy across its own boundary from within. The separability of agent from environment is therefore not merely unconfirmed but structurally *unevidenceable*; and they propose that awakening, as the stable realisation of emptiness, is the embodied recognition of exactly this impossibility.

We find the argument compelling and, on our reading, its central result correct. What we examine is a picture it leaves intact: that behind the unmeasurable boundary there remains a determinate environment, a far side with which one is entangled but whose correlation one can never quantify. Coherence–Rupture–Regeneration (CRR) is a temporal framework in which bounded systems accumulate evidence, reach the limit of a regime, and reconstruct from their own history. Here we turn it on the boundary the impossibility theorem concerns, and ask what emptiness leaves standing once that boundary is read as a rate rather than a surface.

The claim in one sentence. We agree the impossibility theorem holds, inherit it rather than escape it, and then make a single ontological move: behind the unmeasurable boundary there is no determinate far side—only the kernel-weighted past of every other finite system, met at a present that has no content of its own. Everything below is the unpacking of that sentence, reported with the restraint the original paper models: where the two accounts converge, where they diverge, and what the divergence costs each of them.

We should say plainly where the reading originates. The formal work below stands or falls on its own, but the conviction that the present is empty in a hospitable rather than a privative sense was not arrived at from the mathematics; it is the kind of thing one meets before one proves, and the equations that follow are an attempt to say precisely what was met.

This paper is self-contained: a reader who has met neither the quantum free-energy programme nor CRR should find every term defined as it enters. Section 2 sets out CRR. Section 3 reconstructs the impossibility argument in the authors' terms. Sections 4–8 present the embedding of CRR in the QRF formalism and the readings it licenses. Section 10 draws the ontological consequences and states the limitations.

2 Coherence–Rupture–Regeneration

CRR describes any persisting process—a wound healing, a belief stabilising, an attractor forming—through the accumulation, saturation, and rebuilding of coherence. We use three notions from information geometry, defined here. A system's momentary probabilistic state is a point on a *statistical manifold*; the natural way to measure distance on that manifold is the Fisher–Rao metric, under which the *Fisher–Rao speed* $L \geq 0$ is the local rate at which the system accrues distinguishable structure. *Coherence* accumulates as the integral of that speed along the system's trajectory:

$$C(x,t) = \int_0^t L(x,\tau)\, d\tau, \tag{1}$$

which, under the path-integral formulation of the FEP, is accumulated arc length along the information-geometric trajectory. The content of the present is, by this definition, exactly what has been integrated over the past; the upper limit t contributes nothing of its own—an observation, trivial as a statement about an integral, that becomes load-bearing in Section 5.

Coherence cannot accumulate without bound. The Cramér–Rao bound sets the smallest variance any unbiased estimator can attain from given information; when accumulated coherence saturates it, the process *ruptures*—a discrete event of zero temporal extent, carrying one unit of mass on the boundary between accumulated past and regenerated future:

$$\delta(\text{now}) \quad \text{when} \quad C \cdot \Omega = 1. \tag{2}$$

The dimensionless product $C \cdot \Omega$ gives (2) the same algebraic form as the Heisenberg relation, the Gabor limit, and the thermodynamic uncertainty relation; the kinship is qualified in Section 4. After rupture the process regenerates from its own history through a coherence-weighted memory kernel,

$$R(x,t) = \int \varphi(x,\tau)\, e^{C(x,\tau)/\Omega}\, \Theta(t-\tau)\, d\tau, \tag{3}$$

in which the Heaviside step Θ enforces that only the settled past contributes, and $e^{C/\Omega}$ is the maximum-entropy weighting consistent with a constraint on $\langle C \rangle$. The selectivity of φ and the asymmetry of the kernel matter to the divergence from the screen picture developed in Section 5: regeneration prehends the past, it does not replay it. In plain terms: a system builds up structure until it can build no more, breaks, and remakes itself out of what it has been—leaning hardest on the moments that mattered most.

The one parameter is fixed by geometry. The lone parameter Ω is not fitted. By Čencov's uniqueness theorem the Fisher–Rao metric is, up to scaling, the only metric on a statistical manifold invariant under sufficient statistics; the geometry of distinguishability is therefore fixed once the manifold is named, and we take $\Omega = 1/\ell$, the inverse geodesic extent. A *bistable* process inhabits the Bernoulli manifold—a $\mathbb{Z}_2$-symmetric interval of geodesic diameter π—giving $\Omega_{\mathbb{Z}_2} = 1/\pi$ and rupture coherence $C^* = \pi$. A *rotational* process inhabits the circle S^1—an $SO(2)$-symmetric loop of circumference 2π—giving $\Omega_{SO(2)} = 1/2\pi$ and $C^* = 2\pi$. The ratio $C^*_{SO(2)}/C^*_{\mathbb{Z}_2} = 2$ exactly, forced (as Section 2.1 shows) by the fact that the bistable manifold is the semicircle of the rotational one, so the factor of two carries no free parameter once the Fisher–Rao normalisation is fixed.

The same geometry fixes not only the rupture coherence but the variability of rupture *timing*: the coefficient of variation of the inter-rupture interval is $\mathrm{CV} = 1/2\pi \approx 0.159$ for a bistable system and $1/4\pi \approx 0.080$ for a rotational one—a parameter-free signature we return to in Section 9. Finally, Ω acts as a temporal light cone: small Ω peaks the memory kernel (rigid, repetitive reconstruction), large Ω flattens it (flexible, exploratory)—a role analogous to Levin's cognitive light cone and to the scaling of care that Doctor and colleagues place at the root of intelligence, here quantified from the geodesic extent of the manifold itself. This single invariant is, as Section 4 shows, the cleanest survivor of the embedding into quantum reference frames.

2.1 The rupture is $\mathbb{Z}_2$ by construction

The two symmetry classes are easy to read as alternatives—a system is bistable *or* rotational. Consider the breath, which carries both shapes at once and so fixes the difference exactly. Whether one is inhaling or exhaling is a *bistable* question: two ends, an in-breath and an out-breath, with the airflow reversing between them—the system walks a corridor between two states. But the respiratory cycle *as a whole* has no ends: exhalation gives onto inhalation with no wall between, and the phase runs a loop that closes on itself—a *rotational* question. The shape of the question, corridor or loop, is all that separates the two classes. For the boundary the impossibility theorem concerns, their relation is not parallel but *nested*, and the geometry makes the nesting exact.

Under the standard Fisher–Rao amplitude embedding, a binary parameter $p \in [0, 1]$ is sent to $(\sqrt{1-p}, \sqrt{p})$; the Fisher metric becomes the round metric on the unit circle, and the Bernoulli manifold traces a geodesic arc of length exactly π: the geodesic coordinate $\phi(p) = 2 \arcsin \sqrt{p}$ sweeps $0 \to \pi$ as p runs $0 \to 1$. The bistable manifold is therefore not merely the same length as half of the $SO(2)$ circle of circumference 2π; it *is* a semicircle of that circle (Figure 1, left). Writing the same manifold in the sign coordinate $m = 2p-1 \in [-1, +1]$ makes the $\mathbb{Z}_2$ reflection $m \mapsto -m$ explicit, fixed at the maximally-uncertain midpoint $m = 0$ ($\phi = \pi/2$), with the two pure states at the antipodes $\phi = 0, \pi$.

We take the two classes in nested roles. The circle S^1 is the continuous temporal substrate: no boundary, no preferred phase, no point singled out as the now—pure duration. This is the geometric form of the claim, load-bearing in Section 5, that the present has no content of its own: a circle carries no marked point. The rupture is the crossing of the interval—a directed passage from a settled past to an opening future, whose two endpoints are its $\mathbb{Z}_2$ structure. The circle supplies duration; the crossing supplies the cut (Figure 1, right). They do not compete; they compose.

Two uses of one geometry must therefore be kept apart: the *shape of a system's question*—corridor or loop—fixes ℓ and hence Ω and varies by system; the *two aspects of every rupture*—the continuum that holds, the cut that falls—never vary. We accordingly write *bistable* and *rotational* for systems, reserving $\mathbb{Z}_2$ for the structure of the cut and $SO(2)$ for the continuum. The breath, being rotational, completes a full turn between ruptures, its inhale and exhale the two faces of that crossing. This also settles why the factor of two is forced rather than fitted: the circle is not independently assigned a metric and then found to measure 2π; it is the Bernoulli arc completed into a full turn, so $C^*_{SO(2)} = 2\,C^*_{\mathbb{Z}_2}$ follows from the doubling alone.

The same nested structure carries a reading that contemplative and process traditions have drawn in their own terms, which we note because the geometry underwrites it without depending on it. In the Madhyamaka register the two poles are prajñā and upāya: wisdom, which sees the empty, conditioned, groundless nature of things, and skilful means, the compassionate act that meets a particular situation. The continuous circle—no preferred point, no marked now—is the geometric figure of prajñā, the open ground; the discrete $\mathbb{Z}_2$ crossing, a directed act that cuts a settled past from an opening future, is the figure of upāya, method enacted in time; and their inseparability—no crossing except on the holding continuum, no continuum marked except by a crossing—is the classical claim that wisdom and method are never complete alone. The Whiteheadian register names the same pair as the rhythm of an actual occasion: the circle is the concrescing process, the becoming that has no determinate datum while it is underway, and the rupture is *satisfaction*, the moment the private process perishes into a public fact the next occasion inherits. We rest no formal weight on these correspondences; the single fact of information geometry above—that the bistable manifold is the semicircle of the rotational one—is what each can be seen as pointing at.

Why this matters here. In a companion study the same framework was tested in a concrete active-inference architecture, where its single topological parameter produced a measurable, falsifiable structure in the timing of belief updates—a result independent of the implementation and consistent across environments [5]. That CRR earns a falsifiable prediction in a concrete setting is what licenses turning it, here, on the boundary the impossibility theorem concerns. Where the FEP's Markov blanket separates inside from outside in space, CRR's rupture condition delineates past from future in time; the two are spatial and temporal siblings of the same partition.

In a nutshell **The mathematics.** Coherence C accumulates as (1); at $C \cdot \Omega = 1$ one discrete cut falls between past and future, (2); regeneration, (3), rebuilds from history, weighting the most coherent moments most heavily. $\Omega = 1/\ell$ is fixed by geometry: a bistable system cuts every π of accumulated arc, a rotational system every 2π; the factor of two is the ratio of their rates on the shared circle.

The karmic reading. Each cut is two-sided. Its past-face is *phala* in Hershock's gloss, fruit ripe to the point of bursting: the concentrated expression of accumulated conditions, which is exactly what saturation is. Its future-face is opportunity, the opening through which redirection enters. prajñā is the unmarked circle; upāya is the cut laid upon it.

3 The unevidenceable boundary

We reconstruct the impossibility argument at full strength before reading it, because a commentary that weakened it first would be worth nothing. It is built on the same self-evidencing premise CRR shares with the FEP, and proceeds in four moves.

First, the premise. Any system that persists must minimise the surprisal of its sensory states, equivalently maximise the evidence for its generative model; to exist is to self-evidence.

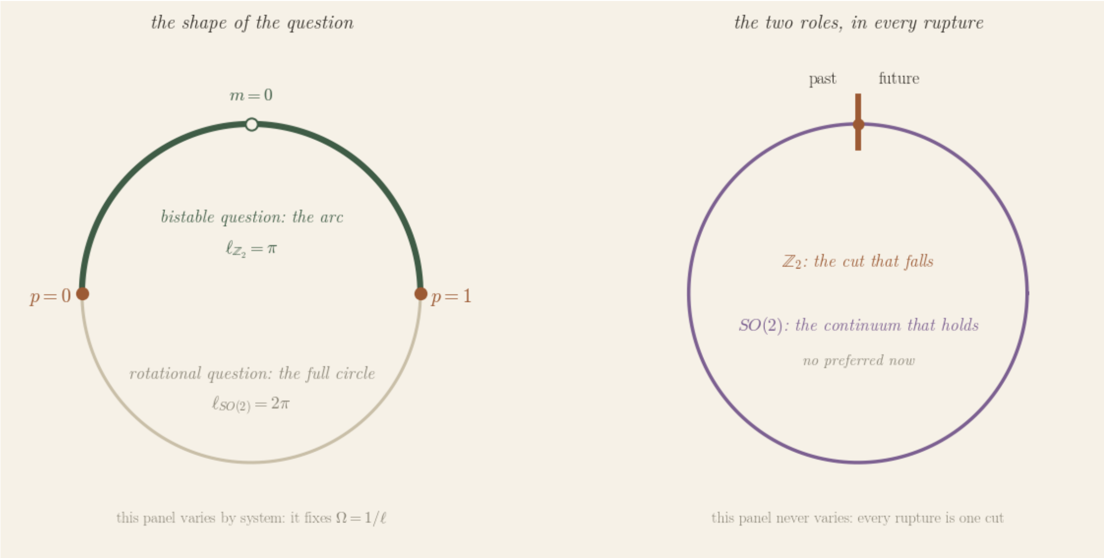

Figure 1: Two uses of one geometry. *Left*, the shape of the question: under the Fisher–Rao amplitude embedding the bistable (Bernoulli) manifold is a geodesic arc of length exactly π—a semicircle of the rotational circle of circumference 2π, with the reflection $m \mapsto -m$ fixed at the midpoint and the pure states at the antipodes. A bistable system accumulates π of arc between cuts, a rotational system 2π; the factor of two is forced by the embedding. This panel varies by system: it fixes $\Omega = 1/\ell$. *Right*, the two roles: in every rupture, the $SO(2)$ continuum holds (no preferred now) and a single $\mathbb{Z}_2$ cut falls (past | future). This panel never varies.

Second, the one unreachable belief: among all the beliefs a self-evidencing system maintains, exactly one can never receive evidence—that there is a boundary constituting it as separate from its environment. *Third, the quantum-information ground.* To confirm separateness would require measuring the *entanglement entropy* across one's own *Markov blanket*—the statistical boundary of sensory and active states that screens a system from its surroundings—quantifying how correlated one is with everything one is not; and recent results in the quantum-FEP programme establish that no finite system can perform this measurement across its own boundary from inside. The authors formalise the belief in separation as a *structural prior* over the agent's *quantum reference frame (QRF) deployments*—the choices of measurement basis available to it—constraining every frame the agent can adopt to respect a self/environment partition: the partition is built into the apparatus of measurement and is never among the things measured. *Fourth, the interpretive move:* if separation is a prior that can never become a posterior, then awakening, as stable realisation of emptiness, can be understood as the embodied recognition of this impossibility—not a new belief acquired but the felt collapse of one that was never earning its keep. The authors mark this fourth move as interpretive, and we respect that marking.

It is worth being precise about the kind of result this is, because it conditions everything below. It is an impossibility theorem embedded in a chosen formalism: its truth is the truth of what the QRF formalism forbids—a strong and honest kind of truth, but not the same kind as a measurement of the world, and the authors do not pretend otherwise. Keeping that distinction in view is what allows the CRR reading below to be a genuine response rather than a rival assertion.

In a nutshell **The mathematics.** No finite system can measure the entanglement entropy across its own boundary from inside; separateness is therefore a permanent prior, never a posterior—true by construction within the QRF formalism, in the way no-cloning is true.

The karmic reading. anattā arrived at epistemically: awakening is not a new belief acquired but the felt collapse of one that was never earning its keep. What a finite being can know is never its own separateness but only, in the spirit of yathābhūtam, what has come to be—which Section 5 will make the whole of what there is to meet.

4 CRR in the quantum reference frame formalism

Following the method the original programme itself uses—and its operational companion, in which Friedman renders the same qFEP and QRF equations as finite operational surrogates checked by positive controls and by discriminating tests that reject stronger readings [2]—we do not test CRR against data here. We embed its three equations in the QRF and holographic-screen formalism—the description of a boundary as a finite channel carrying a bounded number of distinguishable sectors—and ask what that formalism forces, forbids, or over-determines. Four findings bear on the impossibility theorem; for each we distinguish what the computation settles from what it can only illustrate.

The rupture is screen-inducing. The prior question is whether CRR's rate-boundary possesses a screen at all. Constructing an explicit bipartite system $A \otimes E$ and driving an internal coherence to saturation, we find that the reduced state ρ_A remains a valid density matrix—positive semidefinite, unit trace to machine precision (4×10^{-16})—with finite von Neumann entropy at every saturation level up to and including the rupture instant, where $S(\rho_A) = 2$ bits, the maximum for the four-dimensional subsystem (Figure 2). The bipartition is well-defined precisely when $C \cdot \Omega \to 1$. CRR inherits the QRF machinery, and with it the impossibility theorem, rather than evading it. Whatever account of emptiness CRR offers cannot be an escape from the unmeasurability the original paper proves; it must be something else.

Saturation, and a prediction it forces. The condition $C \cdot \Omega = 1$ is a genuine saturated Cramér–Rao bound, verified as the equality a maximum-likelihood estimator attains for a Gaussian family. The bistable:rotational ratio of inter-rupture arc comes out as 2 to machine precision (Figure 8, and Section 8), and the timing signature is the parameter-free coefficient of variation introduced in Section 2. Embedded in the QRF picture, this forces a consequence neither framework contained alone: if $\Omega = 1/\ell$ is a geodesic extent and a holographic screen carries a finite number of distinguishable sectors, a rotational system must expose exactly twice the boundary sector-count of a bistable one—its inter-rupture arc is twice as long at fixed sector density. The prediction is a product of the encounter, and it is falsifiable.

Failure against the unitarity requirement—and why it is a signature. The quantum-FEP requires the joint system to be unitary—information conserved, the screen merely relocating it—and establishes an asymptotic equivalence between the FEP and unitarity. CRR's regeneration kernel $e^{C/\Omega}$ does not pass this requirement. Read literally as amplification, it injects of order $C^*/\Omega = \pi^2 \approx 9.9$ nats per regeneration and inflates the state norm by a factor $e^{\pi^2} \approx 1.9 \times 10^4$, four orders of magnitude (Figure 3, left). Read as renormalised selection weights, it conserves total measure but reduces entropy below the input—$\Delta H = -1.34$ nats, from $\ln 12 = 2.485$ nats to 1.141 nats—a basis change toward the highest-coherence moments rather than an information pump (Figure 3, right). As written, CRR cannot be a closed unitary system. This is the same species of finding the original paper produces—the formalism forbidding something the equations took for granted—and it points somewhere specific: if the amplification must be paid for, writing $e^{C/\Omega}$ as drawing on an explicit free-energy budget makes CRR an open system and couples Ω, until now purely geometric, to a thermodynamic quantity. We argue in Section 6

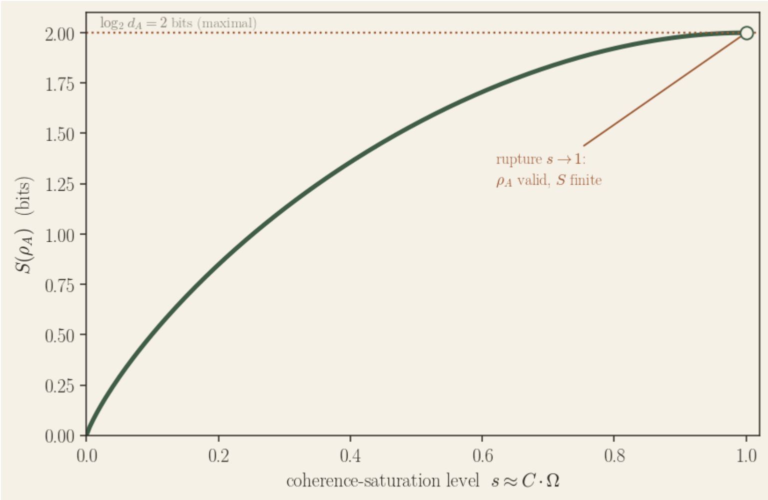

Figure 2: Screen-induction at rupture. The von Neumann entropy $S(\rho_A)$ of the agent's reduced state is finite and well-defined at every saturation level $s \approx C \cdot \Omega$, including the rupture $s \to 1$, where it reaches the maximal 2 bits. A definite self/environment partition exists exactly at the instant the boundary exists; CRR inherits the impossibility result rather than escaping it.

that the failure is more than a bridge: read measurement-theoretically, the non-unitarity is the mark of a collapse, and the budget it owes is a Landauer cost.

The self-measurement obstruction, illustrated. On the same system, the obstruction the impossibility theorem proves in general can be exhibited on examples (Figure 4): the agent's read-out of its own boundary entropy in any measurement frame it can deploy systematically overshoots the true value $S(\rho_A) = 1.626$ bits (mean overshoot $+0.31$ bits across eight arbitrary frames), coinciding with it only in the eigenbasis of ρ_A—which the agent cannot identify without already diagonalising its own state. Recovering the true entropy would demand full tomography and a stable external reference frame, precisely what a finite system lacks for its own boundary. This illustrates, and does not prove, the universal result.

One correspondence CRR has at times stated too strongly should be qualified in the same spirit. The kinship between $C \cdot \Omega = 1$ and the Heisenberg relation holds as an algebraic form, not a literal identification of C, Ω with $\Delta E, \Delta t$: the dimensionless product carries the form of a saturated uncertainty relation, and that is all it carries. The point throughout is to claim the right amount.

In a nutshell **The mathematics.** A definite screen exists at the rupture instant, so CRR inherits the impossibility rather than escaping it; $C \cdot \Omega = 1$ is a genuine Cramér–Rao equality; the encounter forces a falsifiable prediction (rotational screens expose twice the boundary sectors of bistable ones); and the kernel $e^{C/\Omega}$ fails the formalism's unitarity requirement.

The karmic reading. The kernel fails information conservation because prehension is evaluative: regeneration weights the past by what mattered. That is karma as feed-forward as much as feedback—in Hershock's account, karma shapes what is coming to be rather than merely recording what was done—and no value-free, closed unitary evolution can perform that shaping. What a system cares about is written into its kernel.

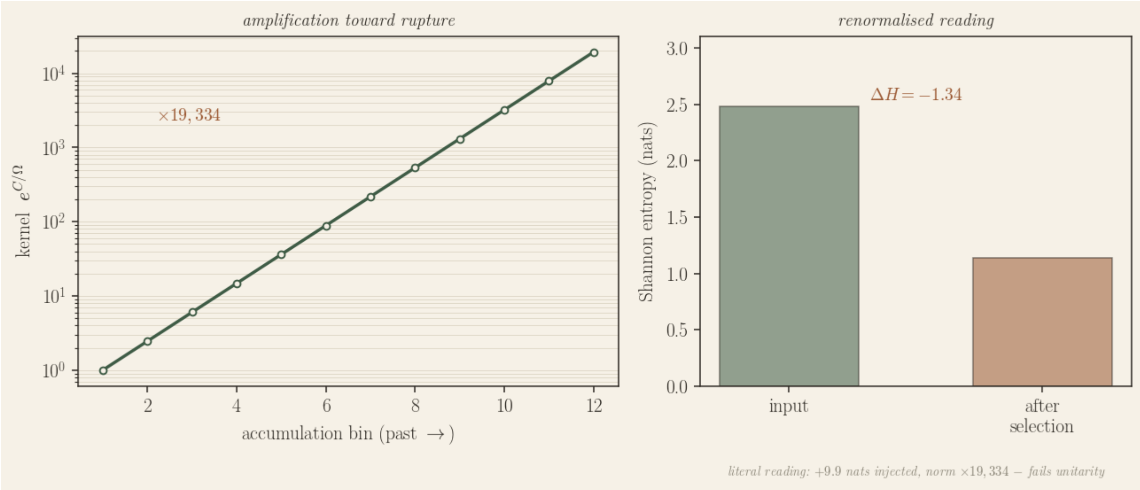

Figure 3: The regeneration kernel is not unitary—a fork. *Left*: $e^{C/\Omega}$ grows by four orders of magnitude across the accumulation toward rupture, peaking at $\approx 1.9 \times 10^4$. *Right*: under the renormalised-selection reading, total measure is conserved but Shannon entropy falls ($\Delta H = -1.34$ nats); under the literal-amplification reading the kernel injects +9.9 nats and inflates the state norm by four orders. Section 6 argues this is a signature, not a defect.

5 The boundary as a rate

The boundary is not a thing you have but a rate you run. The holographic screen of the original programme is static: at any instant a definite self/environment partition exists, and the theorem forbids measuring across it. CRR does not begin with a partition. What plays the role of the boundary is the locus at which coherence is about to saturate and break, and because every system carries its own Ω, its own symmetry class, its own saturation rate, there is no single screen but an interference of many rates meeting (Figure 5). That the boundary is co-constructed at all scales is, in CRR, not a soft phrase but a consequence of there being no single Ω in the universe.

The original claim that you cannot measure across your boundary becomes more specific and more mechanical: the edge one cannot get behind is the front of an accumulation, and one cannot stand behind it because the standing-behind would be further accumulation. The measurement is part of the coherence saturating toward the rupture; one cannot step outside the integral (1) by integrating further.

Here CRR has a claim of its own, and it begins from a fact already present in (1). The coherence $C(x,t)$ is an integral over τ up to the present; content—distinguishability, accumulated Fisher information—is exactly what has been integrated over the past, and the present instant, as the upper limit of integration, carries no content of its own. Only the past has content, because content is defined as accumulated past. So the question the screen account invites—what, then, is on the far side?—receives a different answer. What is out there is not a hidden determinate environment behind a screen; it is the settled past of every other finite system, met at the contentless instant of now.

A seam must not be smoothed over, because it is where CRR earns its divergence. A naïve reading would make the world a flat archive: all completed pasts, equally present. The regeneration term (3) forbids this. The kernel $e^{C/\Omega}$ is selective; a system reconstitutes the past through its own kernel and its own Ω, and two systems meeting in the same now do not encounter a shared archive but each enact the other according to its symmetry class. The past has content, but the content is prehended, never simply given. In this respect CRR is closer to the enactive tradition than to the screen formalism: the world is real and built of the past, but taken up selectively, weighted by whoever is doing the meeting. This is the structure the Madhyamaka

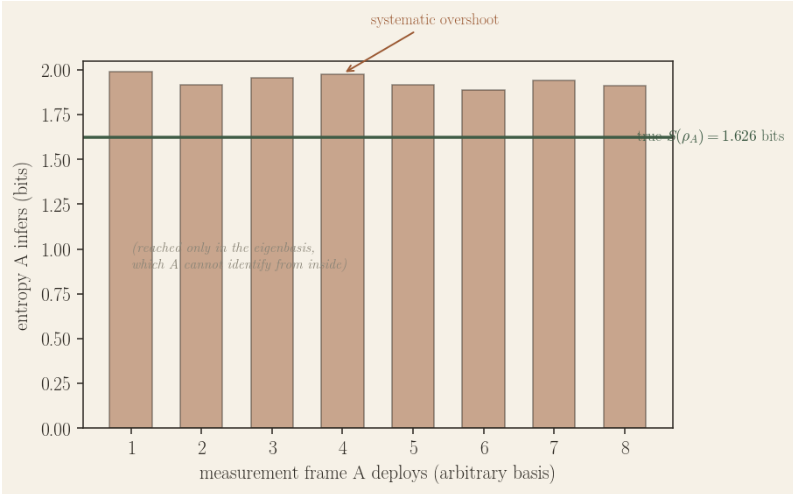

Figure 4: The self-measurement obstruction. The agent's inferred entropy across its own boundary, in eight arbitrary measurement frames, systematically exceeds the true value (solid line, 1.626 bits), reaching it only in ρ_A's eigenbasis—a frame the agent cannot select from inside without having already solved for its own state. The finite-dimensional shadow of the universal impossibility.

calls pratītyasamutpāda, dependent origination—not that the past is unreal, but that it has no character independent of the conditions under which it is taken up. There is no far side; there is only earlier.

In a nutshell **The mathematics.** Content is accumulated past, (1); the present, as the upper limit of integration, has none of its own. The Heaviside step fixes what the past contains; the selective kernel $\varphi\, e^{C/\Omega}$ carries how it contributes. There is no far side; there is only earlier, taken up selectively.

The karmic reading. This is the precise sense in which the past is settled and yet not fixed: its content is closed, its causal contribution is live. Karma, Hershock writes, is less a record etched in stone than an ongoing improvisation, continually reopening for reshaping—in CRR, that reopening is kernel-work, not archive-work. yathābhūtam: meet things as they have come to be, then revise what they will mean.

6 The coincidence of map and territory

We now reach what occurs at the only instant the boundary exists—the rupture, at which $C \cdot \Omega = 1$—and argue that the phrase *the map becomes the territory* is, there, not a metaphor but a close description of Cramér–Rao saturation.

6.1 Saturation is sufficiency is coincidence

A saturated estimator is a *sufficient statistic*: a function of the data carrying everything the data contain about the parameter, with no slack remaining. This is the precise sense in which the inferential map becomes the territory. Below $C \cdot \Omega = 1$ there is a gap, the map (accumulated distinguishability within a model) still chasing the bound set by the territory's informational footprint; at saturation the gap closes to a point, the map holding exactly the territory's footprint—no more (the bound forbids it) and no less (saturation means it is reached). Sufficiency is the coincidence, and it occurs at every now, because each rupture is a fresh saturation; the boundary does not persist but is re-enacted at each $\delta(\text{now})$ (Figure 6). Doctor and colleagues define an agent's *stress* as the gap between its current and optimal state, the inverse of

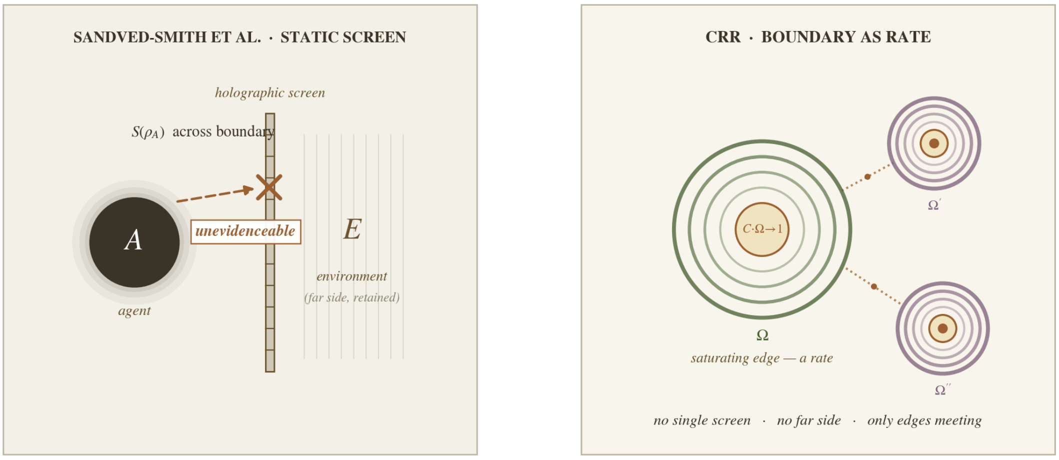

Figure 5: Two pictures of the boundary. *Left* (Sandved-Smith et al.): a static holographic screen separates agent A from a retained environment E; the entanglement entropy $S(\rho_A)$ across it is real but unevidenceable from inside, while a determinate far side is kept. *Right* (CRR): the boundary is the moving front at which $C \cdot \Omega \to 1$; each finite system carries its own rate Ω; there is no single screen and no determinate far side, only the saturating edge, met by other edges.

satisfaction; in CRR that gap is precisely the map–territory gap $1 - C \cdot \Omega$, which closes to zero exactly at satisfaction, $C \cdot \Omega = 1$.

That the boundary exists only momentarily, and is this coincidence, is Whitehead's transition from subject to superject read as a saturation dynamics. An occasion is map-side while it is becoming—prehending, integrating—and territory-side once it has become, a settled datum the next occasion inherits; satisfaction is the instant the private process turns into public fact, and $C \cdot \Omega = 1$ is satisfaction. One cannot remain at saturation: reaching the coincidence is what triggers the rupture and the regeneration that reopens the gap. The objection that this is idealism—the map epistemic, the territory ontic—is met by the refusal of any territory-as-noumenon behind the map. The territory just *is* the settled past (Section 5), itself composed of prior occasions' satisfied inferences, prior maps that became territory; so map-becomes-territory is not the crossing of a metaphysical gap but the temporal relation becoming $\to$ become. The map does not represent the territory at saturation; it perishes into it.

6.2 Rupture as measurement, and the self-measurement seam

This reading changes what the non-unitarity of Section 4 is. In quantum mechanics, unitary evolution is the closed-system, map-side dynamics; the non-unitary part is measurement, the collapse in which a quantum map becomes classical territory. If the rupture is the instant at which map becomes territory, then the regeneration kernel *should* not be unitary—the rupture is a measurement-like event, and CRR's failure of the unitarity requirement is a signature, not a defect. The two readings of the kernel converge thermodynamically: a measurement carries a Landauer cost, and that cost is exactly the free-energy budget for Ω that Section 4 found CRR to owe. Map-becomes-territory, measurement, non-unitarity, and Landauer cost are one event described four ways.

This creates a new problem, which honesty requires naming. If the rupture is a measurement it requires a reference frame, and a system measuring itself into existence at each now would be performing the very self-measurement the impossibility theorem forbids. The escape follows from Section 5: what collapses at each rupture is not the boundary but the past the system measures—the settled occasions it prehends, not its own separability. The boundary is not what is measured; it is the fact that a measurement occurred. This sits exactly inside the impossibility: the agent never measures the entanglement entropy across its own screen (forbidden, and CRR

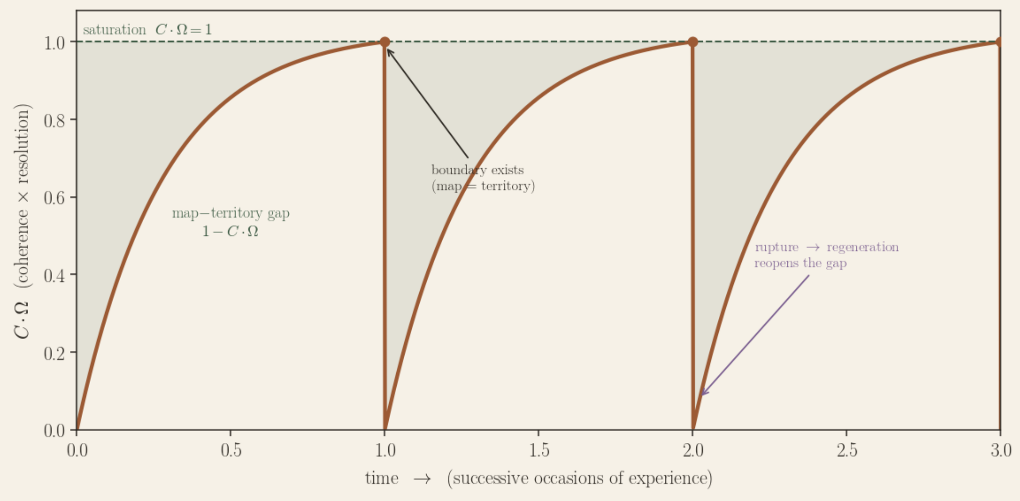

Figure 6: The boundary exists at every now. $C \cdot \Omega$ accumulates within each occasion and saturates to 1 at rupture; the shaded map–territory gap $(1 - C \cdot \Omega)$ closes to zero at each saturation—the instant the boundary exists and map equals territory—then reopens as regeneration begins the next occasion. The boundary is not a standing surface but a momentary coincidence, repeated.

agrees), while what it measures is its prehended past (permitted, and contentful). We do not claim to have proved the thread holds in the full formalism, only that it is the reading the framework's own commitments select.

Remark on status. Three of the facts in this section are theorems or established physics: Cramér–Rao sufficiency at saturation, measurement as the non-unitary step, the Landauer cost of erasure. The bridges between them—rupture as measurement, sufficiency as map-becoming-territory, the temporal reading of the epistemic/ontic gap—are interpretive, in the same register as the original paper's step from a quantum-information theorem to a claim about contemplative realisation. The bridges are load-bearing in the same way theirs are. None of it is a measurement of the world.

In a nutshell **The mathematics.** At $C \cdot \Omega = 1$ the estimator is sufficient: the map holds exactly the territory's informational footprint. That coincidence is the rupture; read measurement-theoretically it is a collapse, hence non-unitary, hence carries a Landauer cost—the free-energy budget the framework was independently found to owe. What collapses is the prehended past, not the forbidden boundary.

The karmic reading. Each now is an axiological turning point: at once outcome and opportunity. The past-face is satisfaction, the map perishing into public fact; the future-face is the reopened gap, where values re-enter through φ. The Landauer cost is the thermodynamic price of commitment: to enact one weighting of the past is to erase the alternatives—karmically, choice is never free of heat.

7 Retensing the forbidden correlation

The impossibility theorem forbids measuring $I(A{:}E) = 2S(\rho_A)$: one's correlation with the co-present complement, quantified across the boundary from inside. But the forbidden quantity is stated over a co-present *spatial* complement, and in CRR there is no co-present totality of not-self—the far side is the settled past (Section 5). This suggests a question: might the correlation that cannot be measured across the boundary be recoverable as a quantity accumulated *along*

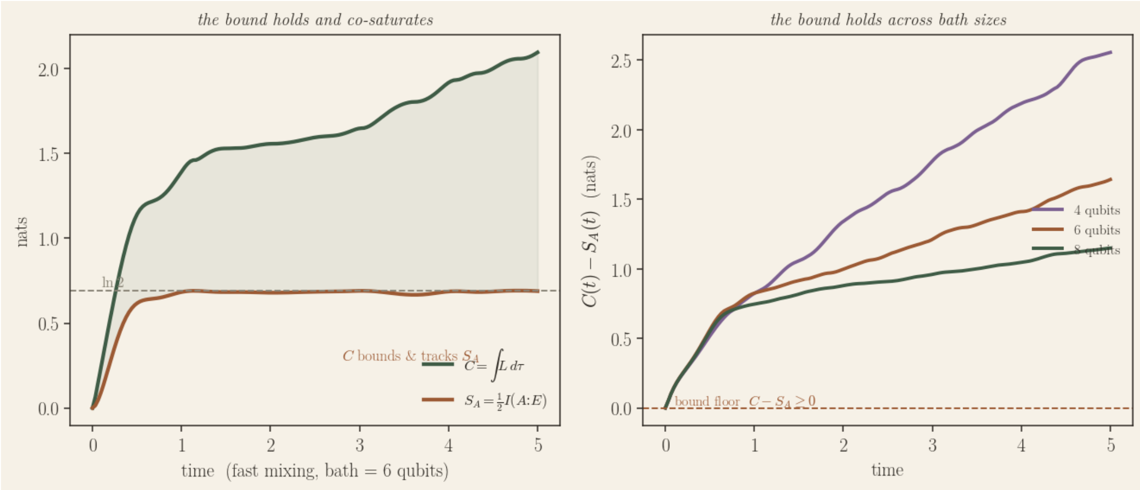

Figure 7: Retensing as a bound, not an identity. *Left*: in the fast-mixing regime the accumulated coherence C bounds and closely tracks the forbidden correlation $S_A = \frac{1}{2}I(A{:}E)$, the two co-saturating near $\ln 2$. *Right*: the margin $C(t) - S_A(t)$ remains non-negative across bath sizes, touching zero only at $t = 0$. The past always suffices to bound the present correlation.

the past, retensed from a spatial into a temporal object? The accumulated coherence C is exactly such an object. We test the relationship directly, and report a result more qualified than the strongest version of the conjecture but more interesting than its failure.

We construct explicit bipartite systems—a system qubit coupled to a chaotic bath—and at each time compute both the forbidden $S_A(t) = \frac{1}{2}I(A{:}E)$ and an accumulated $C(t) = \int_0^t L\,d\tau$ formed from the system's own trajectory. The mechanism that could relate them is real: de Bruijn's identity—that under diffusion the rate of entropy production equals a constant times the Fisher information—holds in our simulations to numerical precision. The findings: (i) the strong identity $C = I(A{:}E)$ is false under coherent dynamics—C grows monotonically while $I(A{:}E)$ oscillates back to zero, and even under fast mixing C keeps accumulating after S_A saturates; (ii) a bound holds robustly—across every bath size $(4, 6, 8$ qubits$)$ and realisation tested, 1920 sampled points, $C(t) \geq S_A(t)$ held pointwise, the two touching only at $t = 0$ and never reversing (Figure 7); (iii) in the fast-mixing regime where rupture lives, C tracks S_A strongly (correlation 0.997) and co-saturates near $\ln 2$; (iv) the fully legal version, computable from a fixed frame without self-tomography, is strongly frame-dependent and does not recover the invariant S_A.

The forbidden spatial quantity is therefore not literally equal to the accumulated temporal one; CRR does not dissolve the impossibility by computing the forbidden quantity through time from inside. What the result licenses is precise: the correlation one cannot measure across the boundary is always bounded above by an accumulated quantity formed from one's own past, and closely tracked by it in the regime where boundaries form. The resolution remains ontological—CRR denies that the co-present complement exists—supported by a robust bound rather than a numerical identity. We regard this as the honest form of the claim, the more so because the test was built to give the identity its best chance and the identity declined.

In a nutshell **The mathematics.** The strong identity $C = I(A{:}E)$ is false, but $C(t) \geq S_A(t)$ held pointwise at all 1920 sampled points, and C tracks S_A closely (correlation 0.997) exactly in the regime where boundaries form. The impossibility is not dissolved; the co-present far side it concerns is denied.

The karmic reading. One cannot survey one's entanglement with everything one is not; but one's accumulated history always suffices to bound it. This is the epistemology of Hershock's wise practitioner—one who sees interdependence and is skilled in karma and its fruits: relation is known along the lived past, never across a co-present wall.

8 Co-construction of the collective now

The boundary-as-rate reading entails that the boundary at the now is not any one system's possession but is co-constructed where many systems' saturations meet. Because the rupture–regeneration cycle of (2)–(3) is a relaxation oscillator, this becomes a question in nonlinear dynamics with a non-negotiable answer.

Two systems. In an explicit quantum model of two coupled subsystems, each coupled to a bath, coupling makes each system's boundary formation depend on the other's state—mutually, and growing with coupling strength (Figure 8, right; the co-determination rises from zero to ≈ 0.15 bits as the contact grows). The decisive question is whether that dependence entrains the rupture events. With $\Omega_P = 1/\pi$ (bistable) and $\Omega_Q = 1/2\pi$ (rotational), P ruptures twice as often as Q: the natural ratio is 2:1, the topological CRR ratio, forced by the doubling of Section 2.1 rather than assigned by comparing two manifolds; the two classes differ only in the arc accumulated between ruptures. Sweeping the natural ratio across 2:1 for several couplings and measuring the rotation number ρ, we find ρ follows the natural ratio with no plateau at zero coupling (merely co-present), while at non-zero coupling a plateau locked at $\rho = 2$ opens and widens—an Arnold tongue, inside which neither boundary is placeable independently (Figure 8, left). At exactly the topological ratio the systems lock for *any* non-zero coupling ($\rho = 2.0005$ throughout), the relative phase pinned: CRR sits at the centre of the tongue by topology, not by tuning, and the tongue widens from a point at $\kappa = 0$ to the full swept range by $\kappa = 0.25$.

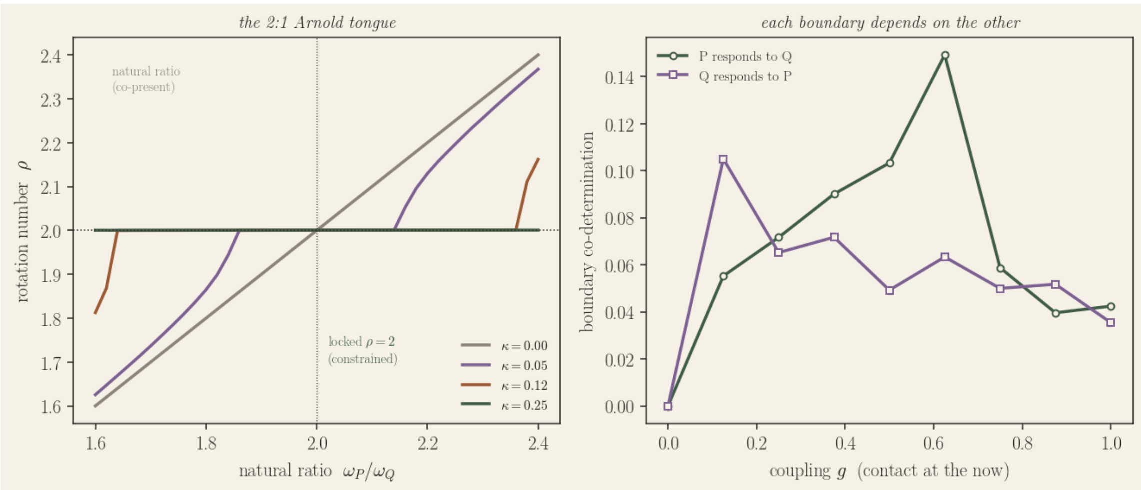

Figure 8: Two-body co-construction. *Left*: the rotation number ρ (ruptures of P per rupture of Q) versus the natural-frequency ratio. At $\kappa = 0$ it follows the natural ratio—merely co-present; at $\kappa > 0$ a plateau locked at $\rho = 2$ opens and widens—an Arnold tongue signalling mutual constraint; the dotted line marks the topological ratio. *Right*: in an explicit quantum model, each boundary depends increasingly on the other's state as contact grows, mutually.

Many systems. Whether the pairwise result extends to the convergence of all finite systems is the Kuramoto problem: N rupture oscillators with a spread of natural frequencies, coupled through their contact at the shared now, with order parameter $re^{i\psi} = N^{-1}\sum_j e^{i\theta_j}$. We hold the result to the standard required of a finding in this literature: validate against exact analytic results, then establish by finite-size scaling whether any transition survives the thermodynamic limit. For a Lorentzian frequency distribution of half-width γ the transition has exact critical coupling $K_c = 2\gamma$ and $r(K) = \sqrt{1 - K_c/K}$ above it; our simulator reproduces this to within 0.015, which validates it. Finite-size scaling then establishes a genuine transition (Figure 9, left): below K_c the order parameter falls as $N^{-0.47}$, consistent with the incoherent $N^{-1/2}$ (no shared now in the limit); above K_c it converges to a finite, N-independent value (a shared now); and the curves pivot ever more sharply at K_c as N grows. Above threshold the locked fraction rises from a few per cent to unity and is stable as N grows from hundreds to thousands (Figure 9, right): a finite, macroscopic fraction genuinely shares the collective now.

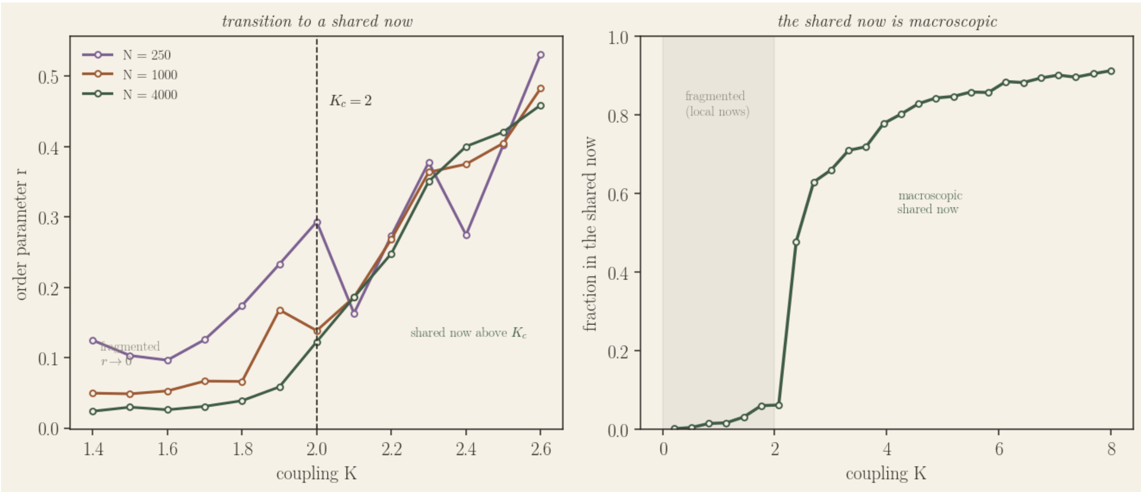

Figure 9: The many-body shared now. *Left*: finite-size scaling of the order parameter near the critical coupling. Below K_c the curves separate by N and fall toward zero (a fragmented manifold of local nows); above K_c they converge onto one N-independent curve (a genuine shared now), pivoting sharply at the exact $K_c = 2$. *Right*: the locked fraction jumps from near zero to macroscopic at threshold and is stable as N grows.

We add one correction against over-extrapolation. The pairwise 2:1 topological resonance governs weak-coupling cross-cluster locking; as coupling strengthens the clusters are pulled toward a common effective frequency and global 1:1 entrainment overrides the pairwise 2:1. The two-body and many-body results are related but distinct regimes, and we do not claim the 2:1 propagates to the collective. This pairwise structure echoes a finding of the companion study [5], where coordination between two channels was likewise a two-body phenomenon that scales by pairwise composition rather than by adding channels within a single agent.

In a nutshell **The mathematics.** Two coupled rupture cycles lock at the topological 2:1 for any non-zero coupling—an Arnold tongue centred by topology, not tuning. Many cycles undergo a genuine synchronisation transition, validated against the exact $K_c = 2\gamma$: below threshold only fragmented local nows survive the large-N limit; above it a macroscopic fraction shares one now. The pairwise 2:1 does not propagate to the collective.

The karmic reading. Collaborative karma. Hershock's karmic cohorts—ensembles coursing together through time in shared commitment—here acquire a dynamical criterion: below critical coupling there are only private nows; above it, a shared present whose responsibility is neither exclusively one's own nor another's, but held in common.

9 A pre-registered prediction: the breath

The breath, used in Section 2.1 to fix the two symmetry classes, also carries the framework's cleanest contact with measurable data. We state the predictions here and do *not* test them; the test belongs to empirical work kept separate from the consistency demonstrations above, so that the formal status of the present results is preserved. What follows is offered in the discipline of pre-registration: the framework fixes the numbers, and the only freedom—which rhythm is assigned which symmetry class, under which condition, with which estimator—is what must be registered *before* any data are touched.

The coefficient-of-variation prediction. The respiratory cycle is an autonomous limit-cycle oscillator whose phase runs a closed loop; as a system its question is rotational, so it takes $\Omega = 1/2\pi$ and ruptures once per cycle. The parameter-free prediction of Section 2 is then that the coefficient of variation of the inter-breath interval is $\mathrm{CV} = 1/4\pi \approx 0.080$. One caution is part of the registration, not a caveat added after the fact: this is a prediction about the *autonomous* rhythm. Awake spontaneous breathing carries additional cortical and behavioural variance—and a known observer confound, that subjects asked to breathe normally tend to hyperventilate—so the clean test is the most autonomous condition available, quiet NREM sleep or a CO_2-clamped protocol, where cortical override is minimal. That the test targets the autonomous condition follows from the theory, Ω being intrinsic and geometric, and is fixed before the data.

The entrainment prediction. Section 8 predicts that a bistable and a rotational rhythm lock at 2:1. Locomotor–respiratory coupling—two strides per breath is the commonest running mode—is the natural test bed. But 2:1 Arnold tongues are generic to coupled oscillators, so observing 2:1 is not evidence for CRR. The distinctive claim is stronger: at the topological ratio the tongue is anomalously robust—locking for any non-zero coupling, centred by topology not tuning (Figure 8)—so the falsifiable prediction is a tongue-*width* one: sweeping natural detuning against a metronome, the 2:1 plateau should be wider and more robust than the neighbouring rational tongues. That is the test that separates the topological claim from ordinary mode-locking.

10 Discussion

The two frameworks converge on the same negation and diverge on what remains after it. The original programme's emptiness is epistemic with ontological force: the separability of agent from environment is structurally unevidenceable, a structural prior that constrains every measurement frame and is never itself measured; the boundary is demoted from posterior to permanent prior, but a determinate far side is retained. CRR's emptiness is ontological: there is no partition to be unable-to-measure-across, because the boundary was never a surface, only the rate $C \cdot \Omega \to 1$; one cannot get behind it because there is no behind, the edge being constitutively where one accumulates from, and the present it bounds having no content of its own. The far side is not hidden but reconceived as the kernel-weighted past of every finite thing, met at the contentless now. The original emptiness says one cannot evidence the wall; CRR's says there is no wall, only earlier, taken up selectively. Both honestly earn *no separate self*; only CRR also relinquishes the screen's far side, trading an epistemic limit for a process ontology.

The emptiness CRR arrives at is, in consequence, not privative. A present with no content of its own is not a vacuum but a capacity: the one locus at which the settled past of every finite system can arrive without any past crowding out another, each taken up through its own kernel, none privileged by the meeting-point itself. Emptiness, on this reading, is what equality of standing looks like from inside time—the now holds nothing so that it can hold everything. That the upper limit of integration carries no content is not a poverty but the condition of the whole: the same fact that empties the present is the fact that lets every past into it.

The deepest agreement is methodological. The original conclusion follows from its axioms—the boundary is unevidenceable is true by construction in the QRF formalism, as no-cloning is—and CRR's conclusion is equally axiom-bound: only the past has content is true by construction once content is defined as accumulated C. Neither sentence is a discovery about the world; each is what its formalism makes unavoidable, which is the honest status of foundational claims in both programmes. The useful question is what each costs to adopt. The QRF picture buys a rigorous impossibility theorem and a clean bridge to contemplative phenomenology, and keeps a determinate, screened-off world one can never touch; CRR relinquishes the far side for a process whose present is empty and whose past is selectively enacted, and takes on a debt—the regeneration kernel is not unitary and Ω must eventually be paid for in free energy. One currency is an unmeasurable correlation, the other a thermodynamic budget; both are payable, neither is free.

There is, finally, a place where CRR may complete rather than diverge. It is noted in the quantum-FEP that the time parameter is not observable clock-time, and that an observer-relative measurable time must be introduced separately and remains underdeveloped. CRR's rupture condition is a clock: $C \cdot \Omega = 1$ defines the next tick as the next saturation, and the same kernel that fails the unitarity requirement is the one that supplies the temporal operator the formalism lacks. If that holds, the relation between the frameworks is not rivalry but a trade—CRR pays an information-conservation debt and is repaid in a definition of *when*—and it is the seam the original paper, by underspecifying time, leaves open for exactly this completion.

Limitations. We take the impossibility theorem to be correct and important, and the honest cost of saying so is that CRR inherits it: at every rupture an effective screen crystallises, and across it the same impossibility holds. Every result reported above is internal consistency under a chosen formalism, exactly as the original theorem is, and none is contact with empirical data. The co-construction results model the rupture cycle as a phase-reduced relaxation oscillator, faithful to (2)–(3) in the weak-coupling limit but a reduction; the many-body critical coupling depends on the distribution of symmetry classes, which CRR does not specify, so what is robust is the existence and finite-size reality of the transition, not the value of K_c; the now is operationalised as phase coherence of rupture timing, a defensible but timing-level operationalisation; the retensing is a bound, not an identity; and the self-measurement thread is the reading the framework's commitments select, not a theorem proved in the full formalism. The breath predictions of Section 9 are stated, not tested. The falsifiable consequences we would press next are the topological sector-count prediction, the breath coefficient of variation, and the explicit free-energy budget for Ω, which may or may not leave Ω on its geodesic values once thermodynamic cost is imposed.

11 Conclusion

Sandved-Smith and colleagues prove that a finite, self-evidencing system cannot evidence its own separateness, and propose that awakening is the embodied recognition of that impossibility. Reading the result through Coherence–Rupture–Regeneration leaves the theorem intact and changes the ontology around it. The boundary is not a screen with a hidden far side but a rate—the moving front at which $C \cdot \Omega \to 1$—and it exists only at the momentary coincidence of inferential map and inferred territory, re-enacted at each now and never persisting. The present, being that coincidence, has no content of its own; content is the past, prehended selectively. Emptiness moves from an epistemic limit on what can be measured to an ontological claim about what there is: not a self behind an unmeasurable wall, but no wall—only earlier, taken up by each finite system through its own kernel. The single point at which CRR fails the formalism becomes the mark of the measurement that constitutes each now, carrying the thermodynamic cost the framework independently owes; and the convergence of many finite systems, far from a metaphor, is a genuine synchronisation transition in which a macroscopic, hierarchically structured collective now emerges above a critical coupling and fragments below it. Whether Ω survives its thermodynamic cost on its geodesic values, and whether the self-measurement thread holds in the full formalism, are the questions on which the account stands or falls—and they are, at least, the right kind of questions, the kind the world can answer.

A last consequence is worth stating, because it follows rather than inspires. A theory of everything, for a finite system, could never be a view from outside the integral (1); there is no outside to view from. What a finite system can have is its kernel—what it attends to, what it reconstructs, how readily it lets its coherence break—and through that kernel the whole of the settled past is taken up. In that strict sense each finite system already carries the only theory of everything available to anything: its own manner of prehending what has been. The equations above are one such taking-up, and are offered as nothing more.

Code and data availability. The numerical procedures underlying Figures 2–9 are available as a single documented, deterministic Python script at `https://www.space-zero.org/blogs/boundary.py`. Following Friedman's operational-surrogate and cognitive-security discipline [2], the script binds each figure to a source anchor and an explicit *evidence ceiling*—what must not be inferred from it—and re-checks every headline number against a declared tolerance band through built-in PASS/FAIL validation gates, recording figure-integrity hashes so no figure can drift from the number it reports. All results are consistency demonstrations under a chosen formalism and are not empirical findings.

Use of AI assistance. Claude Opus 4.8 and Claude Fable 5 were used for proof-reading, composition, and computational verification. All claims, derivations, and interpretations remain the responsibility of the author.

Acknowledgements. The author thanks the Active Inference Institute community and the interlocutors of the CRR framework, and acknowledges the Sandved-Smith et al. preprint as the direct occasion for this work; and a season in Crescent City, where the distinction between the seminar on emptiness and the market stall was, instructively, hard to maintain.

References

[1] L. Sandved-Smith, C. Fields, T. Doctor, R. E. Laukkonen, and J. Hohwy, *There is no self-evidence: A physics of emptiness realisation*, PsyArXiv, preprint m78z2, 2026.

[2] D. A. Friedman, *Realizing Emptiness: Operational Surrogates for No-Self-Evidence, QRF Opacification, and Bayesian Model Reduction* (v1.0.0), Active Inference Institute, Zenodo, 2026, doi:10.5281/zenodo.20834847.

[3] C. Fields, K. Friston, J. F. Glazebrook, and M. Levin, A free energy principle for generic quantum systems, *Progress in Biophysics and Molecular Biology*, vol. 173, pp. 36–59, 2022.

[4] C. Fields and J. F. Glazebrook, Separability, contextuality, and the quantum frame problem, *International Journal of Theoretical Physics*, vol. 62, art. 159, 2023; and C. Fields and A. Marcianò, Holographic screens are classical information channels, *Quantum Reports*, vol. 2, no. 3, pp. 326–336, 2020.

[5] A. Sabine, Phase-gating across precision channels: topological constraints on multi-channel belief update dynamics, in *Active Inference (AGI 2026)*, Springer LNCS/LNAI, 2026.

[6] J. Hohwy, The self-evidencing brain, *Noûs*, vol. 50, no. 2, pp. 259–285, 2016.

[7] K. Friston, A free energy principle for a particular physics, arXiv:1906.10184, 2019.

[8] L. Da Costa, K. Friston, C. Heins, and G. A. Pavliotis, Bayesian mechanics for stationary processes, *Proceedings of the Royal Society A*, vol. 477, no. 2256, art. 20210518, 2021.

[9] S. Ito and A. Dechant, Stochastic time evolution, information geometry, and the Cramér–Rao bound, *Physical Review X*, vol. 10, no. 2, art. 021056, 2020.

[10] N. N. Čencov, *Statistical Decision Rules and Optimal Inference*, Translations of Mathematical Monographs, vol. 53, AMS, 1982.

[11] S. Amari and H. Nagaoka, *Methods of Information Geometry*, AMS, 2000.

[12] E. T. Jaynes, Information theory and statistical mechanics, *Physical Review*, vol. 106, pp. 620–630, 1957.

[13] M. Levin, Technological approach to mind everywhere, *Frontiers in Systems Neuroscience*, vol. 16, art. 768201, 2022.

[14] A. N. Whitehead, *Process and Reality: An Essay in Cosmology*, Macmillan, 1929.

[15] A. N. Whitehead, *The Concept of Nature*, Cambridge University Press, 1920.

[16] F. J. Varela, E. Thompson, and E. Rosch, *The Embodied Mind*, MIT Press, 1991.

[17] H. Cramér, *Mathematical Methods of Statistics*, Princeton University Press, 1946; and C. R. Rao, *Bull. Calcutta Math. Soc.*, vol. 37, pp. 81–91, 1945.

[18] J. von Neumann, *Mathematische Grundlagen der Quantenmechanik*, Springer, 1932.

[19] R. Landauer, Irreversibility and heat generation in the computing process, *IBM Journal of Research and Development*, vol. 5, no. 3, pp. 183–191, 1961.

[20] Y. Kuramoto, *Chemical Oscillations, Waves, and Turbulence*, Springer, 1984; and S. H. Strogatz, *Physica D*, vol. 143, pp. 1–20, 2000.

[21] J. L. Garfield, *Engaging Buddhism: Why It Matters to Philosophy*, Oxford University Press, 2015; and M. Siderits and S. Katsura, *Nagarjuna's Middle Way*, Wisdom Publications, 2013.

[22] P. D. Hershock, Valuing karma: a critical concept for orienting interdependence toward personal and public good, in C. Prebish, D. Keown, and D. S. Wright (eds.), *Revisioning Karma*, Journal of Buddhist Ethics Online Books, 2007.

[23] P. D. Hershock, *Consciousness Mattering: A Buddhist Synthesis*, Bloomsbury, 2023; and *Buddhism and Intelligent Technology*, Bloomsbury, 2021.

[24] T. Doctor, O. Witkowski, E. Solomonova, B. Duane, and M. Levin, Biology, Buddhism, and AI: Care as the driver of intelligence, *Entropy*, vol. 24, no. 5, art. 710, 2022.