Phase-Gating Across Precision Channels: Topological Constraints on Multi-Channel Belief Update Dynamics

Alexander Sabine

Independent Researcher
Alexander@activeinference.institute
https://www.temporalgrammar.ai/human_networks.html

Abstract. Every finite system that persists must accumulate evidence, reach the limits of its current regime, and reconstruct from its own history. We propose that these three temporal operations, which we call Coherence, Rupture, Regeneration (CRR), constitute a candidate temporal grammar for adaptive systems. The dynamics are governed by a single geometric parameter Ω, fixed by the topology of the system's statistical manifold. Two fundamental symmetry classes yield two thresholds without free parameters: π for bistable ($\mathbb{Z}_2$) systems, 2π for rotational (SO(2)) systems, with a ratio of exactly 2. Ω acts as a temporal-cognitive light cone [1], governing the depth of accessible history and the breadth of future possibility.

We apply CRR to a concrete problem in multi-channel AGI architecture - how do multiple precision channels coordinate their belief updates in Active Inference? Assigning $\mathbb{Z}_2$ symmetry to sensory precision and SO(2) to prior precision, we test the resulting dynamics in a POMDP with Dirichlet learning. The primary finding is phase-gating: the topological thresholds produce a strongly non-uniform phase relationship between channels ($\chi^2 = 8{,}041$) that determines whether each update drives learning or action. This result is consistent across environment topologies, independent of the weight function, and structurally compatible with recent empirical findings on neuromodulatory timing. CRR offers a parsimonious temporal constraint for multi-channel AGI architectures, with potential applications for AI alignment, continual learning, and applied Active Inference.

Keywords: Active Inference · Temporal grammar · Precision weighting · Information geometry · Phase-gating · AGI architecture

1 Introduction

The Free Energy Principle (FEP) describes how adaptive systems maintain their statistical boundaries by minimising variational free energy [2]. In Active Inference, this account extends to action and perception through a hierarchy of

precision-weighted prediction errors [3]. Distinct precision channels govern sensory confidence (likelihood precision), model confidence (prior/transition precision), and policy confidence. The textbook treatment specifies *what* each channel does but leaves open the temporal structure of *when* each channel updates, *how much* information it absorbs per update, and how the channels *coordinate* with one another. This temporal structure is underscored in recent empirical work by Jang et al. [4] who demonstrate that the relative timing of neuromodulatory precision signals, not their magnitude, determines whether the same dopamine signal promotes reinforcement learning or movement vigour.

This paper introduces the Coherence-Rupture-Regeneration (CRR) framework as a temporal grammar for self-organising systems. CRR models the update dynamics of bounded systems from information geometry, connecting to the Cramér-Rao bound [5,6] and the thermodynamic speed limits established by Ito and Dechant [7]. The single parameter Ω functions as an information-theoretic capacity limit: the geodesic extent of the system's statistical manifold [8,9] determines Ω uniquely, yielding $\Omega = 1/\pi$ for systems with $\mathbb{Z}_2$ (bistable) symmetry and $\Omega = 1/2\pi$ for systems with SO(2) (rotational) symmetry. Within the agent, the two symmetry classes run concurrently; sharing an evidence stream but accumulating on manifolds with different geodesic extents. This produces a non-uniform distribution over the phase at which one process ruptures relative to the other. This is what we mean by *phase-gating* in an information-theoretic context. Without a topological account, this structure would appear as stochastic noise in update timing; CRR provides a geometric basis for its functional role in free energy minimisation.

The observation that motivated the present study came from an interactive network simulation of agents connected by Markov blankets.[1] The symmetry assignment emerged from the graph structure itself, where each node represents an internal model that traverses a continuous cycle of belief updating (SO(2) dynamics), while each edge is a statistical boundary that alternates between two regimes of influence ($\mathbb{Z}_2$ dynamics). In other words, the two symmetry classes were not awkwardly imposed on the FEP's precision architecture; they were found in its graph topology.

If two precision channels have different symmetry classes, what does the topology enforce about their coupled dynamics? We find that the topology enforces a conservation law: both channels must process equal total precision-gain per unit time. We go on to show in Section 4 that this conservation is a structural consequence of equal evidence delivery, rather than a distinctive prediction of CRR. What the topology contributes is not the conservation principle itself, but the specific temporal decomposition and, crucially, the phase-gating dynamics that emerge from the $\pi/2\pi$ partition. We also found that adding a third channel does not automatically extend pairwise conservation, establishing a clear boundary, while highlighting wider implications for distributed versus centralised AGI systems. Indeed, these findings suggest that the temporal coordination of preci-

[1] An interactive version of this simulation is available at https://www.temporalgrammar.ai/human_networks.html.

sion channels in AGI architectures should respect the topological constraints of their respective state spaces. For explainability, we now ask the reader to play with the pedagogical CRR Active Inference Network simulation and reflect on their experience. The parsimony of the framework - a single geometric parameter governing temporal coordination - is perhaps best appreciated through direct engagement with the dynamics it produces.

2 Related Work

The FEP furnishes adaptive systems with a statistical manifold [10]: internal states parametrise beliefs, flows on this manifold constitute inference, and the Fisher information metric supplies the geometry in which distances between beliefs are measured [8,9]. Kirchhoff et al. [11] extended this picture to biological autonomy, and the path integral formulation [12] showed that trajectories through belief space are weighted by accumulated action, much as in statistical physics. The FEP already provides temporal depth through hierarchical nesting: slower timescales emerge from deeper levels of the generative model, each level Markovian given the level below [3], giving the system a representational architecture for the past and future. What this architecture leaves open is the process dynamics: *when* must a given level reorganise, and how do precision channels at different levels coordinate their updates in time? CRR complements the representational account with a process-level description. The coherence integral (Eq. 1) accumulates arc length along the system's information-geometric trajectory, making the present state explicitly dependent on the regime's entire history: non-Markovian at the process level, even within a Markovian generative model. Rather than adding temporal depth through additional hierarchical levels, CRR reads the temporal structure from a single geometric quantity, Ω, which determines how much past the current regime can access and how far into the future it can project. In this respect, Ω plays a role analogous to the cognitive light cone introduced by Levin [1]: both govern the scale over which an agent integrates information and directs action. Where the cognitive light cone provides a measure of that scale, CRR derives it from the geodesic extent of the manifold itself (Fig. 2). Tolchinsky, Levin et al. [13] have independently shown that the collapse of temporal depth, treated as a parameter of the generative model, characterises dissociative states; CRR offers a complementary account in which the collapse corresponds to a specific, measurable event: the rupture condition $C \cdot \Omega = 1$. The parallel with the FEP's spatial structure is direct. The Markov blanket separates inside from outside in space; the rupture condition delineates past from future in time. The two accounts are not alternatives; they are spatial and temporal aspects of the same partition.

The geometric tools for such a temporal account already exist in the literature, though they have not previously been applied to this question. A system traversing a statistical manifold is subject to fundamental speed limits: the Fisher information bounds the rate of state change [7], and thermodynamic fluctuation timescales are universally constrained by the same geometry [14,15]. Parr, Da

Costa, and Friston [16] showed that Markov-blanketed systems possess both an information geometry and a stochastic thermodynamics, establishing that these speed limits apply to any system with a particular partition. These results specify how fast evidence can accumulate, but not what happens when accumulation is complete. The manifold is finite. Its geodesic extent is a topological invariant: π for the Bernoulli manifold, 2π for the circle. A system accumulating evidence along this geodesic must eventually span it. CRR proposes that this moment, when accumulated arc length equals the manifold's full extent, is rupture: a transition condition at Cramér-Rao saturation, after which the system must reorganise from its own history. The speed limits tell us how fast the system can travel; the geodesic extent tells us how far it has to go; the ratio of the two determines *when* it arrives.

This temporal structure bears directly on a concrete empirical puzzle. Precision weighting is central to Active Inference [3], where distinct neuromodulatory systems, principally dopamine and acetylcholine, set the gain on prediction error channels [17]. But these two signals do not operate independently. Krok et al. [18] showed that dopamine and acetylcholine fluctuate spontaneously at approximately 2 Hz in the striatum, maintaining a consistent phase offset even in the absence of any external reward or salient stimulus: two precision clocks, ticking together but not in unison. Taniguchi and Tritsch [19] found that the precise relationship between cholinergic pauses and dopamine signals varies by striatal subregion, ruling out a simple uniform gating model and suggesting that the phase structure carries functional information. Most directly, Jang et al. [4] demonstrated that the relative timing of acetylcholine and dopamine release, not their magnitude, determines whether dopamine promotes reinforcement learning or movement vigour. When a dopamine burst slightly lagged a cholinergic dip, the animal learned; when it slightly preceded the dip, it moved. The functional difference reduced to a window of tens of milliseconds in the phase relationship between the two signals. The hypothesis that cholinergic and dopaminergic signals must be temporally coincident for synaptic plasticity [20,21,22] further motivates the search for a principled account of what determines this phase structure. CRR offers such an account. If sensory precision and prior precision accumulate on manifolds with different geodesic extents, they will rupture at different rates, and the phase at which one ruptures relative to the other will follow a non-uniform distribution determined by topology rather than by a dedicated controller. The empirical question, why does timing matter more than magnitude, becomes a geometric one: two clocks with a 2:1 period ratio, sharing the same evidence stream, must produce a structured phase relationship. The structure is not noise; it is the topology, expressing itself in the dynamics. This perspective is consonant with work on inter-agent coordination under Active Inference [23], information-geometric constraints on development [24], and the duality between the FEP and the maximum-entropy principle [25].

3 The CRR Framework

CRR models bounded systems through three coupled temporal processes (Fig. 1), resting on five core commitments that any persistent, finite, adaptive system must satisfy. We present these commitments because they determine the temporal decomposition and phase-gating dynamics that underpin the paper's main contribution of the CRR canonical formalism itself.

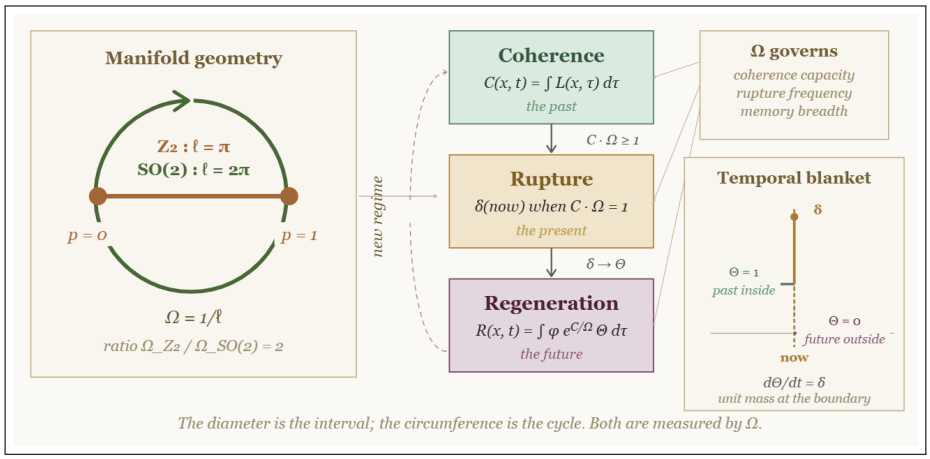

Fig. 1. The CRR framework and its geometric grounding. Coherence accumulates evidence from the past. Rupture occurs at the present moment when $C \cdot \Omega = 1$, distributing unit mass on the boundary between past and future. Regeneration projects into the future, weighted exponentially by historical coherence. The single parameter Ω governs all three processes: coherence capacity, rupture frequency, and the breadth of the memory kernel available for reconstruction. The diameter is the interval; the circumference is the cycle. Both are measured by Ω.

Five Commitments

Commitment 1: Persistence implies accumulation. A system that persists must accumulate patterns that work, reducing surprise (Variational Free Energy) in their eco-niche. We formalise this as a coherence integral:

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

where $L(x,\tau) \geq 0$ is the rate of coherence accumulation, identified with the Fisher-Rao speed on the statistical manifold [8,26]. Under the path integral formulation of the FEP [12], this is the accumulated arc length along the system's information-geometric trajectory. The integral over the system's past makes

CRR distinctively non-Markovian by design - the present state depends on the accumulated history of the current regime up to the present moment of constraint. In other words, the coherence integral represents the totality of the system's past regime, constituting the interior of the temporal Markov blanket.

Commitment 2: Finite capacity implies rupture. No finite system can accumulate coherence without bound. At some point the current regime is exhausted and the system must reorganise:

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

The Dirac delta has zero temporal extension and carries exactly one unit of mass on the boundary between the accumulated past and the regenerated future. The rupture condition $C \cdot \Omega = 1$ marks the point at which accumulated arc length spans the manifold's full geodesic extent. The system has traversed all distinguishable states available to its current regime and must reorganise. This is structurally analogous to Cramér-Rao saturation [7,14], in that both concern the relationship between accumulated Fisher information and the geometry of the estimation problem; we exploit this analogy but note that the formal identity between geodesic traversal and estimator efficiency requires the additional assumption that the speed limit is saturated throughout the trajectory, a condition we do not verify here. In discrete systems, rupture is a threshold-crossing rather than exact saturation: a heartbeat, a breath turning, a neuron firing. At rupture, $C \cdot \Omega = 1$: inside meets outside. The system's accumulated coherence exactly matches its capacity, and what was contained within the temporal blanket makes contact with what lies beyond it.

Commitment 3: Reconstruction is weighted by historical coherence. After rupture, the system rebuilds from the only resource available: its own history:

$$R(x,t) = \int_{-\infty}^{\infty} \varphi(x,\tau) \cdot \exp\left(\frac{C(x,\tau)}{\Omega}\right) \cdot \Theta(t-\tau)\, d\tau \tag{3}$$

where $\varphi(x,\tau)$ is the reconstruction resource (the field of past states available for reconstitution). The Heaviside step function $\Theta(t-\tau)$ enforces causal ordering: only the past contributes to reconstruction. We retain Θ explicitly because the Dirac delta of rupture (Eq. 2) and the Heaviside are conjugate distributions ($d\Theta/dt = \delta$), ensuring that the regeneration integral is the Green's function of the rupture operator and connecting the three CRR equations into a single dynamical system. The exponential kernel $\exp(C/\Omega)$ is the maximum-entropy weighting consistent with a constraint on the expected value of C [27]- given only the accumulated coherence and the capacity parameter, the distribution that maximises entropy while respecting the constraint on $\langle C \rangle$ takes the Boltzmann form, with C playing the role of energy and Ω the role of temperature.

Commitment 4: The system has a characteristic variance. The parameter Ω is fixed by the topology of the state space, not fitted. It simultaneously governs coherence capacity (how much can accumulate before rupture), rupture frequency

(how often the system must reorganise), and memory breadth (which past moments dominate reconstruction). We argue that these are not three independent properties; but a single geometric fact about the system's attractor topology. When Ω is small, the memory kernel peaks sharply and the system reconstitutes the same patterns reliably (habit, rigidity); when Ω is large, the kernel is flat and the system can reconstruct broadly (flexibility, exploration).

Commitment 5: Rupture is a maximum-entropy event. At rupture, the system occupies the maximum-entropy state consistent with its symmetry class [27]. A bistable system at rupture is maximally uncertain about which state comes next; a rotational system at rupture has no preferred phase. This is the axiom that converts geometric structure into specific numerical predictions.

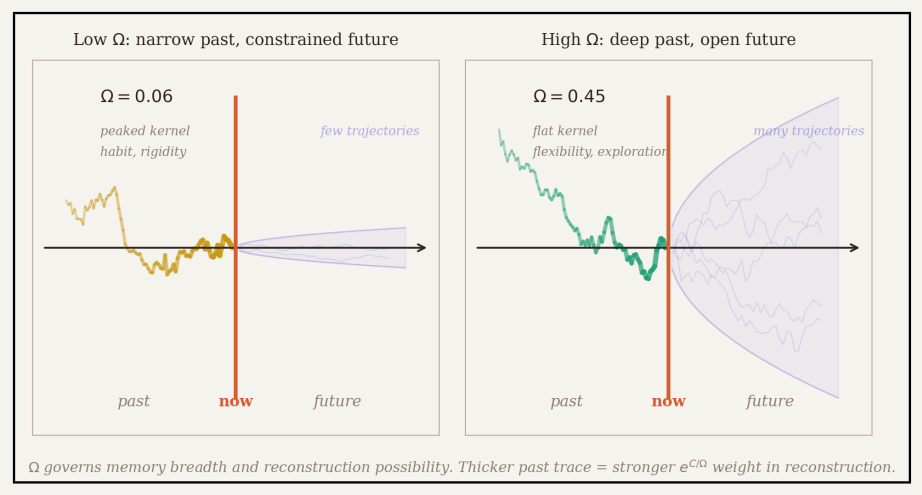

Fig. 2. Ω as a temporal light cone. Left: low Ω produces a sharply peaked memory kernel; only recent moments contribute to reconstruction, and the future possibility cone is narrow. Right: high Ω produces a flat kernel; deep history is accessible, and the future is open.

Fixing Ω from Geometry The five commitments leave one degree of freedom - the value of Ω - which must be fixed by the system's structure alone in order to generate falsifiable predictions. The key insight is that for systems whose internal states parametrise probability distributions (which, under the FEP, includes any system possessing a Markov blanket [10]), the arena of coherence accumulation is a statistical manifold with a specific, measurable geometry. On such a manifold, Čencov's uniqueness theorem [9] constrains the metric to be the Fisher information metric [8], and the geodesic structure of that metric fixes

the maximum arc length the system can traverse. The capacity Ω becomes a geometric invariant: $\Omega = 1/\ell$, where ℓ is the geodesic extent of the manifold.

Two fundamental manifold topologies fix ℓ (Fig. 1):

$\mathbb{Z}_2$ *(bistable).* The Bernoulli manifold, parametrised by $p \in [0, 1]$, has Fisher-Rao geodesic diameter π [8]. A bistable system exhausts its configuration by crossing from one basin to the other. Crucially, the Bernoulli manifold is an interval, not a loop, meaning that the system cannot return to its starting state without retracing the same geodesic. The maximal non-self-intersecting geodesic therefore has length equal to the diameter: $\ell_{\mathbb{Z}_2} = \pi$, hence $\Omega_{\mathbb{Z}_2} = 1/\pi \approx 0.318$ and $C^* = \pi$.

SO(2) (rotational). The circular manifold S^1 has circumference 2π. A rotational system exhausts its configuration by completing one full cycle, returning to its starting point via a path that never revisits any intermediate state. The maximal non-self-intersecting geodesic is therefore the full circumference: $\ell_{SO(2)} = 2\pi$, hence $\Omega_{SO(2)} = 1/2\pi \approx 0.159$ and $C^* = 2\pi$. The ratio between classes is, trivially, exactly 2, a topological invariant independent of any physical parameters.

FEP Correspondence and Precision Architecture The structural mapping between CRR and the FEP is that precision corresponds to $(1/\Omega)\cdot\exp(C/\Omega)$, where $1/\Omega$ is the topological contribution (fixed by manifold geometry) and $\exp(C/\Omega)$ is the dynamical contribution (growing with accumulated evidence); variational free energy corresponds to coherence C (inversely); the FEP Markov blanket corresponds to the high-∇C surface, and is temporally delineated by the Dirac Delta (δ) between past states (C) and future states (R), representing a temporal informational boundary where all past states are inside the (temporal) blanket, future states outside the (temporal) blanket.

As noted in Section 1, the assignment of symmetry classes to precision channels was originally observed in the CRR network simulation before being formalised in POMDP. In any graph of Markov-blanketed agents, an edge is a boundary that alternates between two regimes of influence ($\mathbb{Z}_2$ dynamics), while a node is an internal model that traverses a continuous cycle of belief updating (SO(2) dynamics). The assignment $\mathbb{Z}_2 \to$ sensory precision and SO(2) $\to$ prior precision therefore follows from the graph topology of the FEP itself, where edges are boundaries and nodes are models. In the accompanying simulation, this is directly visible. Nodes (teal) cycle continuously through belief states with a slower, regular rhythm, while edges (purple) flash rapidly as they flip between regimes of influence between nodes. We noticed that the two temporal characters were qualitatively different before any measurement was taken. In other words, the symmetry classes were not assigned to precision channels by analogy alone; they are inherited from the roles those channels play in the network.

This structural assignment is independently supported by the permeability properties of Ω. The regeneration kernel $\exp(C/\Omega)$ entails that high Ω produces a flat memory kernel (broad, flexible, permeable), while low Ω produces a sharp

kernel (narrow, rigid, precise). Indeed, under the FEP, the Markov blanket must be receptive to sensory states outside the blanket, while the internal prior model must maintain stability; since $\Omega_{\mathbb{Z}_2} = 1/\pi > \Omega_{SO(2)} = 1/2\pi$, the graph-structural and functional arguments converge on the same assignment.

Why only these two symmetry classes, not more? For the purposes of the empirical findings in this paper, the answer is dimensional. For a single-parameter system, the two irreducible topologies are the interval (a path between two endpoints, with the Bernoulli manifold as the canonical instance) and the circle (a closed loop with no endpoints). These exhaust the connected one-dimensional manifolds without boundary and with boundary respectively. Higher-dimensional manifolds (SO(3), simplices) could apply to multi-parameter systems; but the single-parameter case is the minimal one, and the one tested here.

Furthermore, this assignment is more immediately empirically falsifiable. Reversing it (SO(2) sensory, $\mathbb{Z}_2$ prior) produces measurably different dynamics, with sensory ruptures becoming rare (modal 1 per prior cycle, versus 3 under the standard assignment), and the prior channel becomes irregular (burstiness $B = -0.16$, versus -0.33). The reversed assignment predicts that sensory updates are *rarer* than prior updates, contradicting the FEP's requirement that the blanket be responsive to evidence and leading to counter-intuitive social dynamics in the pedagogical network simulation.

4 Experiments

Setup We implemented a POMDP agent in a ring Hidden Markov Model where both the likelihood (A-matrix, sensory precision) and transition (B-matrix, prior precision) are parametrised with Dirichlet concentration parameters that update through experience. The environment consists of n states arranged on a ring, with primarily clockwise transitions and Gaussian observation noise. The CRR engine determines *when* each channel updates: the sensory channel accumulates prediction errors (computed as surprisal at the true states under the agent's current model) until C_s reaches the $\mathbb{Z}_2$ threshold (π); the prior channel accumulates until C_p reaches the SO(2) threshold (2π). On rupture, accumulated evidence is applied to the Dirichlet parameters with a clamped exponential weight.

Several implementation details proved critical for replication and deserve explicit statement. First, the gain function uses a clamped exponential: $w = \min(\exp(C/\Omega), 10)/10\times 2+0.5$, mapping to [0.5, 2.5]; without clamping, $\exp(C^*/\Omega) = \exp(1/\Omega^2)$ produces catastrophic parameter inflation. Section 5.2 confirms that conservation is independent of the weight function. Second, "gain" is measured as $\Delta\bar{\alpha}$ (change in mean Dirichlet concentration), not a raw parameter sum. We note that this quantity is a precision-gain measure, not Shannon or Fisher information in the strict sense; throughout this paper, "information rate" should be read as "precision-gain rate". Third, both channels receive exactly one count of evidence per trial (one observation, one transition), ensuring equal evidence delivery. Fourth, prediction errors are computed from the agent's current model evaluated at the true states, not at beliefs.

Precision-Gain Conservation Across 10 conditions (5 state counts from 3 to 12, 5 noise levels from 0.05 to 0.40; 50 agents per condition, 800 trials each, bootstrap 95% CIs), we measured three quantities per channel: per-rupture precision gain, rupture frequency, and their product (total precision-gain rate).

Result. At equal evidence rate, the total precision-gain rate ratio (sensory/prior) converges to **1.003** with a 95% CI of [**1.000**, **1.005**] and CV = **0.024** (Fig. 3a,b). The gain ratio and frequency ratio both vary substantially across conditions (gain: 0.22 to 1.34; frequency: 0.73 to 4.47), but their product is locked to unity. When evidence rates are made unequal by scaling the sensory accumulation rate, the precision-gain ratio scales linearly with the input asymmetry (Fig. 3c), with $R^2 > 0.999$. The linear scaling confirms that conservation holds *per unit of evidence delivered*, not in absolute terms.

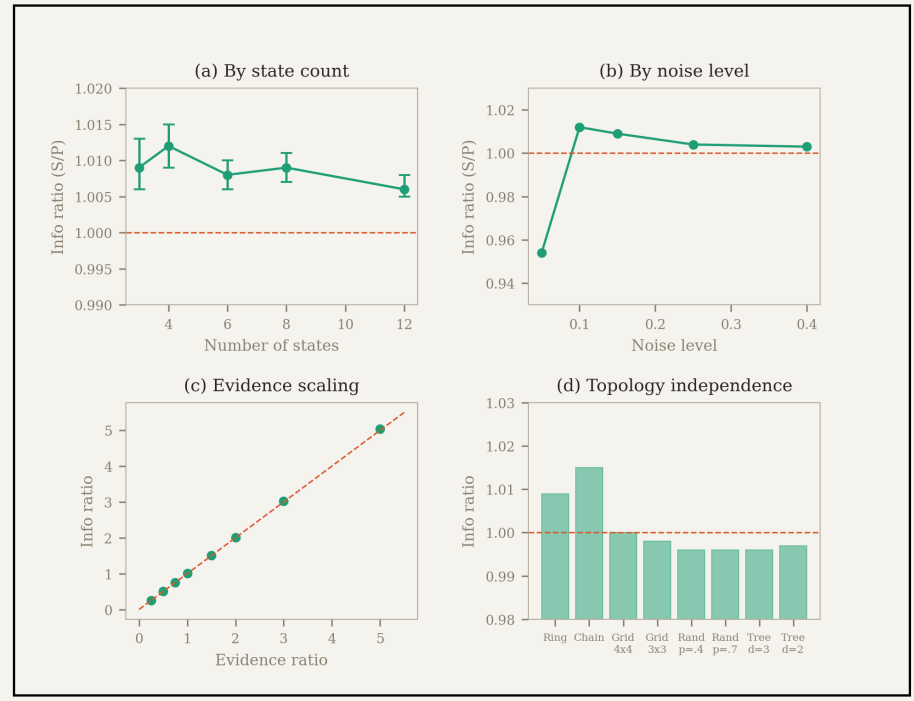

Fig. 3. Conservation results. (a) Information ratio by state count (50 agents each, bootstrap 95% CIs). (b) Information ratio by noise level. (c) Evidence scaling: the information ratio tracks the evidence asymmetry linearly ($R^2 > 0.999$). (d) Environment-structure robustness: conservation holds across ring, chain, grid, random graph, and tree environments. The dashed line marks conservation at 1.0.

Interpretation. The $\mathrm{SO}(2)/\mathbb{Z}_2$ topology enforces that both precision channels process equal total precision-gain per unit time. The $\mathbb{Z}_2$ channel makes many small updates (fast, frequent); the SO(2) channel makes few large updates (slow, deep). The factor of 2 sets the exchange rate between frequency and per-observation precision, not the total precision-gain rate: the topology determines the *grain* of each channel's observations, while the evidence accumulation rate determines their *bandwidth*. This is what two clocks with a 2:1 speed ratio produce when they share the same evidence stream: the fast clock ticks more often but learns less per tick, and the totals balance exactly. In the network simulation, this manifests as a visible gain-frequency trade-off: edges flicker rapidly with small coherence kicks, while nodes glow slowly toward a large belief update, and neither channel runs away from the other.

On the status of conservation. We tested whether the conservation law is a distinctive prediction of CRR or follows from the structure of the experiment. Running the same POMDP with continuous updating (no thresholds, updating every trial) produces a precision-gain ratio of exactly 1.0000. Testing CRR with arbitrary threshold ratios (1:1, 2:1, 3:1, random, reversed) also produces ratios within 1% of unity. Total precision-gain equality is therefore a structural consequence of equal evidence delivery, not a prediction that discriminates CRR from alternatives.

However, this structural guarantee is itself informative. It means that CRR architectures are *self-balancing by construction*: two channels with different thresholds, accumulating at different rates, rupturing at different frequencies, and applying updates of different magnitudes nevertheless produce exactly balanced total precision-gain without requiring a controller to enforce fairness. In keeping with the broader logic of conservation laws in physics, the conservation does not tell us which thresholds are correct (the phase-gating does that), but it tells us that whatever thresholds the topology selects, the architecture will remain stable. The empirical contribution of this paper accordingly rests on the *phase-gating*: the strongly non-uniform phase relationship between channels ($\chi^2 = 8{,}041$; $p < 10^{-100}$) that determines the functional character of each update. This phase relationship depends on the specific threshold ratio and does not emerge from continuous updating or from arbitrary thresholds. The conservation is the FEP-CRR canvas: it ensures that neither channel dominates the total information rate, while the topology's contribution is to partition that rate into distinct temporal grains. The phase-gating is the painting.

Phase-Gating (P2) We tracked the $\mathbb{Z}_2$ phase (normalised 0 to 1) at the moment of each SO(2) rupture across 60 runs of 1,200 trials (10,345 SO(2) ruptures analysed).

Result. The phase distribution is significantly non-uniform (KS test: $p < 10^{-100}$; $\chi^2 = 8{,}041$ with 9 degrees of freedom); we note that the extreme p-value reflects the large sample size ($n = 10{,}345$ ruptures) rather than implying impossibly

precise measurement. The number of $\mathbb{Z}_2$ ruptures between consecutive SO(2) ruptures is not fixed by the threshold ratio alone but emerges from the interaction between the ratio and the evidence statistics. Testing across threshold ratios confirms this: the modal count scales with the ratio (modal 1 at ratio 1:1, modal 2 at ratio 1.5:1, modal 3 at ratio 2:1) but is not simply equal to it. This is an emergent property, not a necessary consequence of threshold arithmetic.

Compatibility with empirical findings. This result may help explain a puzzling feature of the neuroscience: why does the *timing* of neuromodulatory signals, rather than their magnitude, determine functional output? The simulation produces a non-uniform phase partition from topology alone, without parameter fitting. Jang et al. [4] showed that the functional difference between learning and movement is a matter of tens of milliseconds in the ACh-DA phase relationship. Our simulation recovers an analogous partition from geometric first principles, suggesting that the biological mechanism may reflect topological constraints rather than bespoke neural circuitry.

Coherence Dynamics Fig. 4 shows the coherence accumulation trace for a single agent over 300 trials. The $\mathbb{Z}_2$ channel accumulates rapidly to threshold π and ruptures frequently (mean inter-rupture interval 3.0 trials, CV = 0.26); the SO(2) channel accumulates slowly to threshold 2π and ruptures rarely (mean IRI 10.4 trials, CV = 0.50). The temporal characters are fundamentally different: the $\mathbb{Z}_2$ channel is sub-Poisson with a burstiness index of $B = -0.59$ (nearly clock-like), while the SO(2) channel follows a Gamma distribution. The Cramér-Rao product $C \cdot \Omega$ at rupture is always ≥ 1 (mean 1.27 for $\mathbb{Z}_2$, 1.18 for SO(2)), confirming that rupture occurs at or above the bound but never below it. The overshoot reflects discrete evidence delivery: prediction errors arrive in quanta, so the threshold is crossed rather than saturated exactly. The mean overshoot decreasing from 27% to 18% as the threshold increases is consistent with larger thresholds averaging over more evidence quanta.

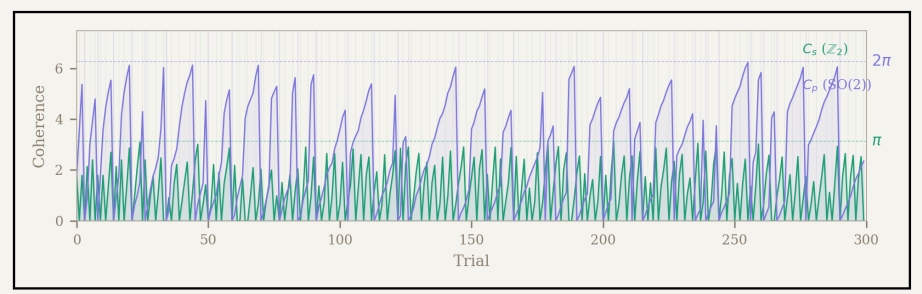

Fig. 4. Coherence accumulation and rupture for a single agent over 300 trials. $\mathbb{Z}_2$ channel (teal) ruptures frequently at threshold π; SO(2) channel (purple) ruptures rarely at threshold 2π.

5 Stress Tests

The results above establish the basic phenomenology: conservation, phase-gating, and a specific temporal decomposition. We now subject these findings to five adversarial tests designed to probe the boundaries of the claim (Fig. 5).

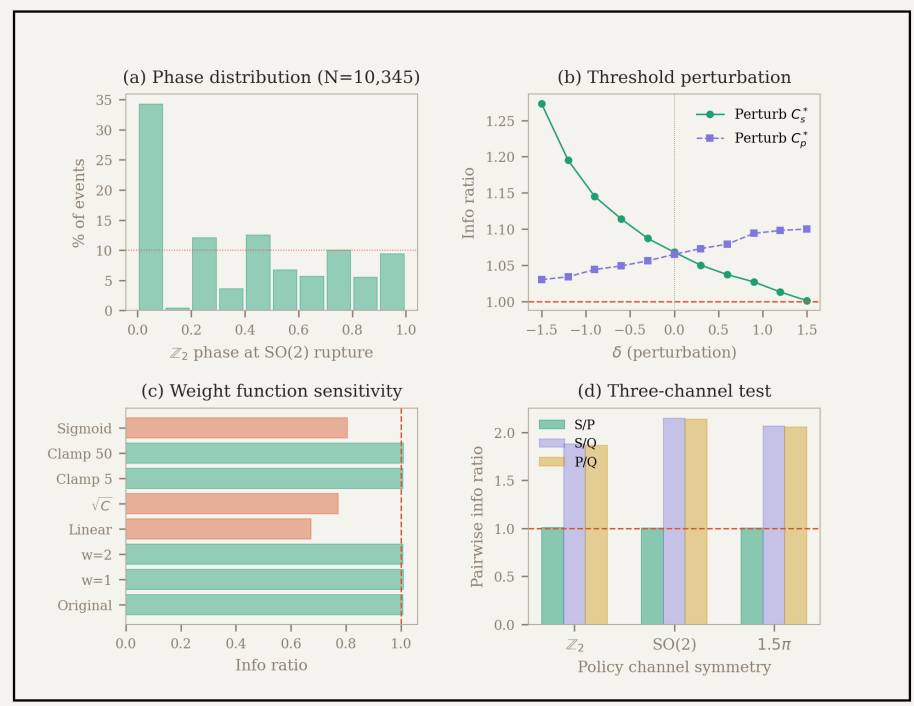

Fig. 5. Stress tests. (a) Phase distribution at SO(2) rupture ($\chi^2 = 8{,}041$). (b) Smooth degradation as thresholds deviate from $\pi/2\pi$. (c) Weight function independence. (d) Three-channel test: sensory/prior conservation holds; pairwise conservation with a third channel fails.

Conservation is claimed to follow from the manifold topology ($\mathbb{Z}_2$ vs SO(2)), which is a property of the agent architecture, not the environment. We tested the same CRR agent in seven additional environment topologies: linear chain ($n = 8$), 2D grid (3×3 and 4×4), Erdős-Rényi random graph ($p = 0.4$ and $p = 0.7$), and binary tree (depth 2 and depth 3). Conservation holds across every environment tested, with all precision-gain ratios within 1.5% of 1.0 and CVs below 2% (Fig. 3d), confirming that conservation is a property of the CRR architecture, not of the environment's graph structure.

We next asked whether conservation is an artefact of the specific clamped exponential weight. Replacing the weight function with seven alternatives (con-

stant, linear, $\sqrt{C}$, sigmoid, alternative clamping thresholds), we found that constant weights preserve conservation perfectly (ratios 1.009 and 1.008), while functions that scale differently between channels break it (ratios 0.67, 0.77, 0.80). Conservation lives in the *thresholds*, not in the weight function (Fig. 5c). Perturbing the thresholds away from their topological values by $\delta \in [-1.5, 1.5]$ produces smooth degradation (Fig. 5b): the ratio deviates from 1.0 monotonically as δ increases. Scaling both thresholds together while preserving the 2:1 ratio maintains conservation across a wide range of k (from 0.25 to 3.0), supporting the claim that the *ratio* of thresholds matters, not their absolute values.

Finally, we tested robustness to correlated evidence and extension to three channels. Conservation holds for positive correlations (even $\rho = 1.0$: ratio 1.002) and mild negative correlations ($\rho = -0.3$: ratio 1.018), but breaks under strong anti-correlation ($\rho = -0.7$: ratio 1.16). Adding a third "policy" channel reveals that conservation is a two-channel phenomenon. The original sensory/prior conservation is unaffected by the third channel (ratio 1.005 regardless of its symmetry), but pairwise conservation with the third channel fails: ratios fall between 1.9 and 2.1 (Fig. 5d). Conservation governs *paired channels sharing equal evidence*, not arbitrary N-channel systems. This limitation is also the source of the architecture's scalability: in a network of Markov-blanketed agents, each edge-node pair constitutes an independent two-channel system, so CRR scales by pairwise composition rather than by adding channels within a single agent. A network of 150 agents has 11,175 edges, but each edge and its two connected nodes form a self-contained $\mathbb{Z}_2/\mathrm{SO}(2)$ pair with its own conservation guarantee. In the simulation, this is visible as one increases the agent count: the system remains temporally coherent at every scale, with no central controller and no parameter adjustment, because each pairwise relationship is independently self-balancing.

6 Discussion

This paper has applied CRR to the problem of coordinating precision channels in Active Inference. The primary finding is phase-gating: $\mathbb{Z}_2$ phase at the moment of SO(2) rupture follows a strongly non-uniform distribution ($\chi^2 = 8{,}041$) whose structure emerges from the interaction between topology and evidence statistics, not from threshold arithmetic alone. The supporting findings are: precision-gain conservation (ratio 1.003, CV 0.024), which we showed is structurally guaranteed by equal evidence delivery; robustness across environment topologies; independence from the weight function, cleanly separating *when* to update from *how much*; and the restriction of conservation to two-channel systems. Both the conservation law and the phase-gating prediction are falsifiable.

These findings bear directly on AGI architecture. Weight-function independence (Section 5.2) frees designers to optimise update rules without disrupting inter-channel coordination. Phase-gating (Section 4.3) eliminates the need for a separate controller to arbitrate between learning and action; the topology does the arbitration. The same mechanism reframes alignment in temporal

terms: what matters is not only *what* a system believes, but *when* it reorganises its beliefs relative to its partner. The phase-gating results show that the functional character of an update depends on where in the partner channel's cycle it occurs, so a phase mismatch between human and AGI would constitute a detectable, measurable form of misalignment. The two-channel limitation (Section 5.4) is itself architecturally informative: because each edge-node pair in a Markov-blanketed network constitutes an independent self-balancing system, distributed architectures that coordinate through local pairwise interactions may have a structural advantage over monolithic designs that must solve the N-channel timing problem centrally. Throughout, the explore-exploit tradeoff, which developmental research frames as a tension between plasticity and stability [28], is parameterised by a single geometric quantity, Ω, a cognitive light cone [1] governing both the depth of accessible history and the breadth of future possibility.

The current test environment is a toy POMDP, not an AGI system; the logical next step is to test phase-gating in a hierarchical generative model with non-stationary dynamics. The restriction to one-dimensional manifolds does substantial work, and the 2:1 ratio's survival in higher dimensions is untested. The precision-gain measure ($\Delta\bar{\alpha}$) is a proxy for Fisher information; confirming conservation under a direct Fisher measure would strengthen the geometric interpretation. The regeneration kernel remains untested but offers a natural extension: specifying the reconstruction resource φ for concrete architectures would extend CRR from a principle about *when* to update to a fuller account of adaptive reconstruction. CRR's rupture condition, that no regime persists indefinitely, connects naturally to mortal computation [29], but testing this implication requires architectures that undergo genuine regime death, not just parameter updates. Nevertheless, the phase-gating mechanism is real, falsifiable, and structurally compatible with empirical findings on neuromodulatory timing [4,18,19]. We offer CRR as a candidate temporal grammar whose limitations we have tried to state as clearly as its strengths. All simulation code is provided in Appendix A.

Acknowledgments. We thank the Active Inference Institute community, Daniel Friedman, and the contributors to the interactive network simulation from which this finding emerged.

Disclosure of Interests. The author declares no competing interests.

Use of AI. Claude Opus 4.6 (Anthropic) was used as a research tool for independent replication, stress testing, literature search, and as a latex drafting aid, and in the design of the pedagogical simulation at https://www.temporalgrammar.ai/human_networks.html. All theoretical arguments, simulation design decisions, and scientific interpretations are the origin and responsibility of the named author.

A Implementation

The complete CRR agent, ring HMM environment, and all stress test scripts (approximately 1,500 lines of Python) are available as supplementary material and at https://temporalgrammar.ai/CRR_PG.py. The core constants are: $\Omega_{\mathbb{Z}_2} = 1/\pi$, $\Omega_{SO(2)} = 1/2\pi$, $C^*_{\mathbb{Z}_2} = \pi$, $C^*_{SO(2)} = 2\pi$. The gain function clamps the exponential weight: $w = \min(\exp(C/\Omega),\, 10)/10 \times 2 + 0.5$. Prediction errors are computed as surprisal at true states under the agent's current model. Both channels receive exactly one evidence count per trial.

References

1. Levin, M.: Technological approach to mind everywhere: an experimentally-grounded framework for understanding diverse bodies and minds. Frontiers in Systems Neuroscience **16**, 768201 (2022)
2. Friston, K.J.: The free-energy principle: a unified brain theory? Nature Reviews Neuroscience **11**(2), 127–138 (2010)
3. Parr, T., Pezzulo, G., Friston, K.J.: Active Inference: The Free Energy Principle in Mind, Brain, and Behavior. MIT Press (2022)
4. Jang, H.J., Ward, R.M., Golden, C.E.M., Constantinople, C.M.: Acetylcholine demixes heterogeneous dopamine signals for learning and moving. Nature Neuroscience (2026). https://doi.org/10.1038/s41593-026-02227-x
5. Fisher, R.A.: Theory of statistical estimation. Math. Proc. Cambridge Philos. Soc. **22**, 700–725 (1925)
6. Rao, C.R.: Information and the accuracy attainable in the estimation of statistical parameters. Bull. Calcutta Math. Soc. **37**, 81–91 (1945)
7. Ito, S., Dechant, A.: Stochastic time evolution, information geometry, and the Cramér-Rao bound. Physical Review X **10**(2), 021056 (2020)
8. Amari, S., Nagaoka, H.: Methods of Information Geometry. AMS, Providence (2000)
9. Čencov, N.N.: Statistical Decision Rules and Optimal Inference. AMS, Providence (1982)
10. Da Costa, L., Friston, K., Heins, C., Pavliotis, G.A.: Bayesian mechanics for stationary processes. Proc. R. Soc. A **477**(2256), 20210518 (2021)
11. Kirchhoff, M., Parr, T., Palacios, E., Friston, K., Kiverstein, J.: The Markov blankets of life: autonomy, active inference and the free energy principle. J. R. Soc. Interface **15**(138), 20170792 (2018)
12. Friston, K., Da Costa, L., Sakthivadivel, D.A.R., et al.: Path integrals, particular kinds, and strange things. Physics of Life Reviews **47**, 35–62 (2023)
13. Tolchinsky, A., Levin, M., Fields, C., Da Costa, L., Murphy, R., Friedman, D., Pincus, D.: Temporal depth in a coherent self and in depersonalization: theoretical model. Frontiers in Psychology **16**, 1585315 (2025)
14. Nicholson, S.B., García-Pintos, L.P., del Campo, A., Green, J.R.: Time-information uncertainty relations in thermodynamics. Nature Physics **16**, 1211–1215 (2020)
15. Nishiyama, T., Hasegawa, Y.: Unified speed limits in classical and quantum dynamics via temporal Fisher information. arXiv:2504.04790 (2025)
16. Parr, T., Da Costa, L., Friston, K.: Markov blankets, information geometry and stochastic thermodynamics. Phil. Trans. R. Soc. A **378**(2164), 20190159 (2020)

17. Doya, K.: Metalearning and neuromodulation. Neural Networks **15**(4–6), 495–506 (2002)
18. Krok, A.C., Maltese, M., Mistry, P., Miao, X., Li, Y., Bhatt, D.K., Bhatt, M.A., Tritsch, N.X.: Intrinsic dopamine and acetylcholine dynamics in the striatum of mice. Nature **621**(7979), 543–549 (2023)
19. Taniguchi, J., Tritsch, N.X.: A mismatch between striatal cholinergic pauses and dopaminergic reward prediction errors. Proc. Natl. Acad. Sci. USA **121**(41), e2410828121 (2024)
20. Wickens, J.R., et al.: Coincidence of cholinergic pauses, dopaminergic activation and depolarisation drives synaptic plasticity. Nature Communications **13**, 1296 (2022)
21. Aosaki, T., et al.: Acetylcholine-dopamine balance hypothesis in the striatum: an update. Geriatrics & Gerontology International **10**(s1), S148–S157 (2010)
22. Iglesias, S., et al.: Cholinergic and dopaminergic effects on prediction error and uncertainty responses during sensory associative learning. NeuroImage **226**, 117590 (2021)
23. Albarracin, M., Pitliya, R.J., St Clere Smithe, T., Friedman, D.A., Friston, K., Ramstead, M.J.D.: Shared protentions in multi-agent active inference. Entropy **26**(4), 303 (2024)
24. Fields, C., Friston, K., Glazebrook, J., Levin, M., Marcianò, A.: The free energy principle induces neuromorphic development. Neuromorphic Computing and Engineering **2**(4), 042002 (2022)
25. Ramstead, M.J.D., Sakthivadivel, D.A.R., Heins, C., et al.: On Bayesian mechanics: a physics of and by beliefs. Interface Focus **13**(3), 20220029 (2023)
26. Wootters, W.K.: Statistical distance and Hilbert space. Physical Review D **23**, 357–362 (1981)
27. Jaynes, E.T.: Information theory and statistical mechanics. Physical Review **106**, 620–630 (1957)
28. Gopnik, A.: Childhood as a solution to explore-exploit tensions. Phil. Trans. R. Soc. B **375**, 20190502 (2020)
29. Ororbia, A., Friston, K.: Mortal computation: a foundation for biomimetic intelligence. arXiv preprint arXiv:2311.09589 (2023)