

The Combination Problem as a Dynamics Problem: How Coupled Systems Come to Share a Point of View

Jeff Wozniak

Independent Researcher [ORCID: 0009-0002-3436-6863](#)

July 22, 2026

Abstract

A composite system can behave as though it had a single point of view, even though each of its subsystems acts relative to an orientation of its own — a perspective. We present a dynamical mechanism for how such unity arises. When one subsystem conditions on another’s record the way a Bayesian agent conditions on evidence, shifting toward what the other registered, their perspectives draw together, and a network of subsystems coupled this way settles into one shared perspective. The belief-update character of the coupling is essential: a coupling that merely registers another’s record, as a measuring instrument does, exchanges information without limit yet never produces alignment. The mechanism carries testable structure. The shared perspective is held more tightly as connections grow dense; it fragments into separate perspectives when connections are severed, a structural analogue of split-brain dissociation; and it cannot form in a one-dimensional chain, where fluctuations grow without bound, first surviving in two dimensions up to a slow growth with network size. A concrete open-system model realizes all three results. The mechanism supplies multi-subject dynamics for projective models of a single perspective, casts binding as a form of active inference, and complements information-based accounts by identifying the form of coupling on which shared perspective turns. Throughout, “perspective” is structural — the orientation in which a subsystem acts — and the results are silent on subjective experience.

1 Introduction

The Projective Consciousness Model represents a subject’s perspective as a point in a projective geometry, moved by the minimization of variational free energy ([Rudrauf et al., 2017](#)). The representation has made the single perspective mathematically tractable, and it presupposes that there is one perspective to model. A brain is many subsystems, each acting relative to an orientation of its own; that these compose one point of view is a fact about the composite that a theory of perspective should eventually explain. Put structurally, this is the binding problem, and its philosophical ancestor is the combination problem that William James pressed under the heading of mind-dust ([James, 1890](#)) and that Chalmers, Goff, and Seager have kept in view ([Chalmers, 1996](#); [Goff, 2017](#); [Seager, 1995](#)). That work has treated the problem chiefly at the level of what could possibly compose. This paper supplies the dynamics.

Give each subsystem a state and a perspective. The perspective is an orientation θ : the basis the subsystem’s next act settles into, and equivalently the question that act asks. Subsystems leave records, and a record carries the orientation in which it was written. One subsystem reads another by conditioning its own dynamics on that record. The question becomes concrete. Which coupling draws the perspectives of coupled subsystems into alignment, so that the composite carries one perspective rather than many?

The paper answers in three moves. A coupling that merely registers another subsystem’s record — the way a measuring instrument registers an outcome — raises the mutual information between the two subsystems while exerting no restoring force on their orientations; competing registrations of this kind average the orientation away rather than align it. We establish this at the couplings where it can be established, and an integration scalar cannot see the difference: a maximally integrated pair can prefer opposed orientations. A belief-updating coupling binds. Reading a record as a Bayesian agent reads evidence, the subsystem shifts toward the recorded orientation; the resulting alignment persists in the slow steady limit, does not depend on how the subsystems started, and has a direction fixed by the coupling itself. Two premises about a single record pin this coupling uniquely. The shared perspective then obeys a quantitative law: it is held more tightly the more densely the network is connected, it fragments into independent perspectives when the network disconnects, and it fails to form at all in a one-dimensional chain, first surviving in two dimensions up to a slow growth with network size.

The construction is spare. It needs only subsystems carrying perspectives, records that carry perspectives, and a read that conditions one subsystem on another’s record. The three results follow from that structure together with open-system and graph-theoretic reasoning. A minimal model (§7) supplies concrete operators, and every result before it stands without them. The setup meets several literatures. The reads are Lindblad channels of open quantum systems (Lindblad, 1976; Gorini et al., 1976; Breuer and Petruccione, 2002). The Projective Consciousness Model already equips a viewpoint with a projective orientation, and the belief-updating coupling supplies it with multi-subject dynamics (Rudrauf et al., 2017). The orientation carried here, a single angle on the circle of preferred bases, is a reduced stand-in for that model’s fuller projective field of viewpoint, and the claim is to add multi-subject dynamics in general rather than to reproduce its perceptual geometry. Active inference casts a system’s dynamics as free-energy minimization over such an orientation, and the coupling that binds is a belief update at the level of the channel (Friston, 2010; Parr et al., 2022). Information-based accounts of unity gain a complement: whether coupled subsystems share a perspective turns on the form of the coupling between them and is a separate question from how much integrated information their joint state carries (Tononi, 2004; Oizumi et al., 2014).

Two things are worth stating at the outset. What binds many subsystems into one perspective is a coupling that is not itself one of the subsystems’ orientation-fixing acts; it lives between them, and its form is fixed by a single premise the surrounding structure does not supply. And throughout, the perspective is structural. It is the orientation in which a subsystem acts, and nothing here claims that a shared perspective is or produces an experience.

What is new here is the distinction between a coupling that registers another subsystem’s record and one that updates on that record as evidence, together with the no-go that separates them: registration raises shared information yet exerts no restoring force on the orientations, while a belief update aligns them. Two premises about a single record characterize the binding coupling uniquely, and the binding coupling is thereby an active-inference update at the channel level. The collective consequences are inherited. The stiffness law and its dimensional threshold are of the classical XY /Kuramoto spin-wave and Mermin–Wagner type (Kuramoto, 1984; Acebrón et al., 2005; Mermin and Wagner, 1966), and the open-system channel formalism is standard (Lindblad, 1976; Gorini et al., 1976; Breuer and Petruccione, 2002). What the paper adds at the network level is the coupling that generates the weighted graph, not the spin-wave dynamics that run on it.

2 Setup: perspectives, records, and reads

A *subsystem* carries two things: a density matrix ρ on a Hilbert space $\mathcal{H}$, and a *perspective*, an orientation θ on the projective circle of preferred bases. The perspective is structural. It fixes the basis the subsystem's next act settles into, and equivalently the question that act asks. We take the plane of orientations to be a qubit, with

$$|\psi(\theta)\rangle = \cos \theta |0\rangle + \sin \theta |1\rangle, \quad |\psi_\perp(\theta)\rangle = -\sin \theta |0\rangle + \cos \theta |1\rangle, \quad (1)$$

and higher-dimensional spaces carry the same plane with inert spectator directions, except at one point (a genuinely two-component perspective) treated in §4. The subsystem's act projects onto the recorded ray; the operator that both performs the act and stores it is the rank-one projector

$$A(\theta) = |\psi(\theta)\rangle\langle\psi(\theta)|. \quad (2)$$

This is a director rather than a vector: $A(\theta)$ is invariant under $|\psi\rangle \mapsto -|\psi\rangle$ and has period π in θ , so the orientation lives on a nematic circle and every coupling we derive is π -periodic.

A subsystem that acts leaves a *record*. The record is stamped with the orientation in force when the act fired, so it carries the value θ_j of the subsystem j that wrote it. A second subsystem i reads the record by conditioning its own dynamics on the recorded content. The read is a completely positive channel on $\mathcal{H}_i$, written in Gorini–Kossakowski–Sudarshan–Lindblad (GKSL) form (Lindblad, 1976; Gorini et al., 1976; Breuer and Petruccione, 2002)

$$\mathcal{L}[\rho] = -i[H, \rho] + \sum_k \left(L_k \rho L_k^\dagger - \frac{1}{2} \{L_k^\dagger L_k, \rho\} \right), \quad (3)$$

with jump operators L_k and a Hamiltonian H built from the recorded orientation θ_j .

Three reads span the cases we need, and the distinction among them carries the results.

- A *projective* read treats the record as a basis to dephase in. Its jump operator is the record's own projector $A(\theta_j)$, a Hermitian and hence normal operator. This is a measurement-like coupling: it registers that subsystem j acted in the orientation θ_j and erases the reader's coherence in that basis.
- A *coherent* read couples the two subsystems through a Hamiltonian built from the same projectors, $H \propto A(\theta_j) \otimes A(\theta_j)$. It is unitary and, like the projective read, blind to direction along the recorded axis.
- An *evidential* read treats the record as evidence and shifts the reader's state toward the recorded value, the way a belief update moves toward what it conditions on. Its jump operator is the transition

$$T(\theta_j) = |\psi(\theta_j)\rangle\langle\psi_\perp(\theta_j)|, \quad (4)$$

a non-normal operator ($TT^\dagger = A(\theta_j)$, $T^\dagger T = |\psi_\perp(\theta_j)\rangle\langle\psi_\perp(\theta_j)|$, $[T, T^\dagger] \neq 0$) whose unique dark state is the recorded ray $|\psi(\theta_j)\rangle$. Its dissipator damps population into $|\psi(\theta_j)\rangle$, an amplitude-damping pull toward θ_j .

Two scalar quantities track two different things, and we keep them apart throughout. The first is *integration*, $\Phi[\rho]$, a mutual-information-type summary of the joint state, here the mutual information under the partition that minimizes it. It is a functional of the density matrix alone

and says nothing about the orientations. The second is the *alignment torque*, the gradient of the orientation free energy

$$F[\rho_i, \theta_i] = -\log \text{Tr}[A(\theta_i) \rho_i], \quad \text{torque} = \frac{\partial F}{\partial \theta_i}, \quad (5)$$

evaluated on the reduced state ρ_i that the coupling produces. The free energy is minimized when the reader’s principal axis aligns with its own orientation; a nonzero torque that depends on the writer’s orientation θ_j is the signature that reading subsystem j has moved subsystem i ’s perspective. A third scalar appears later, the connectivity Φ_λ of the coupling graph (§5), a measure of how densely the network is connected; it is a spectral property of a graph Laplacian and is a different object from the density-matrix integration $\Phi[\rho]$. Integration measures how much the subsystems’ states share; connectivity measures how rigidly their perspectives are tied. The first result is that the two can come apart completely.

The perspective is the projective viewpoint of the Projective Consciousness Model, which represents a subject’s outlook as a projective orientation on a space of possibilities (Rudrauf et al., 2017). The evidential read is a belief update at the level of the jump operator: active inference casts a system’s dynamics as the minimization of a free-energy functional over such an orientation, and treats conditioning on evidence as a pull of the posterior toward the observed value (Friston, 2010; Parr et al., 2022; Friston et al., 2017). The projective read decoheres in the record’s basis; the evidential read performs the posterior pull. The distinction between the two carries what follows, and everything in this section is state, orientation, and channel.

The state space and channels serve here as a general formalism for stochastic dynamics and open-system evolution, and the construction has a classical Bayesian reading. The orientation θ is a parameter each subsystem estimates, a record is the evidence one subsystem leaves for another, and the belief-updating coupling is ordinary Bayesian conditioning, the active-inference update in which a posterior shifts toward what it conditions on. The density-matrix presentation is the general one, and the concrete model of §7 is the realization on which the numbers are computed. The formalism is a general state-space and dynamics language, and the paper makes no quantum-mechanical claim about brains or about the physical basis of cognition.

The mechanism does not depend on quantum mechanics, and the two-subsystem results read directly in classical terms. Write each orientation as a point on the circle. The projective read is symmetric diffusion of that point with no drift: the reader accumulates information about the neighbor’s recorded value while its own orientation wanders, and no term pulls it toward the record, which is the no-go of §3 with the alignment torque set by a drift that is absent. The evidential read adds an Ornstein–Uhlenbeck drift of the orientation toward the recorded value, the sign of the drift fixing the direction of alignment; this is the belief update of §4, the posterior sliding toward the evidence. The single-event weight is the squared overlap $\cos^2(\theta_i - \theta_j)$ of §6, the classical likelihood of a match between reader and record. The network layer is classical already: the stiffness law of §5 runs on overdamped Langevin dynamics of a nematic XY model on the coupling graph, of the Kuramoto synchronization type. The density-matrix formalism states the mechanism in general; these readings are the same mechanism with the Hilbert space removed.

3 Registration without alignment

A projective read couples the two subsystems and raises their shared information. It does not move their perspectives. We establish this at the reads where it holds, the projective read and the coherent read, and prove it in Appendix A; the general question of which couplings bind stays open, since the argument here uses only these two.

Consider subsystem i reading subsystem j 's record projectively while it also continues to act in its own orientation. In the slow steady state, where the state relaxes fully between orientation updates, the reduced state ρ_i is governed by two dephasing channels, one about the reader's own axis θ_i and one about the recorded axis θ_j . Each dephasing channel is symmetric along its axis and has no preferred direction on it. A short calculation on the Bloch ball (Appendix A) gives the steady axis explicitly. When the two dephasing rates are balanced, the steady Bloch vector collapses to the center: the orientation washes out, and ρ_i has no principal axis at all. When the local channel dominates, the steady axis pins at the reader's own orientation θ_i . In both regimes the orientation free energy (5) evaluated at the steady ρ_i is independent of θ_j , so the alignment torque vanishes. Reading subsystem j has changed the reader's coherence and left its perspective where it was.

Proposition 1 (No alignment under registration). *Under a projective record read $A(\theta_j)$, or a coherent read built from the same projectors, the reader's steady orientation is independent of the writer's orientation θ_j , and the induced alignment torque is zero in the slow steady state, independently of the initial state.*

The numerics are clean. On the model operators (§7, Appendix D), the projective read pins the reader's steady axis at exactly θ_i for every writer orientation θ_j across the swept range, with a torque variance of order 10^{-11} ; the coherent read gives the same to a comparable tolerance. The result holds identically across product, correlated, anti-correlated, and maximally mixed initial states. The read integrates the two subsystems, and their mutual information is positive, while the perspective does not follow. Figure 1 shows the two couplings side by side: under the registering read the shared information fills while the alignment stays flat, and under the belief-updating read of §4 both rise together.

What the projective read leaves out is the direction of the alignment. A single example makes this concrete. Take the maximally integrated two-subsystem state, the anti-correlated pair. Its integration is large and positive, $\Phi[\rho] = 5.545$ on the model operators. The orientation configuration this state certifies is the one in which the two subsystems' perspectives are *opposed*. A high integration scalar sits on top of a frustrated orientation relation. Mutual information counts how much two subsystems' states covary; it does not record whether they covary in agreement or in opposition, because it is a function of the state's spectrum and is invariant under the local basis change that flips alignment into frustration. Figure 2 shows the gap directly.

The point is narrow, and it bears on any account that reads a scalar off the joint state and calls it shared perspective.

Whether two coupled subsystems come to share a perspective is fixed by the form of the coupling between them, not by how much integrated information their joint state carries. Alignment turns on the direction the coupling sets, and an integration scalar is invariant under the operation that reverses that direction.

Information-based accounts that summarize unity in a scalar read off the joint state meet this gap directly: the same $\Phi[\rho]$ is consistent with an aligned perspective and with a maximally frustrated one, so what fixes the perspective is carried by the coupling rather than by any such scalar. The result complements such a summary rather than competing with it. Integrated-information theory itself locates unity in a maximally irreducible cause-effect structure that already carries more than

²This figure and those that follow illustrate the classical reduction of the mechanism (§2), in which each subsystem carries a single orientation angle in the 2π (Kuramoto) convention rather than the π -periodic nematic director; the qualitative behavior matches the formal results, and the opposed-orientation case of §3 corresponds to the 2π -opposed pair shown in Fig. 2. The panels are produced by the companion notebook `binding-demos.ipynb`.

Beat 1 - sharing information is not sharing a point of view




Figure 1: Sharing information is not sharing a perspective. Under a registering (projective) read, the shared information between two subsystems fills to its ceiling while their orientations stay scattered and the alignment order parameter R holds near its unaligned floor. Under a belief-updating (evidential) read, the orientations swing parallel and the shared information and the alignment rise together. The contrast is the no-go of §3 against the binding read of §4.²

a state scalar (Tononi, 2004; Oizumi et al., 2014; Tononi et al., 2016); the complement here is to any account that reads shared perspective off a scalar summary of the joint state. The next section exhibits the coupling that carries the direction a scalar discards.

4 The belief-updating coupling and its uniqueness

The evidential read produces the alignment the projective read does not. Its jump operator is the transition $T(\theta_j)$ of (4), and its dissipator damps population into the recorded ray $|\psi(\theta_j)\rangle$. Run symmetrically between two subsystems, each pulled toward the other’s recorded orientation, it produces an alignment torque already present in the exact slow steady state. On the model operators the steady drift points toward the writer’s orientation for every writer orientation, with a fixed point exactly at coincident orientations and a torque variance of 0.9735, roughly eleven orders above the projective read’s floor. The torque rotates the reader’s principal axis smoothly from one orientation toward the other while the axis stays sharply defined, so the effect is a genuine reorientation and not the washing-out of the balanced projective case. The drift is identical across product, correlated, anti-correlated, and maximally mixed initial states, to four digits. The coupling binds, and it binds independently of the state it starts from. In the classical reading this is a conditioning step: the




Figure 2: An integration scalar cannot see the alignment direction. The aligned pair and the opposed pair carry identical mutual information yet sit at opposite ends of the alignment order parameter R . Sweeping the offset from aligned to opposed holds the mutual information flat while the alignment falls from one to zero. What fixes whether two subsystems share a perspective is the form of the coupling, which the belief-updating read of §4 supplies.

reader’s estimate of the orientation moves toward the value the record reports.

The direction of the binding is carried by the operator. The dark target of $T(\theta_j)$ is the recorded state $|\psi(\theta_j)\rangle$, and pulling toward the recorded state aligns the reader with the writer: the coupling is ferromagnetic, with a common target at coincident orientations. Reversing the operator to pull toward the orthogonal state, $T_{\text{away}} = |\psi_{\perp}(\theta_j)\rangle\langle\psi(\theta_j)|$, flips the coupling to anti-alignment. A global phase on T leaves the dissipator unchanged and so leaves the direction untouched. The direction is set by which state the read pulls toward, and by nothing else. This is what supplies what §3 found missing: where an integration scalar discards the alignment direction, the belief-updating operator fixes it, so the anti-correlated configuration that a scalar cannot distinguish is aligned by the read to coincident orientations at every relaxation timescale, including the deep slow regime.

4.1 Two premises

The evidential read is one specific coupling among the belief-updating channels, and two premises about a single record fix it uniquely. Both are stated per rank-one record, a scope made precise below.

Definition 1 (Deposit). The read’s total gain operator is the record’s own act operator: $\sum_k L_k L_k^{\dagger} = A(\theta_j) = |\psi(\theta_j)\rangle\langle\psi(\theta_j)|$. The population the read deposits lands in exactly the state the record certifies.

Definition 2 (Darkness). The recorded state is a dark state of the read: $L_k|\psi(\theta_j)\rangle = 0$ for every k . A reader already at the recorded orientation is inert, and the read writes no record of its own.

Each premise is doing exactly one job, and the two jobs are visible as the exclusion of two distinct families of alternative reads. A read satisfying darkness alone may still deposit into a *rotated* ray: the shear family $L_{\varepsilon} = |\psi(\theta_j + \varepsilon)\rangle\langle\psi_{\perp}(\theta_j)|$ is dark for every offset ε and reaches the recorded ray as its attractor, but its gain is $A(\theta_j + \varepsilon)$, so it overshoots the record. Deposit excludes it. A read satisfying deposit alone may still re-record: the admixture family $\{\sqrt{1-\mu}A(\theta_j), \sqrt{\mu}T(\theta_j)\}$ has total gain $A(\theta_j)$ for every mixing μ , but for $\mu < 1$ its projector term dephases in the record basis, which is the reader writing a fresh record. Darkness excludes it. The two families meet at one point, the transition $T(\theta_j)$, and the following theorem says the conjunction pins it.

Theorem 1 (Characterization). *Among GKSL conditioning channels on a single rank-one record, deposit and darkness force the generator to $\mathcal{L} = \mathcal{D}[T(\theta_j)]$, uniquely up to Kraus and phase gauge and an overall rate. The record is the channel's unique attractor, and the alignment sign is fixed to ferromagnetic; both are consequences of the two premises rather than additional assumptions.*

The proof is short and appears in Appendix B. Deposit forces every jump operator to have its range on the recorded ray, so $L_k = |\psi(\theta_j)\rangle\langle b_k|$; darkness forces each $|b_k\rangle$ orthogonal to the recorded ray, hence proportional to $|\psi_\perp(\theta_j)\rangle$; the two together leave $T(\theta_j)$ and its gauge copies. Record-as-attractor and the ferromagnetic sign then follow, so neither has to be posited. Deposit alone leaves the admixture family, darkness alone leaves the shear family, and attractor uniqueness paired with either one still leaves the other family; only the conjunction of deposit and darkness closes the characterization.

4.2 Deposit as a motivated definition

Deposit stipulates a faithful read, and the stipulation is motivated rather than derived. Among the reads that reach the record, the deposit read dissipates least: its jumps land on the recorded state and stop, while every shear overshoots into a rotated ray and pays for the correction in further acts, and every admixture adds acts by re-recording. Writing the per-act entropy floor as ΔS , the entropy a read deposits is ΔS times its integrated activity, and deposit minimizes that activity at fixed target; the minimum sits at the deposit read on the model operators and stays there under reweighting by the per-act cost. The standing of the premise, why a binding coupling should take the least-dissipation faithful form, is taken up in §8.

4.3 Per-record scope

Deposit and darkness are premises about a single rank-one record. A perspective built from several components is several records at once, one per component, and the read is applied to each. For a two-component perspective with records θ_j^K and θ_j^G on orthogonal sectors, the act operator is the sum of two rank-one projectors, and the bound read is the direct sum $T(\theta_j^K) \oplus T(\theta_j^G)$ of the per-sector transitions. Applied to the sectors separately the characterization is the single-record theorem in each, and the composite is covered. Applied to the lumped rank-two projector as if it were one record, deposit and darkness leave a larger family, since a read may then damp one sector's orthogonal direction into the other sector while still satisfying both premises; the per-record reading removes this freedom. The scope statement is therefore load-bearing wherever the perspective has more than one component, and it is the reason the premises are phrased per record.

5 Stiffness and dissociation

Many subsystems, each reading its neighbors' records, form a network. Place one node per subsystem and one edge per record-exchange coupling. The evidential read fixes a common target on each edge and a nonnegative weight, so the network is a weighted graph G with a nonnegative Laplacian L_G and no frustration. The shared perspective is an elastic medium on this graph. Small departures of each subsystem's orientation from the network mean behave as spin waves, and the quantitative behavior of the shared perspective follows from the spectrum of L_G .

Theorem 2 (Perspective stiffness). *For a connected coupling graph G on N nodes with weighted Laplacian L_G and eigenvalues $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_N$, the Gaussian fluctuations of the orientation*

about its network mean obey

$$\text{Var}(\theta_i - \bar{\theta}) = \frac{D}{4} \frac{1}{N} \sum_{n \geq 2} \frac{1}{\lambda_n(L_G)}, \quad (6)$$

where D is the orientation-noise floor. The mean orientation $\bar{\theta}$ exists and is unique modulo the global orientation symmetry.

The derivation is the standard spin-wave expansion of a nematic XY interaction on a graph, carried out in Appendix C; the collective law is a spin-wave consequence of the classical XY /Kuramoto type (Kuramoto, 1984; Acebrón et al., 2005), and the contribution here is the coupling that generates the weighted graph, not the collective dynamics on it. The noise floor D sets the scale of the variance and is left as a parameter here; its natural value is the per-act entropy floor, and the law fixes the dependence on the spectrum rather than an absolute variance. Stated plainly: the shared perspective fluctuates less the more densely the network is connected. The spectral sum is dominated by its softest term when the graph has a spectral gap, in which case

$$\text{Var}(\theta_i - \bar{\theta}) \propto \frac{1}{\Phi_\lambda}, \quad \Phi_\lambda = \lambda_2(L_G), \quad (7)$$

where Φ_λ is the algebraic connectivity of the coupling graph, the spectral measure of how densely it is connected. The denser the connections, the tighter the shared perspective. This Φ_λ is a different object from the density-matrix integration $\Phi[\rho]$ of §3: $\Phi[\rho]$ measures how much the subsystems' states share, Φ_λ measures how rigidly their perspectives are tied, and the result of §3 is that the first does not control the second.

Fragmentation at disconnection. When the coupling graph splits into disconnected components, $\lambda_2 \rightarrow 0$ and the softest spin-wave mode goes free. The shared perspective stops being one perspective. Each connected component carries its own mean orientation, and the components drift independently. On the model network this is exact: the variance within a component stays constant as the bridge between two clusters is weakened, while the variance of the orientation difference across the bridge diverges as its coupling vanishes. The shared perspective fragments into independent per-component perspectives. This is the model's dissociation limit, the point at which a single perspective ceases to hold across the whole network and becomes two — the structural counterpart of the dissociation seen when the callosal connections between the hemispheres are severed (Sperry, 1968). Figure 3 shows both the stiffening and the tear.

Locking requires two dimensions. A one-dimensional chain of coupled orientations never locks. The spin-wave sum (6) grows with system size in one dimension, so the orientation variance is unbounded and long-range order is destroyed by fluctuations, the lower-critical-dimension result familiar from the Mermin–Wagner theorem for continuous symmetries (Mermin and Wagner, 1966). Two dimensions is the threshold. On a finite two-dimensional network a shared perspective survives, held with a variance that grows only logarithmically in the network size; the same logarithm diverges in the thermodynamic limit, so the locking is marginal, a quasi-long-range order rather than strict order. Size-independent locking, with a convergent spectral sum and a variance that stays bounded as the network grows, sets in only above two dimensions. The shared perspective is a sheet rather than a line for this reason, and the dimensional threshold is a prediction of the stiffness law rather than an input to it. Figure 4 contrasts the two.




Figure 3: A shared perspective stiffens with connectivity and fragments when cut. Left: an intact two-dimensional sheet holds a single shared orientation with small variance. Center: severing the sheet along the dashed line splits it into two components, each settling into its own orientation, the fragmentation limit of this section and its split-brain counterpart. Right: across coupling graphs the orientation variance falls as the algebraic connectivity $\lambda_2 = \Phi_\lambda$ grows, tracking the $1/\Phi_\lambda$ law of (7), so denser connections hold the shared perspective more tightly.

Closure from the channel to the network. The stiffness law is derived at the network level, and its edge weight is the coupling produced by the two-subsystem evidential read. The two levels close. Fitting the exact slow steady drift of the two-subsystem read to a nematic pair coupling gives an edge weight $J(\alpha)$ that grows linearly with the read fraction α . Feeding that weight into the network and comparing the stiffness law (6) to a direct simulation of the orientation dynamics, the two agree to within one to five percent across five graph topologies: a ring with next-nearest neighbors, a two-dimensional torus, a honeycomb lattice, a triangular lattice with dense odd cycles, and the complete graph on twelve nodes. The proportionality to $1/\Phi_\lambda$ holds across a range of couplings with a coefficient of variation of six percent, and the fragmentation limit reproduces as described. The triangular lattice is the graph most apt to expose frustration, and it shows none: the operator-fixed ferromagnetic direction keeps every edge weight nonnegative, so no odd cycle is frustrated. The two-subsystem coupling and the network-scale stiffness are the same physics at two scales.

6 The single-event weight

The read carries a natural weight, and it has the Born form. The single-event deposit weight from a reader at orientation θ_i is the gain contraction $\text{Tr}[G \rho_i]$ with the deposit gain $G = A(\theta_j)$, which for a reader in the state $|\psi(\theta_i)\rangle$ is

$$\text{Tr}[A(\theta_j) \rho_i] = |\langle \psi(\theta_j) | \psi(\theta_i) \rangle|^2 = \cos^2(\theta_i - \theta_j). \quad (8)$$

The complementary weight, the rate of the read's loss channel, is $\sin^2(\theta_i - \theta_j)$, and the two partition unity. Both hold to machine precision on the model operators. The Born form is present by construction: the deposit gain is a rank-one projector, a quadratic object in the orientation amplitude, so the weight it induces is the squared overlap.

The weight is linear in the reader's state because the read is a completely positive channel, and its gain is a positive operator; the weight is therefore a quadratic form in the orientation amplitude,




Figure 4: A line cannot hold a shared perspective; a sheet can. Left: in a one-dimensional chain the shared orientation washes out, the alignment order parameter falling toward its unaligned value. Center: a two-dimensional sheet of the same size holds a shared orientation. Right: the orientation variance of the chain grows with system size, so a line never locks, while the sheet’s variance rises only slowly with size, the marginal logarithmic size dependence of the two-dimensional case, which holds a shared perspective on any finite network short of strict locking in the thermodynamic limit.

hence the squared overlap. On a single plane of orientations the deposit gain can be pinned by two scalar conditions on this weight, a certainty condition that an aligned read fires with unit weight and an orthogonality condition that an orthogonal reader gains nothing, which together with positivity fix the gain to the record’s projector. The operator premise of §4 and this pair of scalar conditions are two statements of the same content, per rank-one record.

This grounds the Born weight in the read dynamics, and it is not a second independent derivation of it. A derivation from the symmetry of the orientation and a derivation from the dynamics of the read rest on the same root, the quadratic structure of the two-norm: the symmetry route selects the squared overlap as the unique quadratic invariant of the orientation, and the dynamics route inherits the same quadratic form from the positivity of the channel’s gain (Gleason, 1957; Busch, 2003; Born, 1926). The literature that forces the exponent by composition or fine-graining imports the same two-norm as its counting measure (Zurek, 2003; Masanes et al., 2019; Wallace, 2012). The two routes agree because they share a root, and the agreement is a consistency check rather than an over-determination.

7 A minimal model

The construction needs a concrete realization that supplies subsystems with perspectives, records that carry perspectives, and a read. The simplest one is a finite-state model, and the numerical operators used above are its operators. Each subsystem is a small finite-state system carrying an orientation θ on the plane of preferred bases, with the two-outcome doublet structure of §2: its act is the rank-one projector $A(\theta)$ of (2), and its orientation is the coordinate that projector selects. A record is a value stamped with the orientation in which it was produced. The registering coupling is the coincidence of two acts, carried by the projector $A(\theta_j)$; the belief-updating coupling is the transition $T(\theta_j)$ of (4), which pulls a reader toward a recorded value without depositing a record of its own. A network is a graph of such subsystems, one node each and one edge per record-exchange coupling. The two-component perspective of §4 is a subsystem carrying two such records at once, which is why the per-record scope is the operative statement there rather than a technicality.

This minimal model is what the reported computations run on. The result of §3, the evidential

read and its characterization, the stiffness law, and the single-event weight are statements about subsystems, perspectives, records, and channels, and each holds for any realization of that structure — subsystems carrying orientations, records, and reads. The finite-state model is the simplest such realization; the computations reported throughout are carried out on it, and a reader who wants only the theorems needs nothing beyond the abstract structure.

8 Discussion

The result of §3 meets accounts that locate unity in an integration scalar, and it meets them as a complement. Whether coupled subsystems come to share a perspective turns on the direction the coupling sets, and an integration scalar is invariant under the operation that reverses that direction, so the same integration value sits over an aligned perspective and over a frustrated one. An account that reads unity off the joint state gains a second ingredient it can name: the form of the coupling, which carries the direction the scalar discards. The quantity a high integration certifies is a real feature of the joint state, and it is a separate feature from the alignment of perspectives; naming the coupling locates what integration tracks rather than displacing it.

What binds many subsystems into one perspective is a coupling that is not itself one of the subsystems’ orientation-fixing acts. The projective read is a subsystem acting: it registers and records, and it does not bind. The evidential read binds, and it is not an act, in that it deposits no record, fires no outcome, and lives between the subsystems as a channel from one’s record to the other’s dynamics. The structure that supplies the acts, the records, and the orientations they carry fixes everything about the binding coupling except one thing: that the read deposits the record’s own state rather than a rotated one. That single premise, deposit, is where the form of the binding coupling is decided, and it concerns the coupling between subsystems rather than any subsystem’s own activity.

Several limits are worth stating. The operators are faithful to a model and are not derived from neuroscience. The perspective is a preferred basis in the structural sense, and nothing here claims that a shared perspective is or produces an experience. The no-alignment result and the binding read are established at the reads named, the projective and coherent reads on one side and the evidential transition on the other, and which further couplings bind is left open. The premises are stated per rank-one record, and a multi-component perspective is handled component by component. Deposit is stipulated as the definition of a faithful read.

The standing of deposit is best stated through what the record structure does and does not fix. A record fixes which outcome, with what weight, and in which state. The gain direction of a read that writes no record of its own is none of these, so nothing in the record structure constrains it, and an attempt to derive deposit from that structure should be expected to come up empty. Deposit is instead a property of the systems that run the coupling. A driven system persists by pumping free energy through a boundary and exporting entropy to its surroundings (Schrödinger, 1944; Friston, 2019; Kirchhoff et al., 2018), and it pays for the activity of its reads out of the throughput that keeps it in existence. Among the reads that reach a given record, the faithful one reaches it and stops, while every alternative overshoots or re-records at additional cost (Friston, 2010). Where that cost is a binding constraint, reads that waste throughput are selected against, so faithful reading is the regularity a persisting reader tends to exhibit rather than a law it obeys — the reading a self-evidencing system does (Hohwy, 2016), in keeping with the active-inference picture of a system that maintains the evidence for its own persistence.

This changes the kind of explanation deposit admits, from a bare stipulation to a selected regularity that comes with a reason, and it does so without deriving the operator identity and

without reducing the assumption count. Three gaps hold it there. First, the argument imports a generic premise from the thermodynamics of computation, that a channel’s dynamical activity is a genuine dissipative cost for a physical implementation (Landauer, 1961); the premise is relocated to the level of the persisting system, not avoided. Second, selection delivers near-faithful coupling, and the exact deposit form only as a strong-selection limit; away from that limit the coupling binds to a slightly rotated target and degrades smoothly rather than failing. Third, and most sharply, the dissipation difference between a faithful and an unfaithful read is, per read, well below the thermal scale, so its selectability rests on aggregation over a system’s history — the regime of dissipation-driven adaptation (England, 2013) — and the biological conditions under which such sub-thermal per-read differences accumulate into a selectable advantage are not established. The free-energy vocabulary here motivates the reading rather than carrying it: the mechanism runs on ordinary selection and stochastic-thermodynamic reasoning, and the account does not rest on treating free-energy minimization as a dynamical law.

This reframes the open problem the construction isolates. Deriving deposit from the record structure is a task the account gives a reason to abandon. What replaces it is a question about biology: whether there is a plausible regime in which selection fixes deposit, which turns on whether a read channel’s activity is a genuine metabolic cost for physical implementations, whether the sub-thermal per-read difference aggregates into a selection pressure above the drift barrier, and whether the strong-selection limit of that pressure delivers the exact deposit form, with near-faithful binding away from the limit.

9 Conclusion

A measurement-like coupling raises the shared information of two subsystems and leaves their perspectives free, and an integration scalar cannot tell an aligned perspective from a frustrated one. A belief-updating coupling aligns them, fixed uniquely by two premises about a single record: the read deposits the record’s own state and writes no record of its own. The shared perspective obeys a stiffness law that closes from the two-subsystem channel to the network mean, holds more tightly as the network grows dense, fragments into independent perspectives when the coupling graph disconnects, and forms only in two or more dimensions. The construction assumes only subsystems with perspectives, records, and a read; one instantiation supplies concrete operators, and the theorems stand without it. What draws many subsystems into one perspective is a coupling that is not itself one of their acts, and the open question is whether a persisting system’s dynamics select the single premise that fixes it.

A Proof of the no-alignment result

We prove Proposition 1. The reader is subsystem i ; it dephases about its own orientation axis $\hat{e}_i$ at rate r_1 from its local activity and about the recorded axis $\hat{e}_j$ at rate r_2 from a projective read of subsystem j ’s record. On the orientation-plane qubit each axis is a unit Bloch vector, and a projective read $A(\theta_j)$ generates pure dephasing about $\hat{e}_j$, since $A(\theta_j)$ is Hermitian and its dissipator $\mathcal{D}[A(\theta_j)]$ annihilates any state diagonal in the recorded basis and damps the off-diagonal coherence.

Steady Bloch vector. Dephasing about a unit axis $\hat{e}$ at rate r acts on the Bloch vector $\mathbf{n}$ as $\dot{\mathbf{n}} = -r(\mathbf{n} - (\mathbf{n} \cdot \hat{e})\hat{e})$, contracting the component of $\mathbf{n}$ transverse to $\hat{e}$. With two dephasing channels the stationary condition is

$$(r_1 + r_2) \mathbf{n} = r_1(\mathbf{n} \cdot \hat{e}_i) \hat{e}_i + r_2(\mathbf{n} \cdot \hat{e}_j) \hat{e}_j = M \mathbf{n}, \quad M = r_1 \hat{e}_i \hat{e}_i^\top + r_2 \hat{e}_j \hat{e}_j^\top. \quad (9)$$

The matrix M is a sum of two rank-one projectors with positive weights, so its largest eigenvalue satisfies $\lambda_{\max}(M) \leq r_1 + r_2$, with equality only when $\hat{e}_i$ and $\hat{e}_j$ are parallel. For $\hat{e}_i \neq \hat{e}_j$ the eigenvalue is strictly smaller, so (9) has $\mathbf{n} = 0$ as its only solution in the balanced case $r_1 = r_2$: the orientation washes out, and the reduced state has no principal axis. Away from balance the stationary state is anisotropic, and its principal axis is the top eigenvector of M . That eigenvector is fixed by M , which depends on θ_j only through $\hat{e}_j$; when the local rate dominates, $r_1 \gg r_2$, the top eigenvector is $\hat{e}_i$ to leading order, so the principal axis pins at the reader's own orientation θ_i .

Vanishing torque. The orientation free energy (5) at the steady ρ_i is $F = -\log \text{Tr}[A(\theta_i)\rho_i]$, and its gradient in θ_i depends on θ_j only through the steady axis of ρ_i . In the isotropic case $\rho_i \propto I$ and F is constant in θ_i , so the torque is zero. In the pinned case the steady axis is θ_i to leading order and carries no θ_j -dependence at that order, so $\partial F/\partial \theta_i$ is θ_j -independent and the orientation acquires no restoring force toward θ_j . In both regimes the projective read exerts zero alignment torque in the slow steady state. A coherent read built from the same projectors generates a unitary that commutes with the recorded-basis populations and likewise carries no preferred direction along the recorded axis, so the same conclusion holds.

The steady state of the projective coupling is degenerate: the record-diagonal subspace is stationary, so the global steady state is not unique. The reader's principal axis and hence the induced torque are nonetheless independent of the initial state, since the steady axis is fixed by M alone. This is what the numerics report, a torque variance of order 10^{-11} across product, correlated, anti-correlated, and maximally mixed initial states. $\square$

B The two-premise characterization

We prove Theorem 1. The channel is a GKSL generator $\mathcal{L} = \sum_k \mathcal{D}[L_k]$ on the orientation-plane qubit, with rates absorbed into the L_k . The record is a single orientation value θ_j ; write $|\psi\rangle = |\psi(\theta_j)\rangle$ and $|\psi_\perp\rangle = |\psi_\perp(\theta_j)\rangle$, and $A = |\psi\rangle\langle\psi|$.

Lemma 1 (Gain-range collapse). *If $\sum_k L_k L_k^\dagger = |\psi\rangle\langle\psi|$, then every $L_k = |\psi\rangle\langle b_k|$ for some $|b_k\rangle$, and $\sum_k \|b_k\|^2 = 1$.*

Proof. Each $L_k L_k^\dagger$ is positive semidefinite. If positive semidefinite operators sum to C , every summand has range in $\text{range}(C)$: for $|\chi\rangle \perp \text{range}(C)$, $0 = \langle\chi|C|\chi\rangle = \sum_k \|L_k^\dagger|\chi\rangle\|^2$, so $L_k^\dagger|\chi\rangle = 0$ for every k , hence $\text{range}(L_k) \subseteq \text{range}(C) = \text{span}\{|\psi\rangle\}$. Thus $L_k = |\psi\rangle\langle b_k|$, whence $L_k L_k^\dagger = \|b_k\|^2 |\psi\rangle\langle\psi|$, and summing gives $\sum_k \|b_k\|^2 = 1$. $\square$

Generator under deposit. With $L_k = |\psi\rangle\langle b_k|$, set $B = \sum_k L_k^\dagger L_k = \sum_k |b_k\rangle\langle b_k|$, a density operator by Lemma 1. Then $\sum_k L_k \rho L_k^\dagger = \text{Tr}(B\rho) |\psi\rangle\langle\psi|$ and $\sum_k L_k^\dagger L_k = B$, so the generator collapses to

$$\mathcal{L}[\rho] = \text{Tr}(B\rho) |\psi\rangle\langle\psi| - \frac{1}{2}\{B, \rho\}. \quad (10)$$

The generator is a functional of B and $|\psi\rangle$ alone; the Kraus decomposition of B into individual $|b_k\rangle$ is invisible to it, and that invisibility is the whole of the operator-set gauge freedom under deposit.

Darkness closes it. Darkness gives $0 = L_k |\psi\rangle = \langle b_k|\psi\rangle |\psi\rangle$, so $\langle b_k|\psi\rangle = 0$ and every $|b_k\rangle \perp |\psi\rangle$. On the qubit that forces $|b_k\rangle = c_k |\psi_\perp\rangle$, hence $L_k = c_k |\psi\rangle\langle\psi_\perp| = c_k T(\theta_j)$. Then $B = (\sum_k |c_k|^2) |\psi_\perp\rangle\langle\psi_\perp|$, and deposit's $\text{Tr} B = 1$ fixes $\sum_k |c_k|^2 = 1$, so (10) gives $\mathcal{L} = \mathcal{D}[T(\theta_j)]$. The residual freedom is the choice of the c_k with $\sum_k |c_k|^2 = 1$ and an overall rate, which leave the generator invariant or proportional.

Attractor and sign as consequences. From (10), the record $|\psi\rangle\langle\psi|$ is stationary iff $|\psi\rangle$ is an eigenvector of B . On the qubit B diagonalizes in $\{|\psi\rangle, |\psi_\perp\rangle\}$, so $B = (1 - \mu)|\psi\rangle\langle\psi| + \mu|\psi_\perp\rangle\langle\psi_\perp|$, and the populations obey $\dot{\rho}_{\perp\perp} = -\mu\rho_{\perp\perp}$ and $\dot{\rho}_{\psi\psi} = +\mu\rho_{\perp\perp}$. For $\mu > 0$ every off-record component decays and $\rho \rightarrow |\psi\rangle\langle\psi|$ from any initial state: the record is the unique attractor. For $T(\theta_j)$ itself, $\mu = 1$, so the record attracts, which fixes the ferromagnetic sign, since the read that instead pulled toward $|\psi_\perp\rangle$ would attract to the anti-record.

The two escape families. Each premise excludes one one-parameter family, and neither excludes both. The shear family $L_\varepsilon = |\psi(\theta_j + \varepsilon)\rangle\langle\psi_\perp|$ is dark for every ε , with attractor the record, but its gain is $L_\varepsilon L_\varepsilon^\dagger = A(\theta_j + \varepsilon) \neq A(\theta_j)$ for $\varepsilon \neq 0$, so deposit excludes it. The admixture family $\{\sqrt{1 - \mu}A(\theta_j), \sqrt{\mu}T(\theta_j)\}$ has gain $(1 - \mu)A + \mu A = A$ for every μ and so satisfies deposit, but for $\mu < 1$ its projector term does not annihilate $|\psi\rangle$, so darkness excludes it. The two families intersect only at $T(\theta_j)$. Attractor uniqueness holds across the whole shear family and the whole $\mu \in (0, 1]$ admixture family, so pairing it with deposit still leaves the admixture, and pairing it with darkness still leaves the shear; only deposit and darkness together pin the generator.

Per-record scope. Lemma 1 uses that $\text{range}(A)$ is one-dimensional. When the record is a rank-two projector $A = P_K + P_G$ on orthogonal sectors, $\text{range}(A)$ is two-dimensional, and deposit and darkness lifted to the lumped projector leave each L_k a map from the orthogonal complement into the two-dimensional record subspace. Deposit then fixes only $\sum_k M_k M_k^\dagger = I_2$, which for a single operator is any 2×2 unitary, so a cross-sector pairing or a relative phase between sectors survives while still binding ferromagnetically. Stated per rank-one record, deposit reads “the K -sector gain is P_K ” and “the G -sector gain is P_G ,” which excludes the cross-sector pairing and the phase mismatch, and the bound read is the direct sum $T(\theta_j^K) \oplus T(\theta_j^G)$. The theorem is the single-record theorem in each sector, and the components of a multi-component perspective are its second and further active records rather than inert spectator directions.

C The stiffness law

We derive Theorem 2. The orientations θ_i on the coupling graph G evolve by the nematic overdamped Langevin dynamics

$$d\theta_i = - \sum_j 2J_{ij} \sin(2(\theta_i - \theta_j)) dt + \sqrt{2D} dW_i, \quad (11)$$

with nonnegative couplings J_{ij} from the two-subsystem read, a preferred offset of zero on every edge, and independent noise of strength D . The π -periodicity of the interaction is the nematic character of the orientation, and the zero preferred offset is the ferromagnetic direction, which makes the interaction frustration-free on any graph.

Spin-wave expansion. Small fluctuations $\phi_i = \theta_i - \bar{\theta}$ about a shared orientation linearize $\sin(2(\theta_i - \theta_j)) \approx 2(\phi_i - \phi_j)$, so (11) becomes the Ornstein–Uhlenbeck process $d\phi_i = -4 \sum_j J_{ij}(\phi_i - \phi_j) dt + \sqrt{2D} dW_i = -4(L_G \phi)_i dt + \sqrt{2D} dW_i$, where $L_G = \text{diag}(J\mathbf{1}) - J$ is the weighted graph Laplacian. In the eigenbasis of L_G , with eigenvalues $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_N$ and normal modes $\tilde{\phi}_n$, each nonzero mode is an independent Ornstein–Uhlenbeck process with stationary variance $\langle \tilde{\phi}_n^2 \rangle = D/(4\lambda_n)$. The zero mode $\lambda_1 = 0$ is the global orientation rotation and carries the symmetry;

it is removed by measuring fluctuations about the mean. Summing the nonzero modes and dividing by N to convert to the per-site variance gives

$$\text{Var}(\theta_i - \bar{\theta}) = \frac{1}{N} \sum_{n \geq 2} \langle \tilde{\phi}_n^2 \rangle = \frac{D}{4} \frac{1}{N} \sum_{n \geq 2} \frac{1}{\lambda_n(L_G)}, \quad (12)$$

which is (6). The mean orientation $\bar{\theta}$ is the projection onto the zero mode and is unique modulo the global rotation, so the shared orientation exists and is unique modulo the orientation symmetry.

Gap-dominated regime. When L_G has a spectral gap, $\lambda_2 \ll \lambda_3$, the sum is dominated by its first term and $\text{Var} \propto 1/\lambda_2 = 1/\Phi_\lambda$, recovering (7). On graphs with a flat or heavy-tailed spectrum the full sum is the robust object and the single-mode approximation is not accurate.

Dimensional threshold. For a d -dimensional lattice the Laplacian spectrum near zero behaves as $\lambda(k) \sim k^2$, so the spectral sum is $\int d^d k / k^2$, which diverges at small k for $d \leq 2$ and converges for $d > 2$; the marginal continuous-symmetry case $d = 2$ diverges logarithmically in the thermodynamic limit, and $d = 1$ diverges as a power. A one-dimensional chain therefore has unbounded orientation variance and never locks, while a two-dimensional sheet on a finite network retains a shared orientation with bounded variance up to the logarithmic size dependence; this is the Mermin–Wagner dimensional threshold for a continuous symmetry (Mermin and Wagner, 1966). The shared perspective is a sheet rather than a line for this reason.

Fragmentation. When G disconnects into components, L_G acquires one zero eigenvalue per component, so $\lambda_2 \rightarrow 0$ and the corresponding $\langle \tilde{\phi}^2 \rangle \rightarrow \infty$: the relative orientation of the two components is unbound. Each component keeps its own zero mode and its own mean orientation, and the variance within a component, computed from that component’s own nonzero spectrum, stays finite and unchanged as the bridge weakens. The shared perspective fragments into independent per-component perspectives, which is the dissociation limit of §5.

D Reproducibility

Every number reported is computed on the operators of §2–§6 and reproduces from a public module. The orientation-plane primitives are $|\psi(\theta)\rangle$, $|\psi_\perp(\theta)\rangle$, the projector $A(\theta)$ of (2), and the transition $T(\theta)$ of (4). The two-subsystem reads are the symmetric record-exchange channel, with local jump operators $A(\theta_1) \otimes I$ and $I \otimes A(\theta_2)$ at rate $(1 - \alpha)\gamma$ and cross reads $T(\theta_2) \otimes A(\theta_2)$ and $A(\theta_1) \otimes T(\theta_1)$ at rate $\frac{1}{2}\alpha\gamma$; the projective control replaces each cross read’s reader-side transition by the projector; and the coherent control uses the projector-built Hamiltonian with the two local dissipators. The steady state is the exact slow steady state, the biorthogonal projection onto the null space of the vectorized Liouvillian, and the alignment torque is the θ -gradient of the free energy (5) at the reduced state.

Reference values. The projective and coherent reads give a steady alignment torque with variance of order 10^{-11} across a swept range of writer orientations and across product, correlated, anti-correlated, and maximally mixed initial states; the maximally integrated anti-correlated pair has integration $\Phi[\rho] = 5.545$ while preferring opposed orientations. The symmetric evidential read gives a torque variance of 0.9735, a drift toward the writer’s orientation at every writer orientation with a fixed point at coincident orientations, and a drift identical to four digits across the four initial

states. The reversed transition gives anti-alignment, and a global phase on the transition leaves the drift unchanged. The single-event weight equals $\cos^2(\theta_i - \theta_j)$ and its loss channel $\sin^2(\theta_i - \theta_j)$ to machine precision. The fitted nematic edge coupling $J(\alpha)$ grows with log-log slope near one in α .

Stiffness closure. The stiffness law (6) is compared to a direct simulation of (11) with the fitted $J(\alpha)$ on the edges, at noise $D = 0.02$. The ratio of measured to predicted variance is within one to four percent on a ring with next-nearest neighbors, a two-dimensional torus, and a honeycomb lattice, and within one to five percent on a triangular lattice with dense odd cycles and on the complete graph on twelve nodes, five topologies in all. The proportionality $\text{Var} \propto 1/\Phi_\lambda$ holds across couplings with a coefficient of variation of six percent. In the fragmentation test, weakening the bridge between two clusters holds the intra-cluster variance constant while the inter-cluster orientation variance diverges as the bridge coupling vanishes. The triangular lattice, the topology most apt to expose frustration, shows none.

Module. The primitives, the two-subsystem reads and their controls, the steady-state solver, the alignment torque, the nematic coupling fit, the graph Laplacian, the stiffness variance, and the connected-component test are provided as a single reference module, and the binding and dissociation panels are generated from it. The figures reproduced here are produced by the self-contained companion notebook `binding-demos.ipynb`.

References

- Acebrón, J. A., Bonilla, L. L., Pérez Vicente, C. J., Ritort, F., and Spigler, R. (2005). The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics*, 77(1):137–185.
- Born, M. (1926). Zur quantenmechanik der stoßvorgänge. *Zeitschrift für Physik*, 37(12):863–867.
- Breuer, H.-P. and Petruccione, F. (2002). *The Theory of Open Quantum Systems*. Oxford University Press.
- Busch, P. (2003). Quantum states and generalized observables: a simple proof of Gleason’s theorem. *Physical Review Letters*, 91(12):120403.
- Chalmers, D. J. (1996). *The Conscious Mind: In Search of a Fundamental Theory*. Oxford University Press.
- England, J. L. (2013). Statistical physics of self-replication. *The Journal of Chemical Physics*, 139(12):121923.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138.
- Friston, K. (2019). A free energy principle for a particular physics.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., and Pezzulo, G. (2017). Active inference: a process theory. *Neural Computation*, 29(1):1–49.
- Gleason, A. M. (1957). Measures on the closed subspaces of a Hilbert space. *Journal of Mathematics and Mechanics*, 6(6):885–893.

- Goff, P. (2017). *Consciousness and Fundamental Reality*. Oxford University Press, New York.
- Gorini, V., Kossakowski, A., and Sudarshan, E. C. G. (1976). Completely positive dynamical semigroups of N -level systems. *Journal of Mathematical Physics*, 17:821–825.
- Hohwy, J. (2016). The self-evidencing brain. *Noûs*, 50(2):259–285.
- James, W. (1890). *The Principles of Psychology*. Henry Holt and Company, New York.
- Kirchhoff, M., Parr, T., Palacios, E., Friston, K., and Kiverstein, J. (2018). The Markov blankets of life: autonomy, active inference and the free energy principle. *Journal of the Royal Society Interface*, 15(138):20170792.
- Kuramoto, Y. (1984). *Chemical Oscillations, Waves, and Turbulence*. Springer, Berlin.
- Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3):183–191.
- Lindblad, G. (1976). On the generators of quantum dynamical semigroups. *Communications in Mathematical Physics*, 48(2):119–130.
- Masanés, L., Galley, T. D., and Müller, M. P. (2019). The measurement postulates of quantum mechanics are operationally redundant. *Nature Communications*, 10:1361.
- Mermin, N. D. and Wagner, H. (1966). Absence of ferromagnetism or antiferromagnetism in one- or two-dimensional isotropic Heisenberg models. *Physical Review Letters*, 17:1133–1136.
- Oizumi, M., Albantakis, L., and Tononi, G. (2014). From the phenomenology to the mechanisms of consciousness: Integrated Information Theory 3.0. *PLoS Computational Biology*, 10(5):e1003588.
- Parr, T., Pezzulo, G., and Friston, K. J. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press, Cambridge, MA.
- Rudrauf, D., Bennequin, D., Granic, I., Landini, G., Friston, K., and Williford, K. (2017). A mathematical model of embodied consciousness. *Journal of Theoretical Biology*, 428:106–131.
- Schrödinger, E. (1944). *What Is Life? The Physical Aspect of the Living Cell*. Cambridge University Press.
- Seager, W. (1995). Consciousness, information and panpsychism. *Journal of Consciousness Studies*, 2(3):272–288.
- Sperry, R. W. (1968). Hemisphere deconnection and unity in conscious awareness. *American Psychologist*, 23(10):723–733.
- Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5:42.
- Tononi, G., Boly, M., Massimini, M., and Koch, C. (2016). Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience*, 17(7):450–461.
- Wallace, D. (2012). *The Emergent Multiverse: Quantum Theory According to the Everett Interpretation*. Oxford University Press, Oxford.
- Zurek, W. H. (2003). Decoherence, einselection, and the quantum origins of the classical. *Reviews of Modern Physics*, 75(3):715–775.