Supplemental Material: When does cortical coherence carry shared content?

Overview

This supplement follows the article from its three requirements through the limits of their observables to the cortical hypothesis and supporting physical models. The core constructions come first: stationary coherence and installed energy, compatible contents, and internal reconstruction. The thermodynamic models describe specified implementations with different state spaces and physical assumptions.

The following routes connect the article to the detailed results:

- **Core constructions:** the phase model in Section 1; microscopic installation and its energy bound in Section 2; gluing, content updates and the phase-encoder bound in Section 3; and reconstruction in Section 5. Table S1 maps their formal statements and the supporting models to Lean declarations.
- **Observations and controls:** finite phase dynamics in Section 7, recovery and mechanism identification in Section 8, plasticity in Section 9, content-agreement estimation in Section 10, and EEG methods and spatial measurement design in Section 11.
- **Physical realization:** cortical field calibration and spatial controls in Section 12, followed by the support comparisons and finite-region embedding in Section 13.
- **Agency and resources:** the boundary result in Section 14 and the erasure bound in Section 15; information and heat balances in Section 16, policy selection and learning in Section 17, and repeated operation, sensor memory and finite supply in Section 18. Section 19 gives the additional identifications used by the conditional composition.
- **Scope:** Section 20 states what the formal endpoint is and is not, Section 21 places it against neighbouring accounts, and Section 22 gives the shape a physical postulate must take here and why.

Witnesses demonstrate their hypotheses on the named models; they need not share a physical implementation. Source code and saved numerical summaries are maintained at `https://github.com/bobaseb/coupling_gated_cortical_coherence`.

1 Macroscopic scaling and Kuramoto phase synchronization

We posit a continuous macroscopic coupling description. We model the local phases θ_i as coupled nonlinear oscillators via the Kuramoto model.

Scope of the continuum description. `mesh_refinement_convergence` proves convergence of one scalar coupling energy. It produces neither a continuum kernel nor field dynamics: `edge_region` has no $M \times M$ product structure from which such a kernel could be recovered, and the vertex-supported candidate is invisible to the continuum functional on an atomless substrate. Both obstructions are to this architecture: `KernelMesh` carries cell pairs of positive mass and its limit is a kernel integral, at the price of declaring the hardware that sets the weights. The symmetric kernel is therefore a modelling input either way. Its identification with the brain's endogenous *electromagnetic* field is a further empirical commitment about cortex, argued in the main text from the ephaptic-coupling literature with explicit calibration and observational limitations. Nothing in this section mathematically establishes either step.

For a system with coupling strength K and thermodynamic noise D, the dynamic evolution is:

$$\dot{\theta}_i = \omega_i + \frac{K}{N} \sum_{j=1}^{N} \sin(\theta_j - \theta_i) + \eta_i(t) \tag{1}$$

Theorem S1 (Stationary coherent branch): For identical frequencies in a rotating frame with positive diffusion, the positive normalized stationary densities of the mean-field Fokker–Planck equation are the von Mises family, and the resulting self-consistency equation has exactly one positive solution when $K > 2D$ and none when $K \leq 2D$. Linearization about the uniform density has a growing mode exactly above that threshold, and the coherent branch carries a strictly dissipative quadratic form transverse to its rotation orbit. This does not assert dynamical selection or noiseless phase locking in the stochastic model.

Lean 4 Implementation (`Phase4_KuramotoDynamics.lean`): We formalize the Lyapunov potential $V(\theta) = -\frac{1}{2}\sum_{i,j} A_{ij}\cos(\theta_j - \theta_i)$ of the coupled field and prove two directions: `phase_locked_minimizes_potential` (the phase-locked configuration attains the global minimum whenever all couplings are positive) and `potential_min_implies_phase_locked` (any minimizer is phase-locked). Note the scope: this is a statement about the *static* potential landscape and holds for all positive coupling matrices. The critical threshold $K_c = 2D$ does not appear in this file. It is formalized separately, in `Phase8_SelfConsistency.lean` (see below), as a statement about the mean-field self-consistency equation, and the density that equation uses is classified in `Phase8_FokkerPlanck.lean` from the stationary Fokker–Planck equation itself. What no file supplies is the passage from the stochastic differential equation to that equation, or a solution theory for its evolution. The numerical controls in Section 7 distinguish identical-frequency diffusion from the separate quenched-frequency model illustrated by `simulations/kuramoto.py`.

Positivity, and the regime these theorems cover. The hypothesis "all couplings positive" is not a technical convenience and it is not free. `KuramotoSystem` carries $A : V \to V \to \mathbb{R}$, symmetric and otherwise unconstrained, so signed — frustrated — couplings are representable; but every result in this development that identifies *where* the dynamics ends assumes positivity. `phase_locked_minimizes_potential` and `potential_min_iff_phase_locked` assume $A_{ij} > 0$; the class field `ThermodynamicCover.A_pos` demands it of every instance, and so of every conclusion of Theorem S2; and `kuramoto_tendsto_global_minimum` requires the strictly stronger uniform bound $0 < a \leq A_{ij}$, since a enters the Łojasiewicz constant. Such a network is an unfrustrated ferromagnet with one minimum up to a global phase shift — exactly not the near-degenerate landscape a memory argument from geometric frustration requires.

By contrast, nothing about the sign of A is used by `is_kuramoto_trajectory_exists` and its uniqueness counterpart, by the Lyapunov identity `dV_dt_le_zero` on which the whole descent argument rests, by `dynamic_potential_antitone` and `dynamic_potential_tendsto`, or by `velocity_sq_tendsto_zero`. Since positivity is a *field* of `ThermodynamicCover` rather than a hypothesis distributed across the development, it enters the chain at one node and not before it. So *that* the motion stops holds for frustrated systems as well; *where* it stops does not. On the continuum side `phase_locked_minimizes_entropy_of_symm` assumes only symmetry and so admits signed kernels, but it takes the existence of a dynamically phase-locked field as a hypothesis, which is precisely what frustration threatens. The main text scopes the frustration claim out of the chain accordingly, rather than repeating it beside theorems that exclude it.

1.1 Stationary self-consistency

Lean 4 Implementation (`Phase8_SelfConsistency.lean`): The self-consistency equation behind $K_c = 2D$ is formalized as real analysis. With the von Mises stationary density $\rho(\theta) \propto e^{a\cos\theta}$, $a = Kr/D$, supplied as an input to this file and classified from the stationary equation in `Phase8_FokkerPlanck.lean` below (`simulations/fokker_planck_von_mises.py` checks the same algebra numerically), `besselRatio` is $I_1(a)/I_0(a)$ written as a ratio of interval integrals, Mathlib having no modified Bessel functions. `vonMisesM_eq_mul_vonMisesS` is the integration-by-parts identity $\int_{-\pi}^{\pi}\cos\theta\, e^{a\cos\theta} = a\int_{-\pi}^{\pi}\sin^2\theta\, e^{a\cos\theta}$, exact because $\sin(\pm\pi) = 0$ kills the boundary term. `vonMisesC2_nonneg` proves $I_2(a) \geq 0$ by folding $[-\pi,\pi]$ onto $[0,\pi/4]$ through the reflections $\theta \mapsto \pi - \theta$ and $\theta \mapsto \pi/2 - \theta$, where the integrand becomes a product of two manifestly non-negative factors; `vonMisesS_le_half_vonMisesZ` restates it as $\mathbb{E}_a[\sin^2\theta] \leq 1/2$ and `besselRatio_le_half_self` as $R(a) \leq a/2$. That factor of $1/2$ is exactly what separates a threshold at D from one at $2D$; `subcritical_fixed_point_eq_zero` is the crude form obtained from $\sin^2 \leq 1$ and is kept to make the point visible.

`subcritical_fixed_point_eq_zero′` is the subcritical theorem. For $D > 0$ and $0 \leq K < 2D$, the only $r \geq 0$ with $r = R(K, r)$ is $r = 0$.

The supercritical theorem is `supercritical_fixed_point_exists`: for $K >$ `critical_coupling` D there is an r with $0 < r \leq 1$ and $r = R(K, r)$. Its proof needs no series expansion. Factoring $R(a) = a\,\mathcal{S}(a)$

with $\mathcal{S}(a) = \mathbb{E}_a[\sin^2\theta]$ — the factorization the integration by parts supplies — it suffices that $\mathcal{S}$ is continuous with $\mathcal{S}(0) = 1/2$: since $D/K < 1/2$, the map overshoots for small $r > 0$, while `besselRatio_le_one` (R is a mean of $\cos\theta$) makes it undershoot at $r = 1$, and `intermediate_value_Icc'` supplies the crossing. Continuity of $\mathcal{S}$ is Mathlib's continuity of a parametric interval integral, applied to the jointly continuous integrand. `critical_coupling_is_threshold` states the two halves together, and §9 of the file discharges the hypotheses of each at concrete values ($K = 3$, $D = 1$ above threshold; $K = 3$, $D = 2$ below it), so neither is vacuous.

Section 8 of the file identifies the equation with the physics it is named for. `vonMisesDensity` is the normalized weight, proved a probability density on $[-\pi, \pi]$ (`vonMisesDensity_integral_eq_one`); `circularOrderParameter` is the average of $e^{i\theta}$ against a density, the continuum analogue of `order_parameter_complex` of `Phase4`; `circularOrderParameter_vonMises` proves that average equals $R(a)$, real because `vonMises_mean_sin` gives $\mathbb{E}_a[\sin\theta] = 0$; and `fixedPoint_iff_selfReproducing` states the equivalence: r solves $r = R(K, r)$ exactly when the density it induces has order parameter r. `incoherent_density_uniform` identifies the solution $r = 0$ with the uniform density, and `supercritical_coherent_density` restates the existence theorem as the existence of a stationary density with positive order parameter.

Strict inequalities determine the threshold case. `vonMisesC2_pos` upgrades the fold argument to a strict inequality for $a > 0$ — `foldB` is continuous, non-negative on $[0, \pi/4]$ and strictly positive on its interior, so `intervalIntegral_pos_of_pos_on` applies — giving $\mathbb{E}_a[\sin^2\theta] < 1/2$ and $R(a) < a/2$ for every $a > 0$. That closes the threshold case: `fixed_point_eq_zero_of_le_critical` covers all $K \le 2D$, so the incoherent and coherent regimes exhaust $K \ge 0$ with no gap, and $K = 2D$ falls on the incoherent side.

`coherent_iff_sRatio_eq` reformulates the coherent branch: for $r \neq 0$, solving $r = R(K, r)$ is equivalent to $\mathcal{S}(Kr/D) = D/K$, so the branch is a level set of $\mathcal{S}$. `coherent_branch_continuous_at_threshold` then proves that for every $\varepsilon > 0$ there is a $\delta > 0$ such that on $(2D, 2D + \delta)$ every coherent solution satisfies $r < \varepsilon$ — the branch emerges continuously from zero, excluding a discontinuous jump. Its proof needs no monotonicity of $\mathcal{S}$: on a compact interval bounded away from the origin the extreme value theorem gives a maximum strictly below $1/2$, and beyond that the crude bound $\mathcal{S}(a) \le 1/a$ suffices.

Strict decrease of $\mathcal{S}$ and strict increase of R give uniqueness and monotonicity of the coherent branch. Section 1.6 presents the folding and crossing argument. The consequences are `coherent_fixed_point_unique` (at most one positive solution), `supercritical_fixed_point_existsUnique` and `supercritical_solution_set` (the non-negative solutions are exactly $\{0\}$ at or below threshold and $\{0, r\}$ above it), and `coherent_branch_strictMono`, which uses both halves: $\mathcal{S}$ decreasing puts the concentrations of two couplings in the order $a_1 < a_2$ through $\mathcal{S}(a_i) = D/K_i$, and R increasing carries that order to $r_1 < r_2$.

Scope: the density is an input to this file rather than an output of it, and the classification that derives it is a separate module; nothing in this file differentiates anything, so $\mathcal{S}$ and R are reached as strictly monotone rather than as smooth, and the rate at which the branch leaves zero is the business of `Phase8_CriticalExponent.lean` below; dynamical selection is not established, and could not be here, since no theorem in the file mentions a trajectory; and the order parameter reached is that of a density, not of a trajectory — no theorem relates `circularOrderParameter` to `order_parameter_complex`, which would be the mean-field limit.

1.2 Stationary onset rate

Lean 4 Implementation (`Phase8_CriticalExponent.lean`): The stationary onset rate follows from a second-order expansion of the integral ratio. `vmMoment` n a is the cosine moment $\int_{-\pi}^{\pi} \cos^n\theta\, e^{a\cos\theta}$, and `hasDerivAt_vmMoment` differentiates it in a under the integral sign, raising the cosine power by one; the interchange is Mathlib's dominated-convergence form for interval integrals, dominated by the constant $e^{|a|+1}$ on the unit ball around the parameter, which is where the compactness of $[-\pi, \pi]$ and the boundedness of cos do their work. Iterating gives `contDiff_two_vmMoment`, and `vonMisesSRatio_eq_moments` rewrites $\mathcal{S}(a) = 1 - $ `vmMoment` $2\,a/$`vmMoment` $0\,a$, so $\mathcal{S}$ inherits two derivatives from a quotient whose denominator is positive.

The moments at $a = 0$ are elementary integrals of powers of cos: 2π, 0, π, 0 and $3\pi/4$. They give $\mathcal{S}(0) = 1/2$, $\mathcal{S}'(0) = 0$ — the first derivative vanishes by the symmetry that kills the odd moments — and $\mathcal{S}''(0) = -1/8$, and Taylor's theorem in Mathlib's `taylor_isLittleO_univ` form then yields `vonMisesSRatio_second_order`: $\mathcal{S}(a) - (\frac{1}{2} - a^2/16)$ is $o(a^2)$ at the origin. The remainder is $o(a^2)$ and not $O(a^4)$; the stronger form would need a fourth derivative and nothing here requires it.

`coherent_solution_critical_exponent` then combines the expansion with `coherent_iff_sRatio_eq`: writing $a = Kr/D$ and rearranging $\mathcal{S}(a) = D/K$ gives $K - 2D = Ka^2(\frac{1}{8} - 2q(a))$ with q the normalized remainder, and letting $K \to 2D$ along any family of positive coherent solutions with $r \to 0$ sends $q \to 0$ and $r^2/(K - 2D) \to 1/D$. The statement is deliberately filter-indexed rather than parameterized by a curve, so it constrains every such family and privileges no path.

Because its hypotheses hold only eventually along the filter, and an eventual clause is vacuously true on the empty one, the file also exhibits a family satisfying it: `coherentBranch` selects the unique positive solution above threshold via `supercritical_fixed_point_existsUnique`, `coherentBranch_tendsto_zero` reads `coherent_branch_continuous_at_threshold` as a limit along $K \downarrow 2D$, and `coherentBranch_critical_exponent` discharges every hypothesis on it.

Scope is inherited unchanged from the file above: this is asymptotics of the solutions of a stationary equation. It establishes no stability, no dynamical selection and no time course, so reading it as a square-root foot on a recovery curve requires the adiabatic separation and a threshold crossing at non-zero speed, both of which the main text states as assumptions.

1.3 The stationary equation and linear stability

Lean 4 Implementation (`PhaseS_FokkerPlanck.lean` and the five modules that build on it, named below): The equation these six files work in is the classical stationary one on the circle. `current` is $J = v\rho - D\,\partial_\theta \rho$ and `operator` is $-\partial_\theta J$; `IsStationary` asks for a differentiable, 2π-periodic, strictly positive and normalized density whose current is itself differentiable and whose operator vanishes. Differentiability belongs in the predicate rather than beside it: `deriv` is totalized, so an equation written with it alone is satisfied by functions that have no derivative where the equation is being asserted.

The density is classified rather than assumed. `integratingFactor_hasDerivAt` is the identity $\partial_\theta(\rho e^{V/D}) = -(J/D)e^{V/D}$ for a gradient drift $v = -\partial_\theta V$. Stationarity makes J constant (`stationary_current_const`), periodicity of $\rho e^{V/D}$ makes its integral over a period vanish, and $\int_{-\pi}^{\pi} e^{V/D} > 0$ then forces that constant to zero: `stationary_current_zero`. The remaining equation is first order, giving `stationary_gibbs`, $\rho = ce^{-V/D}$ with $c > 0$, and with the cosine potential and normalization, `stationary_iff_vonMises`: for the drift $-Kr\sin\theta$ the positive normalized stationary densities are exactly `vonMisesDensity`(Kr/D). `supercritical_stationary_existsUnique` transfers the scalar bifurcation theorem to this statement — above threshold exactly one positive r carries a stationary density whose own circular order parameter is r. No modified Bessel function, spectral theory or operator closure enters; the argument is the one `vonMisesZ_pos` already used, run once more.

The linearization is derived from the equation. `meanDrift` writes the mean-field coupling through the density's two trigonometric moments and `meanDrift_eq_integral` identifies it with $-K\int \sin(\theta - \varphi)\rho(\varphi)\,d\varphi$. `current_expansion` expands the nonlinear current at $q + \varepsilon u$ exactly, as the background current plus ε times `linearCurrent` plus an explicit quadratic term, so `linearOperator` is the derivative of the equation rather than a separately declared list of rates; `selfConsistent_stationary` confirms that a fixed point is stationary for the drift its own density generates.

The incoherent state. On the uniform density the coupling reaches only the first harmonic. `incoherent_mode_rate` gives, for both real Fourier families, rate $K/2 - D$ at $n = 1$ and $-Dn^2$ for $n \geq 2$, and `incoherent_instability_iff` concludes that a mode of positive index grows exactly when $K > 2D$. The constant mode is excluded because a density perturbation carries zero mass. This is a per-mode computation and claims no completeness of the Fourier family in any function space.

The coherent branch. `branchCoupling` $= D/\mathcal{S}(a)$ parametrizes the supercritical branch by concentration; `branchCoupling_supercritical` places it above $2D$ and `branch_stationary` makes its density stationary for its own mean field. `generator_eq_neg_linearCurrent` identifies the primitive generator with minus that linearized current, `generator_symmetric` makes it symmetric in the weighted primitive pairing by integration by parts, and `generator_energy` evaluates its diagonal as $-D$ times the coherent quadratic form. That form is strictly positive away from rotation because of a strict moment inequality, and the inequality is derived rather than posited: differentiating the defining integrals gives the Riccati identity `sRatio_riccati`, which rules out a vanishing derivative at positive concentration (`sRatioDeriv_neg`) and yields `cosine_variance_lt_sine`. With the weighted Poincaré estimate `transverse_poincare`, whose constant is explicit and finite, the result is `coherent_linear_stability`: on regular periodic perturbations of zero mass, orthogonal to rotation in the weighted primitive metric, the form is at most $-$`linearStabilityRate` times the squared norm, and that rate is positive (`linearStabilityRate_pos`). `rotation_generator_zero` exhibits the rotation direction as a nonzero element of zero rate, so the transversality is necessary rather than a convenience: a gap on the whole

space would be false, and stability holds modulo the collective phase.

Two functionals that are not the same functional. `currentDissipation` integrates $J^2/(D\rho)$ over phase space. `stationary_current_separation` evaluates it at the coherent stationary state, where it is zero, against the phase-averaged squared drift at the same state, which is positive. Both are functionals on phase space, so the difference is not a difference of domains; and neither is `sigmaContinuum`, which integrates squared drift over substrate sites at fixed phases and is the object of the plasticity results. No identification among the three is proved or used.

Consumers and witnesses. `IsEMFieldCoupling.coherent_phase_model` carries the stationary state and the gap predicate `HasCoherentLinearGap` to the identified field model of Phase 9, at that model's own coupling and diffusion; it transfers nothing to a heterogeneous spatial kernel. `Examples/Phase8.lean` §21 fixes $D = 1$ and concentration 1, checks that the stationary density is not the uniform one and that the stability rate is not zero, and runs the incoherent criterion at $K = 3$ and $K = 1$ on either side of the threshold.

Scope. These are linear operator and energy estimates for the limiting equation. They construct no solution of the evolution equation, no semigroup and no self-adjoint closure; they establish no nonlinear or orbital stability and no dynamical selection; and the passage from the stochastic differential equation to the Fokker–Planck equation is not formalized anywhere in the development. Every declaration of the eight stationary modules named in this subsection and the two above it depends only on `propext`, `Classical.choice` and `Quot.sound`.

The gap between these stationary-density results and a finite trajectory requires a dynamical mean-field limit with appropriate initial laws and time horizon. General interacting-diffusion and Kuramoto results give relevant settings (Sznitman, 1991; Dai Pra & den Hollander, 1996; Bertini et al., 2010); no theorem in this development establishes the particular limit needed for the finite controls below.

1.4 Order and the cost of probability current

For a smooth positive normalized periodic phase density ρ satisfying $\partial_t \rho = -\partial_\theta J$, with diffusion $D > 0$, define $\sigma_J = \int_{-\pi}^{\pi} J^2/(D\rho)\, d\theta$ and $r = |\int e^{i\theta} \rho\, d\theta|$. Differentiation under the integral gives $\dot{r}^2 \le D\sigma_J(1-r^2)$ wherever $0 < r < 1$. A smooth path staying in that interval therefore obeys

$$\int_0^\tau \sigma_J\, dt \ge \frac{[\arcsin r(\tau) - \arcsin r(0)]^2}{D\tau}. \tag{2}$$

Endpoint order, duration, diffusion and the probability current determine this floor. In an exploratory finite-volume onset calculation, a threshold chosen after inspecting pilot runs gives current cost 0.120 versus a floor of 0.046. These numbers describe the discretized model, not thermal or metabolic power.

The current $J = v\rho - D\partial_\theta \rho$ also satisfies

$$\left(\int_{-\pi}^{\pi} J\, d\theta\right)^2 \le D\sigma_J. \tag{3}$$

A uniform density under constant drive ω attains $\sigma_J = \omega^2/D$. The identical-frequency gradient model has positive self-consistent stationary order above $K = 2D$ with $J = 0$ and $\sigma_J = 0$. Stationary order therefore has no positive current-cost floor. A cortical power claim needs the physical current, reservoir and thermal conversion to be identified. This current functional is distinct from the substrate squared-drift objective used in the plasticity controls.

The current functional in Eq. (3) has a sharp bound that requires no stationarity assumption. Write $m = \int_{-\pi}^{\pi} J$. For continuous J, continuous positive ρ and $\int \rho = 1$,

$$0 \le \int_{-\pi}^{\pi} \frac{(J - m\rho)^2}{\rho} = \int_{-\pi}^{\pi} \frac{J^2}{\rho} - m^2 = D\sigma_J - m^2.$$

Strict positivity on this compact interval ensures that the reciprocal density and the displayed integrands are integrable. The Lean theorem `current_integral_sq_le` proves this expansion using the density's own Fokker–Planck current. The equality witness has $\rho = 1/(2\pi)$ and $v = \omega$, hence $J = \omega/(2\pi)$ and $\sigma_J = \omega^2/D$; `uniform_driven_stationary` and `uniform_current_dissipation` verify both the stationary equation and the cost. The concrete case $D = 1$, $\omega = 2$ has cost four and zero phase order.

Conversely, `supercritical_zero_dissipation` constructs, for every $D > 0$, $K > 2D$, a positive self-consistent stationary order and its density with zero current dissipation. `no_positive_order_`

`only_bound` therefore rules out every universal bound strictly positive at all positive orders. Keeping this density constant in time also refutes such a bound on time averages. This concerns the identical-frequency rotating-frame gradient model; it does not impose detailed balance on cortex. A positive cortical power bound needs an additional drive or current constraint and a physical thermal scale. For a bounded continuous test function g with $g^2 \le 1$, `weighted_current_sq_le` proves $(\int gJ)^2 \le D\sigma_J$ by expanding $\int [J - (\int gJ)g\rho]^2/\rho$ and using $\int g^2\rho \le 1$. For a differentiable periodic current with continuous derivative, integration by parts gives

$$\int_{-\pi}^{\pi} \cos\theta\,(-\partial_\theta J) = -\int_{-\pi}^{\pi} \sin\theta\, J, \qquad \left(\int \cos\theta\,(-\partial_\theta J)\right)^2 \le D\sigma_J.$$

This is `cosine_rate_sq_le_dissipation`. It bounds the rate of the cosine moment for evolving densities when time differentiation under the integral is valid; it constructs no evolution. The uniform density with imposed drift $v = \sin\theta$, $D = 1$ has cosine rate $-1/2$ and dissipation $1/2$, providing a nonzero-rate checked example.

For the rotation-invariant order magnitude $r(t) = |\int e^{i\theta}\rho(t,\theta)\,d\theta|$, let the density and current be positive, normalized and periodic in angle, continuously differentiable as required by the classical continuity equation, and let differentiation of the cosine and sine moments under the integral be justified. At times with $0 < r < 1$, choose the angular direction of the first harmonic. Integration by parts writes $\dot r$ as the current weighted by the perpendicular sine; weighted Cauchy–Schwarz and $\int \sin^2(\theta - \phi)\rho \le 1 - r^2$ give

$$\dot r^2 \le D\sigma_J(1 - r^2).$$

The Lean results `orderMagnitude_eq_norm_circular`, `classical_moment_rates` and `order_magnitude_speed_bound` establish these steps for the stated classical functions. If r and the current cost are smooth on a finite interval and $0 < r < 1$ throughout, the chain rule for $\arcsin r$ and temporal Cauchy–Schwarz give Eq. (2), which is `density_order_endpoint_cost_bound`; it does not assert an absolutely continuous extension through zero order, construct the evolving density, or convert current cost to heat. Strictly positive continuous densities have $r < 1$, but the theorem keeps that domain explicit.

The numerical check evolves a positive periodic density by a conservative finite-volume flux and uses that same edge flux in the current-cost quadrature. Equal-endpoint early, linear and late coupling schedules, an exact uniform zero-current path and ordered diffusion relaxation are saved separately from the artifact-only report. Diffusion, duration, grid resolution and a two-frequency mixture are varied. The report interpolates the first crossing of the exploratory $r = 0.3$ threshold and compares accumulated current cost up to that crossing against the endpoint floor for its elapsed time. The crossing's measured cost exceeds the floor at every resolved grid; its slack remains positive as the grid is refined. Each frequency cohort has its own continuity equation. Their marginal has the aggregate current to which the theorem applies, while the mean cohort cost is a distinct, larger account. Frequency spread removes the identical-frequency drift premise of the homogeneous simulation, not the aggregate continuity equation. The threshold was chosen after inspection of an exploratory grid, so this is a numerical illustration rather than a prospective cortical onset test or a continuum error certificate.

A coupling parameter is a different observable. `arbitrary_coupling_speed_zero_cost` uses $K(t) = e^{at}$, so $K > 0$ and $K'(0) = a$ for any real a, while the uniform self-consistent density has zero current dissipation at every time. `uniform_time_dependent_solution` checks the nonautonomous continuity equation for this constant density under any coupling schedule. Thus an unrestricted bound on coupling speed from phase-current cost alone is false. These examples do not rule out a bound with positive order, a calibrated actuator model or a measured relation between coupling and the changing density.

1.5 Deterministic trajectories and their limits

Lean 4 Implementation (`Phase4_RotatingFrame.lean`): The full potential includes the natural frequency term $-\sum_i \omega_i\theta_i$ that appears in the Lyapunov descent result `dV_dt_le_zero` in the combinatorial thermodynamics module. With any $\omega_i \neq 0$ the full potential is unbounded below (`kuramoto_potential_unbounded_below`), so it has no minimum for a phase-locked state to attain, so the reduced and full potentials must be distinguished.

This file chains them by the standard rotating-frame reduction $\theta_i \mapsto \theta_i - \Omega t$. For a system of identical natural frequencies $\omega \equiv \Omega$: `is_kuramoto_trajectory_rotate` sends a trajectory to a trajectory of the zero-frequency system `sys.reduced` — the same system Theorem S2's `ThermodynamicCover` works with;

`kuramoto_potential_reduced` shows the two potentials coincide on it; `dynamic_potential_descent` transports the descent identity $\dot{V} = -\sum_i \dot{\theta}_i^2$ onto the potential above; and `rotating_frame_chain` runs this through the minimum characterization to phase-locking and $r^2 = 1$. Phase differences, phase-locking and the order-parameter magnitude are proved frame-invariant (`rotate_sub`, `is_phase_locked_rotate`, `order_parameter_r_sq_shift`), so no observable is lost. `Examples` §7 runs the chain end to end on a two-oscillator system, discharging every hypothesis, and exhibits a system whose full potential has no minimum.

Scope: the reduction is exact only for identical frequencies. A spread of frequencies leaves residual detunings $\omega_i - \Omega$ in the reduced system, to which the unboundedness result still applies, and phase-locking then depends on K_c, which this file does not treat — see the `Phase8_SelfConsistency.lean` note above for what is proved about the threshold, and note that it concerns a stationary density rather than the trajectories this file transports.

Trajectories that run. The deterministic initial-value problem admits a unique global solution through each initial state. `kuramotoField_lipschitz` (`Phase3_CombinatorialThermodynamics.lean`) proves the Kuramoto vector field globally Lipschitz with constant $2\sum_{ij} |A_{ij}|$ — sin is 1-Lipschitz and the index set is finite, so there is no smallness caveat and the natural frequencies drop out. `is_kuramoto_trajectory_unique` turns this into Picard–Lindelöf uniqueness on all of $\mathbb{R}$: two solutions agreeing at one instant agree forever, in both time directions.

Existence follows by constructing solutions on arbitrary finite windows and gluing them using uniqueness. Mathlib's `IsPicardLindelof` caps a local solution's lifetime by the radius of the ball on which the field's bounds hold, divided by the field's sup norm there, and carries no global-in-time theorem; but `kuramotoField_norm_le` bounds this field by $\sum_i |\omega_i| + \sum_{ij} |A_{ij}|$ on the *whole* state space, so the cap can be defeated by enlarging the ball after choosing the window, and `kuramoto_exists_on_window` produces a solution on every $(t_0 - T, t_0 + T)$. Uniqueness on an interval makes the family coherent — solutions on nested windows agree where both are defined — so $\theta(t) := \alpha_{\lceil |t-t_0| \rceil}(t)$ is a solution on all of $\mathbb{R}$ (`is_kuramoto_trajectory_exists`), and `kuramoto_cauchy_problem` states existence and uniqueness together as a single `ExistsUnique`: the trajectory predicate is inhabited exactly once for each initial state, rather than only by the solutions we write down by hand.

A two-oscillator solution illustrates the dynamics. On two oscillators with unit coupling and zero natural frequency the phase difference obeys $\dot{\Delta} = -2\sin\Delta$, whose solution $\Delta(t) = 2\arctan(c\,e^{-2t})$ splits symmetrically into an exact trajectory of the full system (`Examples` §15, `pairRelax_is_trajectory`). At $c = 1$ it starts a quarter turn out of phase (`pairRelax_not_locked_at_zero`); the cosine of the phase difference tends to 1, the order parameter tends to 1 through the two-oscillator identity $r^2 = (1+\cos\Delta)/2$, and the dynamic potential tends to its global minimum. By the uniqueness result it is the only solution through its initial state.

Two general statements describe the long run, both for a zero-frequency system, which by `kuramoto_potential_unbounded_below` is the only case in which the potential is bounded below at all. `dynamic_potential_antitone` upgrades the instantaneous descent to monotonicity in t — `dynamic_potential_deriv_nonpos` constrains each instant separately and does not by itself say the trajectory settles — and with the uniform lower bound $-\frac{1}{2}\sum_{ij} |A_{ij}|$ monotone convergence gives `dynamic_potential_tendsto`: the potential converges along *every* trajectory.

Integrating the Lyapunov identity gives the quantitative form, `dissipation_integral_eq`: energy dissipated over $[0, T]$ is exactly the potential dropped, so by `dissipation_integral_tendsto` the improper integral $\int_0^\infty \sum_i \dot{\theta}_i^2\, dt$ converges, to the total drop. The proof uses the derivative identity in `HasDerivAt` form, which includes differentiability and permits application of the fundamental theorem of calculus. `rotating_frame_dissipation` transports both to a uniform-frequency system.

Finiteness of that integral is not the same as $\sum_i \dot{\theta}_i^2 \to 0$: an integrable function can spike forever on ever narrower intervals. Barbalat's lemma supplies the implication, proved here as `tendsto_zero_of_lipschitz_of_integral_le`: a non-negative Lipschitz function with $\int_0^T g$ bounded above tends to zero, because a spike of height ε drags a window of width $\varepsilon/2K$ up with it and only finitely many such windows fit under a finite integral.

Its hypothesis is then checked for the dissipation rate without differentiating it: the velocities are bounded by $\sum_i |\omega_i| + \sum_{ij} |A_{ij}|$, the field is Lipschitz in the configuration, and the configuration is Lipschitz in time because its derivative *is* the bounded field (`trajectory_dist_le`, by the mean value theorem, with no invariant region needed since the bound holds on the whole state space). The conclusion is `velocity_sq_tendsto_zero`: along *every* trajectory of a zero-frequency system, $\sum_i \dot{\theta}_i^2 \to 0$; equivalently each $\dot{\theta}_i \to 0$ and the vector field along the trajectory tends to 0 in the state space. The statement applies to arbitrary trajectories under these hypotheses, and `Examples` §15 applies it to the

relaxing pair, whose velocity starts at -1 and tends to zero.

At zero frequencies, the flow is the negative gradient of its dynamic potential. The identity $\dot V = -\sum_i \dot\theta_i^2$ is formalized by `dynamic_potential_hasDerivAt`. The minimizers of this potential are characterized by phase-locking. `lojasiewicz_estimate` adds the missing ingredient: on the region where every pairwise phase difference lies within a quarter turn, $2a^2 W \le \left(\sum_{ij} A_{ij}\right) \sum_i \dot\theta_i^2$, where W is the excess of the potential over its minimum and $a > 0$ is a lower bound for the coupling. Its proof uses no spectral gap and no convexity, only two extremal indices: at the site m carrying the largest phase every coupling term has $\theta_j - \theta_m \in [-\Delta, 0]$ with $\Delta \le \pi/2$ the spread, so every term is non-positive, and the term from the site carrying the smallest phase is at most $-a \sin\Delta$ (`velocity_max_le`); against $W \le \frac{1}{2}\left(\sum_{ij} A_{ij}\right)(1-\cos\Delta)$ the estimate follows from $\sin^2\Delta = (1-\cos\Delta)(1+\cos\Delta)$ and $\cos\Delta \ge 0$, which is the only place the quarter-turn hypothesis is used.

The initial energy bound keeps the trajectory in this region. Because all terms of the excess are non-negative, a single pair satisfies $\frac{1}{2}a(1-\cos(\theta_j-\theta_i)) \le W$, so an initial excess below $a/2$ — which `dynamic_potential_antitone` preserves — keeps every pairwise cosine strictly positive for all time. A difference escaping $[-\pi/2, \pi/2]$ would have to cross $\pi/2$ in absolute value by the intermediate value theorem, and there the cosine vanishes; so the quarter-turn region is forward invariant (`spread_le_of_small`). Continuity on the connected half-line thus pins the phase differences to that branch.

The differential inequality $\dot W \le -cW$ with $c = 2a^2/\sum_{ij} A_{ij}$ then integrates, in the elementary form "$W(t)e^{ct}$ has non-positive derivative", to $W(t) \le W(0)e^{-ct}$ (`excess_decay`). The rate transfers to the velocity through $|\sin d| \le \sqrt{2(1-\cos d)}$, giving $|\dot\theta_i(t)| \le Be^{-ct/2}$ (`velocity_abs_le_exp`), which is integrable on $(0,\infty)$; since $\theta_i(t) = \theta_i(0) + \int_0^t \dot\theta_i$, the trajectory converges (`phase_tendsto`), and $W \to 0$ forces the limit to be phase-locked, hence a global minimizer with $r^2 = 1$ (`kuramoto_tendsto_global_minimum`). `rotating_frame_tendsto_global_minimum` transports it to uniform frequencies, where what converges is the relative configuration.

The two hypotheses are on the initial data and neither can be removed: splay and twisted configurations are equilibria too, so "every trajectory reaches the phase-locked state" is false, and what is proved is the arc condition version — which the two-oscillator witness satisfies for every c, since $2\arctan$ lands in $(-\pi, \pi)$.

`Examples` §17 runs the theorem on three sites, where the method of §15 is unavailable because the system has no closed-form solution: the trajectory is produced by `is_kuramoto_trajectory_exists` from the initial data $(0, 0, \frac{1}{2})$, which is not phase-locked, and its limit by the convergence theorem as $\theta(0) + \int_0^\infty \dot\theta$; the numerical content of checking the hypotheses is $\cos\frac{1}{2} > \frac{3}{4}$, supplied by `Real.cos_bound`.

1.6 The fold-and-cross proof of strict decrease

Uniqueness of the coherent branch is the injectivity of $\mathcal{S}(a) = \mathbb{E}_a[\sin^2\theta]$ on $(0,\infty)$, and what is proved is the stronger statement that $\mathcal{S}$ is strictly decreasing on $[0,\infty)$ (`vonMisesSRatio_strictAntiOn`). A folding and crossing argument proves the result without differentiation. Differentiating the exponential family in its natural parameter gives $\mathcal{S}'(a) = -\mathrm{Cov}_a(\cos^2\theta, \cos\theta)$, which costs differentiation under the integral sign and leaves a covariance that is not sign-definite pointwise: symmetrizing with an independent copy gives $\frac{1}{2}\mathbb{E}[(X-X')^2(X+X')]$ with $X = \cos\theta$, an integrand changing sign with $X + X'$. The target is kept as the two-point inequality $C_2(a)Z(b) < C_2(b)Z(a)$ for $a < b$, and two moves settle it.

Fold first: the two reflections that proved $I_2 \ge 0$ carry both integrals onto $[0, \pi/2]$, replacing $e^{a\cos\theta}$ by $W(a,\theta) = 2\cosh(a\cos\theta)$ and confining $\cos\theta$ to $[0,1]$, after which $\cos 2x - \cos 2y$ and $W(b,x)W(a,y) - W(a,x)W(b,y)$ have the same sign, both being governed by $\cos x$ against $\cos y$. This is the symmetrization $X \mapsto -X$ that the covariance form needed the tilt to perform, done once at the start; its analytic content is one identity, $2\cosh X \cosh Y = \cosh(X+Y) + \cosh(X-Y)$ with cosh increasing in $|\cdot|$.

Then cross: instead of integrating that inequality over a square, let α be the mean of $\cos 2x$ under $W(a,\cdot)$ and $x_0 \in [0, \pi/2]$ a point where $\cos 2x_0 = \alpha$, supplied by the intermediate value theorem; evaluating at $y = x_0$ and integrating in x alone gives the result, since the term carrying $W(b, x_0)$ vanishes by the choice of α. No product measure and no Fubini appear.

Run with $\cos\theta$ in place of $\sin^2\theta$ the same argument needs no fold — $R(a) = \mathbb{E}_a[\cos\theta]$ is the mean of the variable the family is tilted by — and gives `besselRatio_strictMono`: R is strictly increasing on all of $\mathbb{R}$, the monotonicity of the Bessel ratio.

2 Installed coupling and its resource condition

The mode construction connects a declared hardware configuration to its kernel and stored energy. The scalar threshold applies only after identifying the kernel average with the phase model's coupling.

2.1 Installed energy bounds mean coupling

A spatial model replaces the finite sum by integration against a kernel. One declared implementation carries response modes with continuous profiles ϕ_i, installation prices u_i and nonnegative occupancies $c_i(\Xi)$ in configuration Ξ:

$$K_\Xi(x,y) = \sum_i c_i(\Xi)\phi_i(x)\phi_i(y), \qquad U(\Xi) = \sum_i u_i c_i(\Xi). \tag{4}$$

The same configuration sets the reciprocal kernel and stored energy. An executed transition changes the kernel through its changed occupancies and requires work $U(\Xi') - U(\Xi) + Q$, with Q positive into the reservoir. The installed kernel is averaged over the resulting configuration law $p(\Xi)$. Cell-pair reconstruction on refining spatial partitions converges to its product-measure integral under the specified compactness and sampling assumptions; the refinement index is independent of the number of control updates.

Normalize the substrate measure μ to one. Write $m_i = \int \phi_i \, d\mu$ and declare $\kappa > 0$ with $m_i^2 \le \kappa u_i$ for every mode. Factorizing each mode's integral gives

$$\begin{aligned} U_{\mathrm{inst}} &= \sum_\Xi p(\Xi) U(\Xi), \\ K_{\mathrm{mf}} &= \iint K(x,y)\, d\mu(x)\, d\mu(y) = \sum_\Xi p(\Xi) \sum_i c_i(\Xi) m_i^2 \le \kappa U_{\mathrm{inst}}. \end{aligned} \tag{5}$$

If this mean coupling is the K of the scalar model, $\kappa U_{\mathrm{inst}} \le 2D$ excludes a coherent stationary branch and any growing Fourier mode about the uniform density. Thus coherence on that branch requires $U_{\mathrm{inst}} > 2D/\kappa$. The condition is necessary: energy above it need not produce sufficient coupling because the bound can be loose.

The theorem concerns stored energy in the declared actuator modes. The estimated cortical signalling expenditure and extracellular electrostatic energy of Section 2.3 measure different quantities; neither identifies the occupancies, mode prices or their relation to metabolic supply. Consequently the bound supplies a necessary condition for a declared actuator, but no cortical value of κ or demonstration that cortex meets it. Applying the scalar threshold to heterogeneous or signed spatial dynamics also needs a justified reduction.

2.2 Microscopic modes and a spatial kernel

Formalization: *Phase3_LocalActuator.lean*.

`LocalActuator` carries mode occupancies c_i on configurations, fixed spatial profiles ϕ_i and installation prices u_i; its `kernel` and `storedEnergy` are the two sums of Eq. (4). `kernel_update` computes a transition's kernel change from its changed occupancies alone, and `kernel_update_local` leaves a pair untouched when every changed mode has zero profile at one of its two points; `kernel_symmetric` and `kernel_nonneg` record reciprocity and, for nonnegative occupancies and profiles, positivity. `pathWork` is $U(\Xi') - U(\Xi) + q$ along one path, and `installation_first_law` evaluates its mean on the executed step as the change in expected stored energy plus that step's mean heat. `expected_kernel_update` states the ensemble kernel change as the mean of those same pathwise changes, so an unrelated final law or a prospective policy evaluation does not substitute for the executed one.

For the installed kernel, refining cell-pair reconstruction gives the limit

$$\mathcal{K}_n \longrightarrow \iint K \, d\mu \, d\mu. \tag{6}$$

under the continuity, compactness and sampling conditions specified below.

`KernelMesh` (`Phase2_KernelMesh.lean`) is a finite measurable partition with one sample per cell, whose reconstruction extends each sampled weight throughout the whole cell pair. `energy_eq_integral` proves the cell-mass weighted finite energy to be exactly that reconstruction's product-measure integral;

it needs finite mass and measurable cells, and neither continuity nor refinement. `kernel_tendsto` and `energy_tendsto` then give Eq. (6) for a continuous kernel on a compact second-countable substrate of finite mass, with `snap_tendsto_of_dist` reducing the sample hypothesis to a shrinking distance bound. `KernelArrangement` bundles the hardware, its configuration law, the substrate measure and the partition sequence, and its `coarseGrains` is the limit. In `Chain.lean`, `e45Active_of_localActuator` discharges the active edge for an arrangement holding a named step's final law, and `e45_of_localActuator` discharges the passive edge for one holding a declared predictive system's joint law.

The witness in `Examples/MicroscopicCoupling.lean` is one response mode on $[0,1]$ with profile $(4/3)(1+x)$ and unit price, held by the two-bit thermal channel: the mode is occupied exactly when controller and switch agree. Before the channel the installed kernel is $(8/9)(1+x)(1+y)$ and after it $(4/3)(1+x)(1+y)$, so executing the channel strictly raises a pairwise value; closing the switch takes the pathwise kernel at the origin from 0 to $16/9$. The mean external work is $(1+\log 3)/4 > 0$, which the first law splits into the increase of expected stored energy from $1/2$ to $3/4$ and the channel's $(\log 3)/4$ of heat, and a zero-work claim is rejected for those same paths. Clamped left-endpoint partitions into $n+1$ cells have sample error at most $1/(n+1)$; the first two coupling energies are $4/3$ and $25/12$, and the limit is 3. That limit matches the existing field witness, so `chain_active_microscopic_jointly_satisfiable` composes the active branch through a kernel on the product space, and a zero target is rejected for the same sequence. The unit interval is atomless, so these energies are cell masses rather than vertex weights. Reading §18's predictive law on the same hardware installs $(16/9)(1+x)(1+y)$ with limit 4, so the configuration law is load-bearing, and §18's two systems differ by their whole dissipated work while discharging the identical passive edge.

Mode profiles, prices, occupancy readout, the reservoir convention and the substrate measure are declared hardware data. The model accounts for one prepared update: fabricating the hardware, preparing its initial configuration law and supplying continuing power are separate. The spatial index is independent of the number of executed updates, and convergence in one is not convergence in the other. Nothing here identifies the substrate with cortex or the modes with a physical field, and no argument makes predictive efficiency select this hardware.

2.3 Installed energy and the stationary threshold

In `Phase9_InstalledCoupling.lean`, the theorem `continuumEnergy_eq_mean_field_coupling` in `KernelArrangement` proves the Fubini identity in Eq. (5). Compactness, finite substrate mass and continuous profiles supply integrability; the identity itself needs no probability normalization. Its hypotheses equate the arrangement's measure and kernel to the field's. The scalar threshold additionally uses a probability substrate. `continuumEnergy_eq_modes` factorizes the integral into the configuration mean of $\sum_i c_i(\Xi) m_i^2$, with $m_i = \int \phi_i \, d\mu$. `PricedArrangement` carries a positive conversion factor κ, nonnegative occupancies and the mode condition $m_i^2 \le \kappa u_i$. Positive conversion and this condition also force nonnegative prices. `continuumEnergy_le_installedEnergy` multiplies each mode condition by its occupancy and averages over the actual configuration law, giving $0 \le K_{\mathrm{mf}} \le \kappa U_{\mathrm{inst}}$.

`no_coherence_of_installedEnergy` composes the bound with $K_c = 2D$: if $\kappa U_{\mathrm{inst}} \le 2D$, `exhibits_phase_transition` fails, zero is the only nonnegative self-consistent order parameter, and no positive-index Fourier mode of the uniform state grows. `stationary_eq_incoherent_of_installedEnergy` uses the stationary-density classification to identify any positive classical self-consistent density with $1/(2\pi)$. These are results for the scalar cosine-drift model in coordinates with mean direction zero, without a theorem transferring stability to heterogeneous spatial dynamics. The bound concerns installed energy rather than dissipated heat: the reversible register–bath gate transfers energy with zero path entropy production and does not satisfy the proposed transition-log heat identification. The conversion factor is declared hardware data, independent of the noise threshold; the mode condition checks its admissibility, not when it was chosen. Exceeding $2D/\kappa$ is necessary and does not guarantee coherence. Table S3 lists the connecting hypotheses E12–E89 of the formal composition. E56's identification of the kernel with a cortical field remains a physical commitment. In `Chain.lean`, `installedEnergy_required_by_chain` derives $U_{\mathrm{inst}} > 2D/\kappa$ from E56 and E67 for the arrangement's own convergent energy sequence; it supplies neither edge.

The bound against a cortical energy budget

The condition $U_{\mathrm{inst}} > 2D/\kappa$ concerns energy stored in declared actuator modes. A metabolic budget supplies a different quantity. Take a sphere whose radius is one LFP decay length, 0.2 mm, containing 1676 neurons. At 3.29×10^9 ATP s^{-1} per neuron for a mean action-potential rate of 4 Hz in rodent

grey matter (Attwell & Laughlin, 2001), of which synaptic mechanisms take 55% (Harris et al., 2012), and 57 kJ per mole of ATP hydrolysed at cytoplasmic concentrations (Nicholls & Ferguson, 2013), each neuron draws 3.1×10^{-10} W and the region's signalling power is 5.2×10^{-7} W, of which 2.9×10^{-7} W is spent at synapses. Over a declared interval of one phase-diffusion time $1/D = 2$ s at $D = 0.5$ rad^2/s, the metabolic expenditure is 1.0×10^{-6} J. Another interval rescales that expenditure by its ratio to this one. For comparison, a stored-field estimate using LFP gradients of 1–5 mV/mm, radii 0.1–0.3 mm and an illustrative low-frequency relative permittivity of 10^5 gives electrostatic energy 1.9×10^{-20}–1.3×10^{-17} J (or 4×10^0–3×10^3 kT). Tissue impedance and dielectric measurements motivate these input ranges (Barbour, 2017; Logothetis et al., 2007; Gabriel et al., 1996); the static-capacitor approximation remains a modelling choice.

The estimates differ by many orders of magnitude, but neither is a measurement of U_{inst}. Continuous metabolism may replenish an actuator; a relation between its power, storage and coupling must be specified before that claim can be used in the theorem. The small electrostatic estimate alone does not show that the theorem's actuator energy is too small, because its priced modes have not been identified with the static field energy.

If a physical actuator model equated the expenditure over this interval with energy installed in its priced modes, the conditional requirement would be $\kappa > 9.6 \times 10^5$ J^{-1}s^{-1}. The expenditure is 2.4×10^{14} times kT at body temperature, but that large number cannot discharge an installed-energy premise. A mode is a priced rank-one contribution to the kernel; cortical field generation has no specified decomposition into such modes. The estimate therefore provides neither a cortical conversion factor nor a verdict on the threshold.

`Examples/InstalledCoupling.lean` uses the existing profile $(4/3)(1 + x)$ with spatial mass 2 and unit price, fixing the sharp factor $\kappa = 4$ before specifying diffusion $D = 1$. Prepared agreement probabilities 1/4, 1/2 and 3/4 give installed energies 1/4, 1/2 and 3/4 and mean couplings 1, 2 and 3. The first two are excluded, including the threshold; the executed third law is above it. A mode with the same profile, price 2 and occupancy 1/2 gives $\kappa U_{\mathrm{inst}} = 4$ but $K_{\mathrm{mf}} = 2$, rejecting the converse. A negative price and an underpriced mode fail the mode condition; every admissible conversion factor places the executed arrangement above the energy line, rejecting a factor fitted to force exclusion. Finally, adding a zero-response mode with occupancy -1 cancels the stored energy of one occupied original mode while leaving mean coupling 4: its prices and mode conditions pass but its nonnegative occupancy condition fails, and the claimed exclusion would be false. These are exact formal witnesses, with preparation mechanisms and physical calibration supplied by no result here.

The class field in `Phase3_KLBound.lean` is separate from Gibbs' nonnegativity theorem. The witness `boolDetuned` attains the assumed bound with $D_{KL} = \sigma = \log 2$, showing that a positive divergence is permitted.

The fixed-phase squared-drift functional in `Phase8_ContinuousField.lean` is a quadratic functional of an affine kernel map. The bounded operators `kernelCLM` and `continuumDriftCLM`, together with the Hilbert-space adjoint, give $\nabla\sigma = (2/D)A^*(AK + \omega)$. The theorem `structural_resonance_decreases_sigmaContinuum` proves descent for the specified gradient flow, with an atomless witness on the unit interval. The identifier `entropy_production_rate` names this squared-drift functional in the source; the name does not establish its identification with thermodynamic entropy production. That interpretation requires probability-current information absent from this calculation.

For a symmetric kernel, antisymmetry of $K(x, y)\sin(\theta_y - \theta_x)$ makes its double integral zero. Thus the mean drift is independent of the phase field. A dynamically phase-locked field attains the associated squared-drift lower bound under the stated hypotheses. The theorem `phase_locked_minimizes_entropy_of_symm` expresses this minimization. A one-way two-site coupling supplies a counterexample without symmetry. These are statements about the objective and its constrained mean, not convergence of joint noisy phase/plasticity dynamics.

2.4 Discrete coupling and mesh refinement

From continuous to discrete coupling: The coupling matrix A_{ij} above is not postulated. A triangulation of the substrate assigns to each edge (u, v) a region of the manifold, and the coupling weight is the Lebesgue–Bochner integral of the scalar magnitude of the stress-energy tensor over that region (`edge_weight`); symmetry and non-negativity of A are then proved rather than assumed. The bridge back is mesh refinement: as the regions shrink, the discrete energy $\frac{1}{2}\sum_{u,v}\mu(R_{uv})\,f(x_u)$ should converge to the continuous energy $\int_S f$. Stated as a conjecture over the bare triangulation class, that is false: nothing there forces the regions to cover the manifold, empty regions have diameter zero, and the everywhere-

empty triangulation refutes it. Stated correctly it is a theorem (`mesh_refinement_convergence`), quantified over a sequence of triangulations of a region S whose regions are measurable, pairwise disjoint across distinct edges, covering, and anchored to the embedded vertices, with mesh size tending to zero. The fixed-mesh error is bounded by $\varepsilon\,\mu(S)$, with ε a modulus of continuity of the integrand, and the symmetric double sum is proved equal to the midpoint rule over unordered edges.

Two limitations follow. The theorem gives convergence, not the $O(1/N^2)$ rate reported by `mesh_refinement.py`. And `TriangulatedManifold` itself is bare data: it pairs a complex, an embedding and a family of edge regions without relating them, so the geometric conditions live in a separate predicate `IsRegularTriangulation`. Because those conditions are in tension — cells disjoint yet covering, every vertex anchored to each of its edges — `Examples` §6 exhibits an instance: the uniform partition of $[0, 1)$ into N half-open cells, on which the convergence theorem is instantiated and the approximations are checked to genuinely change with N.

The discretization structure `DiscreteThermodynamics` *carries* that predicate as a field, which is what makes its edge weights integrals over a geometrically constrained family rather than over arbitrary sets, with three consequences for the coupling weights. `weight_self`: a vertex has no edge to itself, so the coupling matrix has zero diagonal. `face_of_weight_ne_zero`: a nonzero weight $w(u, v)$ forces $\{u, v\}$ to be an actual face of the complex, so the discretization cannot invent couplings between vertices the triangulation does not join. And `total_weight_eq_setIntegral`: $\frac{1}{2}\sum_u \sum_v w(u, v) = \int_S |T|$, i.e. the total coupling strength is exactly the scalar stress-energy of the region covered — nothing counted twice (disjointness), nothing dropped (covering). That last identity is the precise sense in which the coupling matrix is derived from the continuum rather than posited, and it is exactly what the everywhere-empty triangulation violates.

`Examples` §6 witnesses the structure on the same grid: each consecutive edge carries weight $1/N$, and the weights sum to 1 for every N. The witness is one-dimensional and its stress-energy tensor is constant, so it establishes satisfiability, not a physically interesting discretization.

One field of the structure is deliberately left free and is marked `[MODELLING]`: `scalar_magnitude`, the scalar read off the stress-energy tensor, constrained only to be non-negative. A $(0, 2)$ tensor has no canonical scalar — the trace, an observer's energy density and the tensor norm are different functions, and the structure carries neither a metric nor a frame — so imposing one would pin a choice rather than derive it, and the three statements above hold for every non-negative choice, so it would prove nothing new.

That the choice is nonetheless a physical input is stated as a theorem rather than left as a remark: §6 exhibits a second instance over the same triangulation, tensor and regions, differing only in that field, for which the partition identity again holds while every edge weight differs (`gridThermo_magnitude_matters`). The remaining constraint, non-negativity, is the discrete shadow of an energy condition and is what makes the induced couplings a legal Kuramoto coupling.

3 Global sections and the unity of consciousness

Phase synchronization does not imply overlap compatibility. The model represents global consistency by gluing compatible local descriptions; its identification with conscious unity remains a separate stipulation.

Theorem S2 (Sheaf-Theoretic Unity): In the sheafification of finite spatial measures over a substrate X, local sections agreeing on all spatial intersections $U \cap V$ uniquely glue into a global section. On a finite open cover, compatible finite-measure representatives have a unique finite-measure extension whose sheaf image is that section. The overlap agreement is an independent physical hypothesis; phase-locking does not imply it. This is not a theorem about arbitrary probability marginals over decoded contents, for which compatible overlaps need not determine a unique global joint distribution (Abramsky & Brandenburger, 2011).

Lean 4 Implementation (`Phase5_GlobalSection.lean`): The headline gluing theorem separates two inputs. The class field `ThermodynamicCover.thermodynamic_equilibrium` places the phase field at the potential minimum, while `LocalSectionSynchronization.section_agrees_of_phase_eq` independently requires equal-phase patches to agree on overlaps. The second supplies the compatible family consumed by this sheaf theorem. The vector-content models below give conditions for agreement of their own local fields; they do not discharge this class field for an arbitrary cover.

What the second one says can be stated without the sheaf: on a discrete substrate two patch sections have the same restriction to their overlap exactly when they assign the same mass to every site they share (`restrict_eq_iff_densityOn_eqOn`), so the hypothesis is that patches at a common phase carry a common local density where they meet. It excludes data the rest of the structure admits: a two-patch

cover of the three-site substrate at one phase, with profiles differing at the shared site, satisfies every other requirement and fails this one (`overlap_agreement_fails`). It does not require patches to agree away from their overlaps — the three-patch witness carries three pairwise-different profiles.

The equilibrium condition can be obtained from a convergence result for a class of initial data: `Phase5_EquilibriumBridge.lean` supplies `ThermodynamicCover.ofConvergentTrajectory` from `kuramoto_tendsto_global_minimum`, and `Examples` §17.1 is a three-patch cover built over the resulting phase field.

The cover class also has an independent instance (`Examples`, §4), so the theorem is not conditional on a structure of unknown realizability. Its three-site substrate has two overlapping patches meeting at one site; local sections are germ families of the phase-indexed measure $\mu_\theta = (1 + \cos\theta)^+ \, \delta_{\mathrm{mid}}$ in the sheafification of the probability presheaf. The family is provably non-constant in θ (`phaseMeasure_not_const`), and `thermodynamic_equilibrium` follows from `phase_locked_minimizes_potential`.

Two further results specify the phase and section data. The theorem `potential_min_iff_phase_locked` (`Phase4_KuramotoDynamics.lean`) characterizes the minimizers of the reduced potential in both directions, from which `ThermodynamicCover.phase_locked` follows for every instance. The witness in `Examples` §13 carries a non-constant phase field and a non-uniform invariant measure, with patch-local densities read from the sections through `densityOn`; agreement at the shared site is an equation between separately computed numbers. The limit is structural: phase-locking equates phases, while agreement of the associated sections remains the independent class field named above.

The gluing construction assumes local overlap agreement and derives both existence and uniqueness of the global section. The lemma `probability_glue_unique` takes a compatible family to its unique global extension; `ThermodynamicCover.invariantMeasure` denotes that section. The construction uses local data, not a global state supplied as a premise.

Representation by an actual measure. `SpatialMeasure.finite_glue_unique` (`Phase1_MeasureGluing.lean`) proves finite-measure gluing directly on a finite measurable cover. Order the patches, remove the parts already assigned to preceding patches, and sum the local measures restricted to those disjoint pieces. Overlap agreement gives the prescribed restriction on every patch, and equality of those restrictions gives uniqueness. Each local measure can be extended by zero to the ambient space without changing its restriction (`measureOnOpen_map_subtype`). `ThermodynamicCover.existsUnique_measure_representation` identifies the resulting measure's sheaf image with the cover's `invariantMeasure`. It requires compatible finite-measure representatives of the supplied local sections; these are additional hypotheses, with no global measure supplied. No equivalence between arbitrary sheafified sections and finite measures on arbitrary spaces is asserted. The real-line witness uses two overlapping open intervals with distinct local measures and counts their shared atom once. A countably infinite cover requiring unit mass at every natural-number site has no finite-measure extension. These results impose no normalization and choose no interpretation of the masses as content.

3.1 Approximate overlap agreement

Theorem S2 admits no tolerance: the sheaf condition is an equation. Under an ε-tolerance neither existence nor uniqueness of a global section survives, and what replaces them is a selection rule with an explicit constant. The results are stated for a finite substrate whose local states are nonnegative site masses, with the uniform distance $\max_{x\in U} |f(x) - g(x)|$ on the sites a pair of patches shares (`ApproximateGluing.profileDist`, `Phase5_GlobalSection.lean`).

Let $\{U_i\}_{i\in I}$ be a finite cover, let s_i be the local profiles with $\mathrm{dist}(s_i, s_j) \le \varepsilon$ on $U_i \cap U_j$ for *every* pair (`Compatible`), and let $w_i \ge 0$ vanish off U_i and satisfy $\sum_i w_i(x) = 1$ at every site (`IsPartition`). The selection is

$$(\mathrm{select}\ w\ s)(x) \;=\; \sum_i w_i(x)\, s_i(x).$$

It depends on the profiles only through their values on their own patches (`select_eq_of_eqOn`), so the bookkeeping values a patch carries elsewhere do not enter. Four estimates follow. The selection is within ε of each patch on that patch (`select_close`); two partitions subordinate to the same cover select states at distance at most ε (`select_dist_le`); any two profiles fitting every patch to tolerance δ lie within 2δ of each other (`approximate_diameter_le`); and at $\varepsilon = 0$ the selection is the unique profile fitting all patches (`select_glue_unique`).

The constant in the second estimate is $C(N) = 1$ at every cover multiplicity N, rather than the $N\varepsilon$ a term-by-term count would give: the weights sum to one rather than to the number of patches active at a site. This is a statement in the uniform mass metric and not in total variation. A feature F that

is L-Lipschitz on profiles therefore satisfies $\mathrm{dist}\big(F(\text{select } w\, s), F(\text{select } v\, s)\big) \le L\varepsilon$ (`feature_dist_le`), which is the resolution condition an interpretation at scale δ must meet; the theorem chooses no decoder.

The three-site witness exhibits both halves (`Examples/Phase5.lean` §21). Its three two-site patches have multiplicity two, and the family `apxW` is §17.1's data with one patch reporting ε more mass everywhere it can see. Two partitions — each site assigned to the lowest, respectively the highest, indexed patch containing it — select states at distance exactly ε (`apx_select_dist`), so the constant is attained and cannot be improved at this multiplicity. For every positive ε that same family has no exact global section at all (`apx_no_global_section`): two patches see the site `right` and disagree there, so selection stands in for a gluing that does not exist. The corresponding sections fail the class field of Theorem S2 exactly as the two-patch fence does (`apxSection_overlap_fails`), and the tolerance cannot be dropped to zero (`apxW_not_compatible_zero`).

Checking a spanning subfamily of the overlaps is not checking the family. The family `chainW` agrees exactly on the overlaps $(0,1)$ and $(1,2)$ and is off by 2ε on $(0,2)$; its two selections are 2ε apart (`chainW_not_compatible`, `chainW_select_dist_gt`), so the hypothesis quantifies over all pairs. On the same witness the total-mass readout moves by 2ε between the two selections, within the 3ε its Lipschitz constant permits (`totalMass_apx_dist`, `totalMass_apx_dist_le`): the selection dependence of a readout is a real effect and not an artefact of a loose bound. At $\varepsilon = 0$ the three-patch family of §17.1 satisfies the relaxed hypothesis and the selection returns its glued profile (`trioW_compatible`, `trioW_glue_unique`).

3.2 Contents driven by observation

The agreement above is a hypothesis on a supplied family. What a mechanism would have to do is change the family, so the contents are carried as fields $s_i : A \to \mathcal{C}$ in a real normed vector space $\mathcal{C}$ on the patches of a declared cover. Disagreement at a shared site is $\|s_i(x) - s_j(x)\|$, and each patch's field is mixed with that patch's observation at gain $\eta \in [0,1]$, $s_i \mapsto (1-\eta)s_i + \eta o_i$. Disagreement across an overlap then obeys

$$\varepsilon' \le (1-\eta)\varepsilon + \eta\delta, \tag{7}$$

with ε the contents' mismatch and δ the observations' own (`update_residual`, `Phase5_ContentDynamics.lean`). Observations that agree on their overlaps preserve exact agreement, and contract any initial mismatch to $(1-\eta)^n\varepsilon$ after n of them, even when the value they agree on moves (`update_preserves`, `run_residual`). Observations agreeing only to tolerance δ contribute that tolerance to the upper bound in Eq. (7).

A readout names which patch to read at each site. Where that choice covers the site, the readout is within ε of every patch that sees it (`readout_residual`); at $\varepsilon = 0$ it extends every patch and is the only field that does (`readout_agrees`, `readout_unique`). Coverage is a hypothesis of those theorems rather than a property of a pre-existing global field, which is the same discipline the selection results above follow. The nonnegative mass profiles are the scalar case $\mathcal{C} = \mathbb{R}$, read through the coercion at the same tolerance (`ApproximateGluing.compatible_pointwise`), so the two agreement predicates coincide on those profiles.

`compatible_real` and `update_real` identify the scalar instances. A common nonexpansive linear map preserves compatibility (`compatible_map`); the converse fails. `Examples/VectorContent.lean`, §22, supplies a nonconstant two-component family with a unique extension and two vectors whose first projections agree while their second components disagree.

Observations agreeing only to δ are covered by the same recursion. Iterating Eq. (7) bounds disagreement by $(1-\eta)^n\varepsilon + \delta$ (`run_residual_floor`). The geometric term decays for $\eta > 0$, while the bound retains δ. This is an upper bound; it does not assert that the actual disagreement reaches that value. The next subsection supplies a condition under which the observations themselves agree.

The contents remain vector-valued fields identified with no neural variable and with no section of the probability sheaf.

The selection results concern nonnegative spatial mass profiles on a finite substrate, while these updates act on normed-space-valued local fields. Using the scalar instance as a dynamics of mass profiles requires observations and initial profiles that preserve nonnegativity. Neither construction concerns matching marginals of decoded contents (Abramsky & Brandenburger, 2011), or selects a cover, decoder or weights. Theorem S2, the `ThermodynamicCover` instances and the composed chain retain their exact agreement hypothesis.

3.3 The decodability cover

The cover of Section 4 of the article has content variables as its base rather than sites (`Phase5_Content Cover.lean`). A `DecodingSetup` declares a relevant family of states, the true value $v(s)$ of each variable v, what each region i reads of a state, and region i's decoder for each variable. Region i's domain at tolerance δ is the set of variables it decodes within δ at every relevant state (`DecodingSetup.domain`); two regions overlap on the variables both decode, and restriction forgets variables. The worst-case tolerance idealizes the article's statistical criterion, cross-validated decoding above a permutation null, and does not formalize it.

Four of the spatial results transfer. Give the variables the discrete topology: a section over a set of variables is then a real function on it, Mathlib's presheaf of functions is a sheaf, and the gluing theorem used for the spatial cover, `sheaf_glue_unique`, applies unchanged (`decode_glue_unique`). At $\delta = 0$ the glued section is the true content (`decode_glue_value`). The decodes form a family of contents in the sense of Section 3.2, and on this cover their compatibility is a theorem rather than a hypothesis: two regions that decode v within δ satisfy

$$|d_i(v) - d_j(v)| \le |d_i(v) - v| + |d_j(v) - v| \le 2\delta$$

(`decode_compatible`), which is the accuracy confound of the article's disconfirming comparison (Section 10.2) in formal form. The weighted selection of Section 3.1 holds for signed and vector-valued contents on any index of variables: it is within ε of each patch and two selections differ by at most ε, at every multiplicity (`ContentCover.select_close`, `ContentCover.select_dist`). On the decodability cover it is within δ of the true value (`select_decode_near_value`). The decoder class needs an upper bound. A region whose readout is injective on the relevant family decodes every variable exactly with some decoder (`exists_exact_decoder_of_injOn`), so with unrestricted decoders every such region overlaps every other on everything. A region that reads two relevant states identically decodes no variable on which they differ by more than 2δ, whichever decoder it uses (`not_decodes_of_read_eq`).

`Examples/ContentCover.lean`, §37, has three variables and two regions reading $\{0,1\}$ and $\{1,2\}$, each decoding the shared variable with an offset b in opposite directions. For $0 \le b < \frac{1}{2}$ the cover at tolerance b is exactly those two domains, covers, and has a connected nerve (`twoRegions_domain_zero`, `twoRegions_covers`); at $b = 0$ the regions glue to the true content (`twoRegions_exact_glue`). At $b = 1/10$ the decodes of the shared variable differ by exactly $2b$ at every state, so the bound is attained (`twoRegions_disagreement`) and no exact global section exists (`twoRegions_no_exact_glue`); equal weights cancel the offsets and select the true content.

The probability presheaf, its measure representation and the phase-to-content bounds of the next subsection do not transfer. A section over a set of variables is a list of values rather than a distribution, and the $\sqrt{N}$ bounds and their nerve and spectral forms bound disagreement between spatial patches of oscillators, where on this cover disagreement is bounded by decoding error. Nothing here relates a region's phase order to what it decodes; that relation is the bridge assumption E78. Nothing here selects the candidate variables, the tolerance or the decoder class either.

3.4 Agreement acquired from coherence

For declared encoders that agree as functions on overlaps, phase coherence bounds observation disagreement. The shared-encoder hypothesis specifies how two patches describe the same quantity; coherence constrains their inputs.

Write N for the number of oscillators and r for the order parameter. The identity

$$N^2 r^2 = \sum_{i,j} \cos(\theta_i - \theta_j) \tag{8}$$

(`order_parameter_r_sq_eq_mean_cos`, `Phase4_KuramotoDynamics.lean`) rewrites $N^2(1 - r^2)$ as $\sum_{i,j}\big[1 - \cos(\theta_i - \theta_j)\big]$, a sum of nonnegative terms. A pointwise bound is sharper when read about the mean rather than across pairs. With $z_k = e^{i\theta_k}$ and $m = N^{-1}\sum_k z_k$, so that $|m| = r$,

$$\sum_k |z_k - m|^2 = N\big(1 - r^2\big) \tag{9}$$

(`sum_normSq_sub_order_parameter`). For $i \ne j$ the two terms $|z_i - m|^2$ and $|z_j - m|^2$ are at most this sum, and $|z_i - z_j|^2 \le 2(|z_i - m|^2 + |z_j - m|^2)$, so

$$1 - \cos(\theta_i - \theta_j) = \tfrac{1}{2}|z_i - z_j|^2 \le N\big(1 - r^2\big) \tag{10}$$

for *every* pair (`chord_sq_le_of_coherence`, `cos_gap_le_of_coherence`). The factor N comes from bounding two nonnegative deviations by their total. Its order cannot be lowered in general: one oscillator in antiphase among N has $1-\cos = 2$ against its partners while $N(1-r^2) = 4(N-1)/N$.

A phase is an angle, so a content that is a function of phase must be a function of the circle point $(\cos\theta, \sin\theta)$; a model reading θ as a real number is not 2π-periodic and is not a function of the phase at all. In the chord metric on that circle, $\|\cdot\|^2 = 2(1-\cos\Delta)$, and Eq. (10) becomes

$$\operatorname{chord}(\theta_i, \theta_j) \le \sqrt{2N\big(1-r^2\big)} \tag{11}$$

(`chord_le_of_coherence`). Patch i reads oscillator $v(i)$ through an encoder $e_{i,x}$ taking circle points to the content space $\mathcal{C}$. Require one uniform Lipschitz constant L and

$$x \in U_i \cap U_j \quad \Longrightarrow \quad e_{i,x} = e_{j,x}.$$

These are `UniformLipschitzEncoder` and `SharedEncoder`. Observations can vary across sites:

$$o_i(x) = e_{i,x}(\cos\theta_{v(i)}, \sin\theta_{v(i)}).$$

The resulting family is compatible at residual $L\sqrt{2N(1-r^2)}$ (`compatible_of_shared_coherence`, `Phase5_ContentDynamics.lean`). The residual is computed from the phases using the supplied equality of encoders on overlaps. At $r^2 = 1$ it is zero and the observations agree exactly (`compatible_of_shared_phase_locked`). The one-encoder theorems are special cases with $e_{i,x} = e$. A uniform bound on the phase residual over successive observations gives `run_residual_of_shared_coherence`.

The population in Eq. (11) need not be the whole array. Nothing in the derivation uses more than a finite site type carrying an order parameter, and a declared patch P supplies one: the order parameter of the subtype of P is exactly P's resultant, by the two facts that the sum over the subtype is the sum over P and that the two counts agree (`order_parameter_complex_patch`). Instantiated there, Eq. (11) reads

$$\operatorname{chord}(\theta_i, \theta_j) \le \sqrt{2|P|\big(1-r_P^2\big)} \qquad (i, j \in P) \tag{12}$$

(`chord_le_of_patch_coherence`), and the content statement follows with it whenever every patch reads a site of P (`compatible_of_patch_coherence`). The factor that can make the global estimate weak is the population's count, so reading through small patches replaces it by a small one. What this does not do is improve the global bound: Eq. (12) constrains the pairs inside one patch and says nothing about a pair straddling two. Section 12.6 gives the state that separates the two readings furthest.

Three results decline the step that puts the count there. It has one source: Eq. (9) is a sum of nonnegative terms, and the two terms for the pair in question are bounded by the whole sum — which permits all the disorder in a population to sit at the two sites the conclusion is about.

The first asks how many terms can be large rather than how large one can be, which is Markov's inequality on the same sum. For any set S of pairs whose chords are all at least c,

$$\frac{|S|}{N^2} \le \frac{2\big(1-r^2\big)}{c^2} \tag{13}$$

(`sum_gap_eq`, `card_gap_le`, `card_chord_le`, `chord_fraction_le`), with no population size in it at all. The bad set is supplied as an arbitrary `Finset` of pairs rather than filtered, which is strictly more general — the filtered set is the largest such — and keeps a decidability instance for a real inequality out of the statement. Reading the same sum by rows gives the site form: if every site of T has at least δN partners at chord c or more, then $|T|/N \le 2(1-r^2)/(\delta c^2)$ (`card_site_gap_le`, `site_fraction_le`). Where Eq. (11) improves like $\sqrt{1-r^2}$, Eq. (13) improves like its square. What neither does is name a pair or a site: both bound counts, the exceptional set they permit is unlocated, and a cover's overlaps are a set of sites fixed before the state is known. Reading either as agreement on an overlap requires the exceptional set to miss that overlap, which is a joint fact about the state and the cover that no coherence hypothesis supplies; the only route to it is to choose the cover in the light of the state, which is the move that empties a gluing claim.

The second keeps the coherence hypothesis and drops the common patch. The nerve of a family of patches has an edge wherever two patches share a site (`patchNerve`, `SimpleGraph.fromRel` on that relation, so symmetric and irreflexive by construction). The chord is a metric (`chord_triangle`), so a bound β on the separation inside every patch composes along a walk:

$$\operatorname{chord}(\theta_x, \theta_y) \le (k+1)\,\beta \qquad (x \in P_a,\ y \in P_b,\ \operatorname{dist}(a,b) \le k) \tag{14}$$

(`chord_le_of_patch_walk`), with β read off each patch's own resultant through Eq. (12) (`chord_le_of_patch_walk_coherence`) and the content statement following with it (`compatible_of_patch_nerve`). The constant is $(k+1)\beta$ and not $k\beta$ because both endpoint sites pay a step inside their own patch, the shared sites of the walk being interior to the chain. No patch need contain both sites, which is the situation of every pair of territories far enough apart. The global resultant does not appear: a state whose global resultant is zero still constrains content across a connected cover. What is supplied is the connectivity: the nerve's diameter is a property of the cover, nothing here produces a walk, and a cover whose nerve is disconnected supports no conclusion between its components. Accumulation is linear and nothing makes it sublinear, so on the sheet of Section 12.6 the reach is 20 patch diameters rather than a hemisphere.

The third reads no resultant. A configuration is frequency-locked to a common rotation Ω when $\omega_i + \sum_j K_{ij}\sin(\theta_j - \theta_i) = \Omega$ at every site; this is not an extra postulate, being equivalent to the rigid rotation $t \mapsto \theta + \Omega t$ solving Eq. (1) without noise (`is_frequency_locked`, `is_frequency_locked_iff`). Write $Q(x) = \frac{1}{2}\sum_{ij} K_{ij}(x_i - x_j)^2$ for the coupling-weighted Dirichlet form (`couplingForm`), which on the adjacency matrix of a simple graph is the graph Laplacian's (`couplingForm_eq_lapMatrix`), and declare a spectral gap λ_2 by the Rayleigh inequality $\lambda_2\|x\|^2 \le Q(x)$ on mean-zero fields (`SpectralGap`); turning K up on a fixed graph scales it (`SpectralGap.scale`). Pairing the balance equation with the configuration symmetrizes (`sum_mul_coupling_sin`), Jordan's inequality applied termwise keeps the sine and pays $2/\pi$ rather than linearizing (`couplingForm_le_pairing`), and Cauchy–Schwarz closes:

$$4\lambda_2^2 \sum_i \big(\theta_i - \bar{\theta}\big)^2 \le \pi^2 \sum_i \big(\omega_i - \Omega\big)^2 \tag{15}$$

(`spread_le_of_frequency_locked`), read at a named pair as $2\lambda_2^2\,\mathrm{chord}(\theta_i,\theta_j)^2 \le \pi^2\sum_k(\omega_k-\Omega)^2$ (`chord_le_of_frequency_locked`) and giving a content residual of $L\pi\|\omega - \Omega\|/(\lambda_2\sqrt{2})$ with no population count, no patch and no cover geometry in it (`compatible_of_frequency_locked`). The bound improves as the coupling graph becomes better connected, which is the dependence a resultant cannot express, r being blind to who is coupled to whom. Where it is informative is fixed by the hypothesis that licenses it: confinement to a quarter turn puts every chord below $\sqrt{2}$ before any coupling datum is read, so Eq. (15) constrains the pair exactly where $\lambda_2 > \pi\|\omega - \Omega\|/2$, and the comparison that matters is against $\sqrt{2}$ rather than against the maximum chord. Four things are declared and none is discharged here. Mathlib carries the graph Laplacian and its positive semidefiniteness but no Fiedler value, so λ_2 is an obligation in the standing of κ, computed in `Examples/Phase4.lean` §34 for the complete graph and for no other. Existence of a locked configuration is assumed; the critical coupling appears nowhere in the derivation. The quarter-turn confinement is assumed and is not implied by locking. And the detuning enters in the population's ℓ^2 norm, which is extensive, so a population whose frequency spread grows with its size pays for that growth through the data: what is removed is the unconditional factor, not the physics of heterogeneity. The balance equation is the identical-frequency spatial model's, so reading Eq. (15) at the scalar threshold still needs the reduction of Section 12.4.

Witnesses run the third estimate on two oscillators with *different* natural frequencies locked at $\pm\pi/12$, every hypothesis discharged, the configuration provably not phase-locked and the resulting chord bound $\pi/4$, which is below the $\sqrt{2}$ the quarter-turn hypothesis gives for nothing and so constrains the pair (`chord_le_two` records the weaker comparison against the maximum chord); and the second on four sites covered by three patches in a line, bounding the pair that no patch contains at $3\sqrt{2-\sqrt{3}} \approx 1.553$, where the direct patch bound is not merely weak but unavailable (`Examples/Phase4.lean` §34–35).

Neither route weakens `SharedEncoder`. Two patches meeting at a site must still encode that site's quantity the same way, no amount of phase agreement supplies that, and it is as true across a dozen hops as across one.

The result is fenced on both sides. Two antiphase oscillators have $r^2 = 0$, computed through Eq. (8) rather than asserted, and at that coherence the two patches demonstrably disagree (`Examples/AgencyFoundations.lean`), so the derivation is not returning a vacuous zero residual. The bound is attained there: it gives $\sqrt{2\cdot 2} = 2$, which is the actual disagreement, so the factor N of Eq. (10) is not an artefact of the proof.

The shared-encoder obligation is substantive: `Examples/VectorContent.lean`, §23, supplies a site-dependent family that agrees at perfect locking, and a locked population whose reports disagree by a unit when that obligation is dropped. The encoders, uniform Lipschitz bound and common content coordinates are declared hardware data. Identifying these vector contents with neural variables, and testing whether their agreement matters for experience, remain the empirical work of Section 3.2's closing scope.

3.5 Phase offsets, winding and content

Twisted gluing does not recover content from phase frustration (`Phase5_TwistedGluing.lean`). We also tested the natural weakening of exact overlap agreement: local states may differ by a phase on each overlap. `gluesUpToPhase_of_isCoboundary` proves that coboundary offsets glue after one rotation per patch, while `isCoboundary_of_gluesUpToPhase` proves the converse when the phase action is free on overlaps. Any Kuramoto configuration, including a twisted or splay state, supplies absolute patch phases ϕ_i, so its offsets $\phi_i - \phi_j$ are automatically a coboundary (`isCoboundary_of_phaseField`) and the twisted family glues (`gluesUpToPhase_of_phaseField`). The formal `PhaseObstruction` quotient is deliberately not called Čech cohomology because its total cochains still assign irrelevant values to patch pairs with empty overlap.

The data give a narrower positive answer. (Townsend et al., 2015) identify phase singularities in multielectrode primate LFP recordings by an integer winding number around electrode loops, and the associated complex waves modulate spiking; (Xu et al., 2023) find task-dependent spiral centers and rotation directions in human fMRI phase maps. Around such a zero-amplitude point the circle-valued phase need not admit one real-valued lift. Local lifts differ on overlaps by integer multiples of 2π, supplying measured transition data without an arbitrary gauge field. The same integers are available from a state rather than from a recording: on a ring the q-fold winding's affine lifts drop q whole turns across the wrap and none elsewhere (`ringTransition`, `winding_succ`), which is what connects the obstruction below to the stationary states of Section 12.6. `Phase5_PhaseLifts.lean` formalizes the finite electrode loop: `loopWinding` sums its integer transitions, `loopWinding_eq_zero_of_hasGlobalLift` proves that transitions from one global lift telescope to zero, and `threeElectrodeWinding_not_hasGlobalLift` gives a winding-one witness. This does not rescue `PhaseObstruction`: that quotient uses circle-valued differences of absolute patch phases, while the measured obstruction uses integer transitions between local lifts.

The cover-level specification has cortical surface M, analytic signal $z = Ae^{i\phi}$, singularity set Z_t, punctured domain $M_t^* = M \setminus Z_t$, and measurement map $q_t = z/|z| : M_t^* \to S^1$. A trial-independent electrode or source-space cover of M_t^* supplies local real lifts; their differences divided by 2π are integer transitions on connected overlaps, and nerve loops are the inputs to `loopWinding`. In finite-resolution data the amplitude threshold defining Z_t, the loop geometry and the required cross-scale persistence are fixed using training or baseline data, never the content labels of test trials.

The proposed content-sensitive test compares nested preregistered classifiers. The control model contains stimulus, task, behavioural response, eye movement and arousal; the augmented model adds signed winding, source-localized singularity position, rotation direction and lifetime. Evaluation uses held-out participants and sessions, with label permutations within task, stimulus and arousal strata. The content hypothesis requires replicated incremental out-of-sample discrimination, including under a changed or absent motor report. This controls the known task dependence rather than mistaking it for content. A positive result would be an empirical association, not an identification of a topological sector with phenomenal content; a null increment rejects this proposed role. The construction and test remain outside the conditional chain, and a full nerve-indexed integer complex is deferred until the measurement protocol warrants it.

Lean 4 Implementation (`Phase8_ContinuousField.lean`): The continuous-field file defines $K_c = 2D$ and a predicate for a substrate above that threshold. Two cautions. First, the predicate is not connected to `is_continuous_kuramoto_trajectory`, `entropy_production_rate`, or the order parameter of `Phase4`. What it is connected to is the mean-field self-consistency equation $r = I_1(Kr/D)/I_0(Kr/D)$, whose bifurcation at $K = 2D$ we formalize in both directions (see the next paragraph): `exhibits_phase_transition_coherent` shows that a substrate satisfying the predicate admits a solution with $0 < r \leq 1$, and `exhibits_phase_transition_unique_coherent` that it admits exactly one. So the threshold is not a bare stipulation, but what is proved is a statement about a stationary density rather than about a trajectory. Second, the predicate requires the substrate's `volume` to be a probability measure, and this hypothesis is load-bearing rather than decorative. It compares `mean_field_coupling`, the kernel averaged over both arguments, against $2D$; on a substrate of total mass m that average carries a factor m^2, so without normalization the comparison tracks the substrate's size rather than its coupling and can be satisfied by any kernel at all. `Examples` §9 exhibits a system a factor of two below threshold whose unnormalized double integral nonetheless clears $2D$. Under the hypothesis, `mean_field_coupling_const` gives that a constant kernel has strength exactly K and `exhibits_phase_transition_const_iff` reduces the predicate to $K > 2D$, witnessed in §9 on both sides of the threshold.

3.6 Local communication and timely responsiveness

`Phase6_Locality.lean` defines a synchronous `Network`: each update uses the site's own state and messages from its declared neighbourhood. Its causal past $B_T(v)$ contains the sites whose initial states can reach v in at most T rounds. `run_eq_of_agree_on_ball` proves that initial configurations agreeing on $B_T(v)$ give the same state at v after T rounds. Locality follows from the update construction. The regional version uses the union of the sites' causal pasts.

A fresh intervention changes the initial state at w while leaving the baseline elsewhere fixed. If $w \notin B_T(v)$, all such interventions yield the same code at v. Reading a report from that code therefore cannot reconstruct two possible intervention values more than 2ε apart to tolerance ε (`not_reconstructs_of_outside_past`). This consumes the reconstruction inequality in Section 5.6. The report reads only the local state, and accuracy is required over the declared intervention family.

`Examples/Locality.lean`, §25, gives a three-site line: a value is unavailable at the reading site after one round and exact after two. A cut path never delivers it; a constant report fails responsiveness over separated values. A report correct for one fixed world remains possible. Thus the theorem limits guaranteed response by a deadline and leaves static gluing intact. Physical communication paths, round durations and relevant deadlines require a separate model of the implementation.

4 Content bounds beyond the uniform patch estimate

What Eq. (5) of the main article bounds is a chord, which never exceeds 2, so the estimate constrains anything only where $\sqrt{2N(1-r^2)}$ falls below that. The winding sweep of Section 8.3 fixes both ends of that range on one sheet. Read on the sweep's own patch, 625 sites at patch-local order 0.9015, the bound is $15.3\,L$ and constrains nothing; holding that order fixed, the count alone brings it below the maximum chord only at 10 sites or fewer. A nearest-neighbour pair of the same winding, the smallest patch there is, gives $0.0981\,L$. Shrinking patches therefore thins the overlaps at the same rate as it tightens the bound, and the two ends are not reached by the same cover. At a columnar grain the estimate constrains agreement inside one locally synchronous population and is vacuous across the distant territories whose agreement unity is about.

That vacuity is the uniform bound's; Section 3.4 derives the three results that decline the step putting N there, and this sheet fixes what each buys. Read on the patch above, where the uniform bound gives $15.3\,L$, the pair count (`chord_fraction_le`) allows at most a fraction 0.37 of pairs to disagree by more than L; a count names no pair, so a guarantee about a named overlap still needs the uniform bound at its own grain. The nerve chain (`chord_le_of_patch_walk_coherence`) reaches territories no single patch contains, but its accumulation is linear: at the sheet's nearest-neighbour value $\beta = 0.0981\,L$ it stays inside the maximum chord for 19 hops, which is 20 patch diameters and 0.312 mm of that sheet. The territories whose agreement the compatibility claim is about are separated by centimetres, more than an order of magnitude beyond where this chain saturates; the route reaches further only on a cover whose patches are larger or more coherent, and both are measurements. The spectral-gap route (`chord_le_of_frequency_locked`) reads no resultant and so takes no number from this sweep; its four declared inputs are stated with its derivation.

Structural connectivity is a declared graph and is what supplies λ_2. Measured functional connectivity supplies it no better: it is an estimate, so it inherits the decoder failure modes of Section 8, a shared prior shrinking estimates toward agreement and regional bias manufacturing disagreement, and feeding an estimate to a hypothesis strengthens no conclusion drawn from it. Using it to choose the cover is worse than neutral, being the move of Section 4 that empties a gluing claim, in different clothes.

5 Reflexive topology and the self

Self-representation is an additional relation between the glued global state and an encoding within it. The model specifies an encoding and readout whose composite fixes that state; contraction gives uniqueness. A dissipation bound imposes neither the encoding nor its accuracy.

Theorem S3 (Conditional Auto-Resonance): On a nonempty complete metric space of global sections, a self-prediction map has a unique fixed point if it is Lipschitz at the independently assumed rate $\Lambda = e^{-(K-K_c)\tau/2} < 1$. A state with reconstruction residual at most ε is within $\varepsilon/(1-\Lambda)$ of that fixed point. Coherence supplies $K > K_c$; it does not supply the Lipschitz law. To identify this fixed point with the section glued on a given cover, that section must itself be fixed by the map.

5.1 States and their metric

The construction in `Phase6_ReflexiveTopology.lean` applies Banach's theorem to a complete metric space of global sections. The three-site witness in `Examples`, §10, equips these sections with a metric derived from their spatial measures.

On a discrete substrate, every point has a smallest open neighbourhood. Restriction preserves point masses, so a germ at a site determines a mass there. The map `massEquivOn` identifies sections over an open set U with mass profiles on U: injectivity follows because restriction to a point determines the germ, and surjectivity because a finite measure realizes every profile on the finite substrate. The global case is `massEquiv`. This correspondence also constructs the patch-local sections used by the compatibility witnesses.

The distance between two global sections is the largest difference between their site masses. Completeness follows by constructing the section represented by the limit of a Cauchy sequence of profiles. On this finite substrate, the metric is equivalent to total variation within a factor of the number of sites, but is not equal to it. Distinct sections at distance $1/2$ witness its nontriviality (`exists_dist_lt_one`).

The metric matters to the fixed-point argument. Under the discrete $0/1$ metric, every contraction is constant (`contracting_implies_const`). The witness instead uses uniform distance between densities. It makes no claim of equivalence with the Lévy–Prokhorov metric: on a discrete substrate, thickening by $\varepsilon < 1$ leaves a set unchanged, and averaging a measure with a baseline need not contract that metric.

5.2 Encoding, readout and information loss

A `ReflexiveBoundary` carries an encoding into a local region, called the avatar, and a readout back to global sections. Its predictive map is their composite. This ties reconstruction to the information passing through that region. A map that reads the whole field cannot serve as a witness unless it factors through the stated local encoding.

The separate predicate `IsRestrictionResonance rb` requires the encoding to be presheaf restriction to the avatar region. Under this condition,

$$F(s) = \mathrm{readout}(s|_{\mathcal{A}}), \qquad s_* = \mathrm{readout}(s_*|_{\mathcal{A}})$$

at a fixed point. The second equality is `self_eq_readout_restrict`. The more general `self_mem_range_readout` places a fixed point in the range of the readout even without restriction resonance. These statements express what the local encoding reconstructs; restriction resonance is an assumption on the boundary, not a consequence of contraction.

Local readings can also determine a global state collectively. If a family of resonant avatar regions covers the substrate, equal readings at every avatar imply equal global states (`eq_of_avatar_eq`). The proof uses the same gluing uniqueness as Theorem S2. A single avatar can discard information away from its region, so collective reconstruction from a cover must be distinguished from reconstruction through one region at a fixed point.

5.3 The contraction hypothesis

For $\tau > 0$, the assumed factor is

$$\Lambda = \exp\left(-\frac{K - K_c}{2}\tau\right), \qquad K_c = 2D.$$

The lemmas `resonanceRate_lt_one` and `one_le_resonanceRate` split the parameter plane at the threshold. Above it, a map Lipschitz at this factor has a unique fixed point (`self_of_supercritical`). The bridge `self_of_coherent_order_parameter` obtains the required inequality from the existence of a coherent order parameter. At or below threshold, the specified factor fails the strict bound required by `ContractingWith` (`not_contractingWith_resonanceRate`); this argument is unavailable there.

The rate is motivated by the mean-field linearization near the incoherent state, $\dot{r} = \frac{1}{2}(K - K_c)r + O(r^3)$. Linearizing the cubic normal form about its coherent branch gives relaxation rate $K - K_c$, twice as fast. The assumed factor Λ is therefore larger than the coherent-branch relaxation factor. Both cross one at $K = K_c$. This comparison motivates a choice of bound; it does not derive a Lipschitz law for a neural encoding or readout from phase dynamics.

Banach's theorem establishes uniqueness for the specified map. E89 additionally requires the section glued on the specified cover to be fixed by that map, making it the unique fixed point. The theorem `chain`

uses this equality and the contraction law but does not require restriction resonance. The interpretation through $s|_{\mathcal{A}}$ needs that additional predicate.

5.4 Approximate reconstruction

Let $F(s_*) = s_*$ and $d(s, F(s)) \le \varepsilon$. The triangle inequality and the Lipschitz law give

$$d(s, s_*) \le d(s, F(s)) + d(F(s), F(s_*)) \le \varepsilon + \Lambda\, d(s, s_*),$$

hence $d(s, s_*) \le \varepsilon/(1-\Lambda)$. The theorem `ReflexiveBoundary.approximate_self_bound` applies Mathlib's `ContractingWith.dist_le_of_fixedPoint`. An estimate about a supplied fixed point requires neither completeness nor a separate existence proof; those hypotheses enter the existing Banach construction. The approximate states can form a set while the exact fixed point remains unique.

There is a stronger joint restriction when the *same* map reconstructs an entire relevant family. For any two members s, t reconstructed within ε, the triangle inequality and Λ-Lipschitz law give $(1-\Lambda)d(s,t) \le 2\varepsilon$, hence $d(s,t) \le 2\varepsilon/(1-\Lambda)$ (`Encoding.dist_le_of_contracting_reconstructs`). At zero error the family has at most one metric-distinct state. This theorem applies to a fixed map with a fixed readout. The code-separation bound below remains valid for any encoding, but it does not remove this diameter restriction when its composite is E89's contraction.

A readout conditioned on an input is a different interface, stated in `Phase6_ConditionedReconstruction.lean`. Each relevant state s carries a declared input u_s, the scene it is a state for, and the map applied to it is $F_u(s) = R(u, \Phi(s))$ (`ConditionedEncoding`). Suppose each F_u is a Λ-contraction on the family and $d(F_u(s), F_v(s)) \le L\, d(u, v)$. Splitting $d(s,t)$ at $F_{u_s}(s)$, $F_{u_s}(t)$ and $F_{u_t}(t)$ gives

$$d(s,t) \le \frac{2\varepsilon + L\, d(u_s, u_t)}{1-\Lambda}$$

(`dist_le_of_conditioned_contracting`): the family can be as broad as its inputs, and members given one input still obey the fixed-map limit (`dist_le_of_same_input`). On a complete space each F_u has a unique fixed point s_u (`selfState`). Mathlib's perturbation estimate for fixed points gives $d(s_u, s_v) \le L\, d(u,v)/(1-\Lambda)$ (`dist_selfState_le`), and each reconstructed member lies within $\varepsilon/(1-\Lambda)$ of s_{u_s} (`dist_selfState_input_le`). These are the statements a self that follows a scene would satisfy.

The interface has a degenerate solution. A readout that returns a state stored for each input and ignores the code (`storedReadout`) is a contraction with $\Lambda = 0$ and reconstructs exactly every family in which the state is a function of the input (`stored_reconstructs_iff`); it encodes nothing. The code criterion therefore binds within an input fibre and nowhere else: members of one fibre pairwise more than 2ε apart need distinct codes (`card_le_card_codes_fibre`), and a constant encoder reconstructs only families whose fibres have diameter at most 2ε (`dist_le_of_encode_const`). A claim that a system reconstructs its own state rather than its scene is a claim about a family with several separated states per input. `Examples/ConditionedReconstruction.lean`, §36, exhibits both halves on four real states in two scenes, $\{0, 3\}$ and $\{10, 13\}$. A one-bit code with a scene-conditioned readout reconstructs all four within $3/4$ by maps that halve distances on the family. No fixed $\frac{1}{2}$-contraction can, since the family's diameter is 13 and the fixed-map limit is 3 (`no_fixed_map_reconstructs_scenes`); each scene attains that limit, needs both codes, and defeats every constant encoder (`no_constant_code_reconstructs_scenes`). An affine readout that contracts the whole line has self-states $u + 3/2$, which attain the fixed-point bound (`sceneAffine_selfState_dist`). The input, its metric, L and Λ are declared, and no physical readout or cortical scene variable is shown to meet the hypotheses.

If $d(t, s) \le \delta$, the same fixed map satisfies

$$d(t, F(t)) \le \varepsilon + (1+\Lambda)\delta,$$
$$d(t, s_*) \le \frac{\varepsilon + (1+\Lambda)\delta}{1-\Lambda}.$$

These are `prediction_residual_of_dist_le` and `approximate_self_bound_of_dist_le`. The errors add before division by the contraction gap, whose inverse amplifies the bound near threshold. The displacement bound is supplied; this result does not derive approximate gluing or weaken Theorem S2's overlap hypothesis.

For $K > 2D$, $\tau > 0$ and $\varepsilon \ge 0$, `approximate_self_bound_of_supercritical` instantiates the estimate at the prescribed rate, and `resonance_error_radius_bounds` proves

$$\frac{2\varepsilon}{(K-2D)\tau} \;\le\; \frac{\varepsilon}{1 - e^{-(K-2D)\tau/2}} \;\le\; \frac{2\varepsilon}{(K-2D)\tau} + \varepsilon.$$

The proof uses $1 - e^{-x} \le x$ and $x/(1+x) \le 1 - e^{-x}$ for $x = (K - 2D)\tau/2 > 0$, both consequences of $1 + x \le e^x$. Thus the radius has leading behaviour $2\varepsilon/((K - 2D)\tau)$ and diverges as $K \downarrow 2D$ at fixed positive ε and τ. This concerns the tolerance bound, not an actual state's error. Zero residual still gives zero distance at every supercritical coupling. The map, metric and Lipschitz law remain specified inputs; no result here models a changing readout or constructs a biological encoding.

5.5 A nonconstant witness and its controls

On the three-site witness, the avatar is one site. The encoding reads its mass (`avatarRead`), and `avatarReadout` reconstructs a profile by averaging the reading with a baseline. The resulting map is nonconstant, has a unique fixed point, and contracts by one half of the difference visible at the avatar (`cortexPredict_not_const`, `cortexPredict_fixed_unique`, `cortexPredict_dist`). Differences elsewhere are invisible to this map. The whole-field averaging map `relax` therefore cannot replace it as a witness for a one-site encoding.

The witness takes $K = 3$, $D = 1$ and $\tau = 2\ln 2$, so $\Lambda = 1/2$ (`cortexResonanceRate`). These are parameters chosen for the example, not calibrated cortical values. The boundary `cortexReflexive` performs restriction and satisfies resonance. Its avatar distinguishes the silent field from the baseline, as shown by `cortexReflexive_restrict_ne`. One avatar per site gives a covering family and reconstruction from the three readings (`cortex_eq_of_avatar_eq`). The same map at $K = 1$ exercises the failure of contraction at the prescribed subcritical factor (`cortexSubcritical_not_contracting`).

Replacing the encoding by a constant makes the composite constant. It still has a unique fixed point, namely the readout of that constant (`constResonance_existsUnique_self`). In the blinded control, this fixed point differs from the original baseline (`cortexBlind_self_ne`), and distinct states have identical blinded readings (`cortexBlind_not_determined`). A constant encoding fails restriction resonance whenever two states differ on the avatar region (`constResonance_not_isRestrictionResonance`). The control thus distinguishes fixed-point uniqueness from sensitivity to the encoded state.

The approximate bound is sharp on the same one-site-avatar witness. The silent profile has $d(s, F(s)) = 1$ and $d(s, s_*) = 2$, attaining $\varepsilon/(1 - \Lambda)$ at $\varepsilon = 1$, $\Lambda = 1/2$ (`cortexSilent_predict ion_residual`, `cortex_approximate_self_sharp`). At $\Lambda = 1$, `cortexIdentity` instead takes the whole substrate as its avatar and uses identity encoding and readout. It is nonexpansive, and the silent and baseline profiles are distinct fixed points (`cortex_nonexpansive_not_unique`). This control establishes the failure of nonexpansive uniqueness on the same metric; it is not a model of compression into a proper local region.

5.6 Bounded-error reconstruction over relevant macrostates

`Phase6_Reconstruction.lean` separates the data of an `Encoding`—a relevant state family, encoder $E : S \to C$ and readout $R : C \to S$—from the predicate `Reconstructs` ε: $d(s, R(E(s))) \le \varepsilon$ for every relevant s. If $E(s) = E(t)$, the triangle inequality gives

$$d(s, t) \le d(s, R(E(s))) + d(t, R(E(t))). \tag{16}$$

`dist_le_add_error` proves the sum bound and `half_dist_le_max_error` puts at least half the separation into one of the two errors. Consequently the encoder is injective on any reconstructed finite family pairwise more than 2ε apart. If it has M members and the code alphabet has at most 2^b elements, then $M \le 2^b$ (`Encoding.card_le_two_pow`).

For a `ReflexiveBoundary`, take $E =$ `auto_resonance` and $R =$ `readout`. Their composite is `predict`, so `dist_le_prediction_residuals` constrains the existing self-prediction map. No contraction, completeness or restriction-resonance premise is used in this inequality. A constant encoder or constant readout fails reconstruction on two relevant states separated by more than 2ε, despite the composite having a unique fixed point.

`Examples/Reconstruction.lean`, §24, reconstructs the three-site witness's two extreme states within tolerance one using a nonconstant encoder and a contracting map. Its blinded boundary fails at tolerance $1/2$. A separate three-state example rejects a two-code alphabet at that tolerance and succeeds with three codes. These are additional operational criteria on candidate implementations, not consequences of `Self` or `UnifiedSelf`. The relevant macrostates, metric and fidelity are declared, and b counts distinguishable codes; physical bits, channel capacity and delivery times need further hardware assumptions.

5.7 Measurement obligations for a digital candidate

The results of the article's Section on unconditional limits hold for any candidate, and for a large language model running on graphics processing units four of them are checkable rather than merely stateable. They fix what a claim of unity or self-representation for such a system commits its maker to measuring; none of them excludes a digital realization.

A fixed synthetic task, reported with the spectral bound in Section 5.9, exercises these obligations on linear, attention and phase-code decoders under declared CPU fits. Its results rank no architecture family and supply no hardware-energy or experiential claim.

Codes, not reports. Reconstruction across a declared family requires the encoder to distinguish states more than 2ε apart when the target error is ε (Eq. (9)). Verbal reports alone do not establish that condition. For a proposed self-model, declare the family and fidelity, identify the internal code under intervention, then test whether it separates the family. A token sequence is one possible code; an activation region is another and requires its own measurement.

The region, not the readout. If the code is a restriction of the candidate's state to a region, that region must distinguish the relevant family before a readout is chosen. In one forward pass of a decoder-only transformer, the causal mask restricts position t to positions at or before t; an intervention outside that causal past cannot be reconstructed from its activation alone. Across passes, prior output adds a separate channel. The bounded-fan-in communication result of Section 5 of the article applies to a declared physical interconnect and deadline; a full-support attention graph does not satisfy its sparse-neighbourhood premise. The graph bounds and their hardware scope are given with the locality results of Section 3.6.

Agreement about what. Unity requires the candidate's local descriptions to agree on independently identified shared quantities (Section 4 of the article). Common timing, a shared residual stream and fluent output supply none of that, for the reason the article's Section 2 gives: coordination bounds content disagreement only through encoders already related on overlaps. Which quantities two parts of a model describe in common is a decoding question, carrying the failure modes of Section 8 of the article. The cover is a choice, and choosing it to secure agreement empties the claim.

Physical scope. Digital representations can satisfy compatibility and reconstruction. The coupling-gated route additionally requires a calibrated physical coupling, its noise regime and an independently identified relation from phase order to shared content. An attention matrix or a shared clock supplies none of these; the cortical proposal carries the same physical obligations. Finite-resolution capacity depends on code-space volume, resolution and readout gain, so continuity alone favours no substrate. The capacity and hardware-support comparisons appear in Table S1. Neither the installed-energy condition nor those comparisons exclude a digital realization.

5.8 Information carried by a finite token output

For a finite joint probability law μ on an input state X and output alphabet V, `mutualInfo_le_output_entropy` proves $I(X;V) \le H(V)$ using the same KL-based mutual information as the thermodynamic information results. On an atom of positive joint mass, $p(x,v) \le p_X(x)$ gives $\log[p(x,v)/(p_X(x)p_V(v))] \le -\log p_V(v)$. Summing proves the inequality; zero joint atoms contribute zero, so neither strict positivity nor deterministic emission is needed. The finite entropy bound then gives $I(X;V) \le \log|V|$ nats, or $\log_2|V|$ bits (`mutualInfo_bits_le_log_card`).

`token_information_bits_le` applies this to the whole sequence space V^k, giving $k\log_2|V|$ without assuming independence between emissions. `no_tokens_no_information` proves exact zero at $k = 0$. Checked binary laws attain the one-bit ceiling for a perfect copy and zero for an independent output with identical marginals. The input space is finite in these statements. They count a specified output interface, not a model's internal activation or cache capacity. Additional accessible state, timing or a variable-length termination signal changes the output variable and must be counted explicitly.

The finite-resolution packing result in Table S1 applies the same logic to a continuous code space. At a fixed resolution its distinguishable codes form a finite alphabet; unbounded real precision would be needed for an unbounded capacity claim. This framework attributes no computational power beyond ordinary finite-precision computation to an analog medium.

5.9 Spectral error through a linear code space

Let a finite spatial coupling operator K have orthonormal eigenmodes e_i with $|\lambda_1| \geq \cdots \geq |\lambda_N|$. Every linear approximation A of rank at most d satisfies

$$\sum_{i=1}^{N} \|Ke_i - Ae_i\|^2 \geq \sum_{i=d+1}^{N} \lambda_i^2. \tag{17}$$

Retaining the first d modes attains equality. A linear encoder and readout factoring through d real coordinates therefore pay this error. Applying the bound requires a spectrum for the target and a rank constraint on the actual approximation. Query/key dimension alone supplies no such constraint on softmax attention: even a two-position rank-one logit matrix can yield rank-two attention. The result prices an explicitly linear bottleneck in a finite representation of a spatial kernel; it gives no universal gap between transformers and continuous fields. The derivation and the softmax counterexample follow.

`Phase6_AttentionRank.lean` proves the finite-dimensional bound in Eq. (17). It accepts an orthonormal eigenbasis $Ke_i = \lambda_i e_i$ and bounds the error of an arbitrary linear map A; A need not preserve the eigenmodes. Let P project onto range A and set $p_i = \|Pe_i\|^2$. Parseval gives $0 \leq p_i \leq 1$ and $\sum_i p_i = \operatorname{rank} A \leq d$. Projection onto the orthogonal complement of the range then gives

$$\|Ke_i - Ae_i\|^2 \geq \lambda_i^2(1 - p_i).$$

For the d largest squared eigenvalues, the unit-capacity weights p_i retain at most their sum: weight outside these modes can be exchanged for unused weight inside without decreasing the retained sum. Summing yields `spectral_tail_le_error`. The ordered-index corollary is `spectral_tail_le_error_of_antitone`. The explicit truncation has range dimension at most d and exactly the tail error (`spectralTruncation_rank`, `spectralTruncation_error`). A two-mode witness with eigenvalues three and one pays exactly one at rank one, and every rank-one competitor pays at least one.

In `Phase6_Reconstruction.lean`, `linear_bottleneck_error` derives the rank hypothesis when $A = \mathcal{D}\mathcal{E}$ factors through a code space of dimension at most d. This is a bound on summed squared errors over basis inputs, not a worst-case reconstruction theorem or a finite-bit memory bound. The target spectrum and finite spatial representation are inputs; no Mercer expansion or physical kernel spectrum is constructed.

The rank budget cannot be inferred from query/key dimension after softmax. The logits $L = \left(\begin{smallmatrix} 0 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ come from scalar queries and keys $(0, 1)$. They have rank one, but row softmax gives

$$\begin{pmatrix} 1/2 & 1/2 \\ 1/(1+e) & e/(1+e) \end{pmatrix},$$

whose determinant is $(e-1)/(2(1+e)) > 0$. `softmax_rank_counterexample` proves its rank is two. The spectral bound therefore applies to an actual linear rank constraint, including a declared factorization, and supplies no universal softmax-attention exclusion. An effective numerical rank additionally needs a specified error tolerance.

For a supplied input second-moment factor $C = RR^*$, expected squared error of a linear decoder A is the squared Frobenius norm of $(K - A)R$. `covariance_root_singular_tail_le_error` (`Phase6_CovarianceRank.lean`) bounds every rank-limited A by the omitted squared singular values of KR; the theorem takes the singular bases and factor as hypotheses. The diagonal-weight case is `input_weighted_spectral_tail_le_error`, and `Phase6_SampledRank.lean` realizes its floor for a declared eigenbasis-sampling input law. None of these statements estimates a distribution from data, constrains a nonlinear decoder or promises transfer when the test family's input law differs from training.

The CPU comparison fixes a four-coordinate binary global state, two independently observed three-coordinate local views with one shared scene and one shared self coordinate, and a target containing the global coordinates and their declared self-relative relation. The code region is the actual internal code: the linear bottleneck, the last causal attention token, or the readout node of a phase graph. The attention control has two tokens, width four, one head and query/key dimension one; the phase control has a two-node coupling graph and fixed communication rounds. The saved artifacts contain training and held-out state combinations, seeds, fitted matrices and nonlinear parameters, actual train and test codes, effective code rank at a prespecified relative singular-value tolerance, output size, causal past, rounds, parameter count, dense storage and CPU times. An independent report recomputes the scores and ranks from those artifacts without fitting.

On an even-parity training and odd-parity held-out split, the rank-two linear decoder pays a predicted input-weighted training floor of 2.000 and a held-out error of 2.000 in mean summed squared target error. A second, prospectively fixed split trains where the scene coordinates agree and tests where they disagree. In the saved seed-24 noiseless run, the full linear map has training error 0.000 but held-out error 2.667: the training covariance lacks a direction needed by the test family. The declared task map, supplied independently of the restricted fit, has noiseless held-out error 0.000. Thus the fitted map's error measures failure to identify that direction from the training family; the task itself admits exact linear reconstruction. The attention and two-node phase candidates have held-out errors 3.648 and 3.662 under the declared short CPU training budget. A protocol-only commit fixed a separate confirmation batch's seeds, scoring rule and source hashes before that batch was run; its saved artifacts passed the audits and its declared full-linear distribution-shift criterion at every noiseless confirmation seed. The noiseless training and test arrays are identical across those seeds, so the deterministic linear fit repeats one numerical calculation; varying nonlinear initializations probes only fitting variation on those same states. The noisy confirmation run changes its sampled observations. These figures assess the fitted candidates; they neither rank architecture families nor measure hardware energy. Independent noise, local-view disagreement, and shifts to the encoded self coordinate holding the scene fixed are scored separately. A shuffled code and an incompatible target readout check whether a candidate's code supports the intended decoder. Equal resultant magnitude with different phase codes checks that order alone does not identify content.

A three-node chain on the same task puts one local view two graph edges from the encoding node and a shared coordinate one edge away. A synchronous nearest-neighbour phase update carries the one-edge perturbation to the code after one round and the two-edge perturbation after two; before those deadlines the respective code is exactly unchanged. Readouts are fitted on the training family at each deadline and scored on the same held-out family. This is a causal-past control, not an inference from an attention label or a measured physical communication network. The toy task has no independent experience endpoint, and its digital and phase codes establish no substrate preference.

5.10 When restriction resonance fails

Restriction resonance — that the encoding a region carries is the global state's own restriction to that region — is a premise of how the fixed point is read, not a consequence of contraction. Stated on bare maps it is agreement between an encoder and a declared reference, `Resonates` E ρ, and `isRestrictionResonance_iff_resonates` records that the boundary's predicate has that shape with $E =$`auto_resonance` and $\rho =$`restrictToAvatar`.

Its failure has a criterion that evaluates neither map. Resonance forces the two to confuse the same pairs of states, so a pair the encoder identifies and the reference separates refutes it (`not_resonates_of_confuses`), and so does a pair the encoder separates and the reference identifies (`not_resonates_of_splits`). The first, applied to a boundary whose encoder happens to be constant, gives `not_isRestrictionResonance_of_const_avatar`. The constant control above refutes resonance for a boundary built by blinding another one; this refutes it for any boundary at all, from a hypothesis about its encoder rather than about its construction.

The second combines with accuracy, and it is the half that puts a condition on the substrate rather than on the model. Reconstructing two relevant states more than 2ε apart to within ε forces their codes to differ, by Eq. (16); so a boundary whose avatar region reads those two states the same way is not resonant (`not_isRestrictionResonance_of_reconstructs`). Read forwards: a resonant boundary that reconstructs a pairwise separated family restricts injectively on it (`restrictToAvatar_injOn_of_isRestrictionResonance`). The resolution demanded is the region's. Which states a region reads identically is fixed by where the region sits, before any encoding is chosen, so no read-out repairs it, and E89 — which supplies a Lipschitz law and identifies the glued section with a fixed point — supplies no such resolution. The implication runs one way: a region that does resolve the family carries no resonant avatar by that fact, and nothing here constructs one.

`Examples/Reconstruction.lean`, §24, exhibits both sides on the one-site avatar. The baseline and the uniform profile at the same level agree at the shared site and differ by two elsewhere, so the resonant boundary of §10 does not reconstruct that pair at tolerance 1/2 (`cortexReflexive_not_reconstructs_offAvatar`); a boundary over the same region whose encoder names the state reconstructs the pair exactly and fails resonance (`cortexCheat_reconstructs`, `cortexCheat_not_resonant`).

The deadline version is the same criterion with the causal past of Section 3.6 supplying the confusion, and it needs no metric, tolerance or report. A reading region whose causal past excludes an intervened

site holds the same state in both worlds (`regionReading_eq_of_outside_past`), so its contents are not any declared reference that tells the two worlds apart (`not_resonates_regionReading_of_outside_past`). On the three-site line the region fails that reference at round one and meets it at round two (`no_resonance_at_one`, `resonance_at_two`).

5.11 What a biological interpretation must add

The readout, its contraction law and the restriction-resonance condition are choices or hypotheses of a model. No field dynamics here produces them. The squared-drift objective is absent from the fixed-point result, so its minimizer is not identified with the self-model's fixed point.

The article's cup example illustrates the further content assignment: the represented state would include the scene and the perceiver's relation to it, and the encoding would need a demonstrated role in representing that relation. The three-site witness establishes mathematical realizability, not such a role. A nonconstant reconstruction and a unique fixed point alone do not establish a perspective or experience. Interventions on a candidate neural encoding would have to test the functional interpretation separately from the equality defining its fixed point.

6 Claims and their Lean identifiers

The table below is the claim-by-claim map, kept here rather than in the main text. It is here because it is a reference rather than an argument: each row names the Lean identifier that carries a claim and states the scope limitation of that identifier, and no step of the main text depends on the reader working through it. Figure 1 in the main text separates the three requirements and their inputs; Table S3 records the eight connecting hypotheses of the formal composition, and the article's Table 1 the four it argues from. The present table records node-level results and their implementation scope. The cover's compatibility is a supplied field. The boundary's restriction-resonance condition is a separate predicate, needed for the interpretation of self-encoding as restriction but not consumed by `chain`.

Table S1: Mapping of physical claims to Lean 4 formalization status. The development declares **no axioms**: `#print axioms` on any result in the development reports only `propext`, `Classical.choice` and `Quot.sound`, which `Audit.lean` verifies of every declaration the library adds to the environment, on every build. A physical postulate is an obligation on each instance rather than a global assertion, which is the shape a standalone axiom cannot have (Section 22); two of the postulates a reader might expect to find are not postulates at all, but consequences of the local data (Section 3). "Theorem" still does not mean "physically established": the postulates live in hypotheses and instance obligations, and every such class is inhabited by a witness in `Examples`, so no row is vacuous.

Conceptual Claim	Lean 4 Implementation	Status
Continuing finite agent	`FiniteProtocol` sequences different operations on one evolving law. `ContinuingProcess` bounds total heat and joint entropy reduction by a declared work store plus the system's mean energy drop. `ContinuingAgent` executes actuation, observation-driven learning and reset; `Examples/ContinuingAgent.lean` proves persistent above-chance reward, positive recurring cost, a funded first cycle and exclusion of 22 cycles at its declared allowance. Blinding removes the reward advantage; omitting reset reduces later reward. `Chain.activeBound_of_continuing` derives stage budgets when every stage has nonnegative mean heat. The store constrains expectations at every prefix, not individual paths; the prior, store and composite-channel reservoir model are supplied. The measurement's own register is carried separately below; battery dynamics, fabrication, register-ledger identification and cortical interpretation remain separate	Theorem (conditional)

Conceptual Claim	Lean 4 Implementation	Status
A store on the trajectory	`PathwiseStore` (`Phase3_ContinuingAgent`) carries the store's reading as a state coordinate, with a draw and a declared supply per transition, so solvency is a claim about the states the protocol reaches. `solvent_of_funded` derives it from a stagewise funding condition through `exists_pred_of_reachable`, with no positivity and a possibly random draw. `mean_balance_eq` averages the same ledger and `totalDraw_le_of_solvent` derives the expected bound, one way only. `solvent_forall_of_replenished` and `horizon_le_of_net_cost` are the sufficient and necessary fences on sustained operation. `Examples/PathwiseStore.lean` witnesses a store whose expected allowance never runs out and which a path of probability 1/4 overdraws, and a deterministic spend-and-recharge cycle solvent at every horizon whose cumulative draw exceeds any allowance and whose supply is load-bearing. The generic supply is an input; finite-source accounting and the measurement's register appear below. Initial-law preparation remains outside the model	Theorem (conditional)
Finite supply and delivery	`PathwiseStore.SourceLedgered` identifies supply with a source coordinate's energy loss on the executed paths. `withSource_ledgered` cancels internal replenishment from the combined resource boundary; `totalDraw_le_initial_resources` bounds draw by the initial combined mean, and `horizon_le_of_finite_source` bounds a uniformly positive-cost horizon. `Examples/FiniteSupply.lean` discharges the identities with energy-conserving reversible gates on source, buffer and load bits, controlled by a uniformly prepared selector. Half the paths deliver a unit and half leave it in the source; an empty source delivers nothing. A forced unit-draw ledger fails, and the unbounded charger admits no nonnegative finite source. Initial preparation, gate control, external refuelling and identification with the learning agent's actuator remain separate	Theorem (conditional)
A sensor memory and its erasure	`memoryHeat` is the mean log-ratio heat of a memory channel on a declared memory law. `memoryHeat_const` (`Phase3_SensorMemory`) proves that an erasure landing on one law ν costs $k_B T[H(\mu) - H(\nu)] + k_B T\, D_{KL}(\mu\|\nu)$, with `memoryHeat_ge_entropy_drop` the Gibbs corollary and `memoryHeat_self` the matched case. `MemoryAgent` runs act, record, learn and clear four-periodically on $R \times (M \times Env)$ as a `FiniteProtocol`; `learn_register_independent` proves the register reads the memory alone, `memoryLaw_clear` makes the memory autonomous under the erasure while `memoryLaw_record` does not, `clear_stageHeat_eq` allocates the stage's heat to its three coordinates, and `clear_memoryHeat_const` is the identity on the agent's executed paths. `Examples/SensorMemory.lean` bounds the memory's mass in $[7/16, 9/16]$ at every clearing from channel masses alone, giving a cost between $(3/16)\log 3$ and $(5/16)\log 3$ whatever has been learned, and $(1/4)\log 3 = [(3/4)\log 3 - \log 2] + (1/2)\log(4/3)$ in the first cycle. A blind measurement pays the same, a drift in place of the erasure pays nothing and forgets nothing, a matched memory is free and an over-ordered one absorbs $(1/4)\log 3$. Gate fabrication remains separate; prior preparation is treated below	Theorem (conditional)
A funded clearing	`PathwiseStore.balance_le_of_stage_cost` generalizes the reading bound to a stage-dependent cost, with the constant case derived from it; `sum_period_indicator` and `horizon_le_of_periodic_cost` give $mc \le b$ for a charge paid once per period, and `periodic_horizon_le_of_finite_source` states it with the source inside the boundary. `net_draw_eq_zero_of_positive` proves that full support admits no net draw at all. `MemoryAgent.clearings_le_of_source` applies this to the clear stage of the sensor-memory agent. `Examples/FundedMemory.lean` carries a charge of three units in the world the same agent acts on, spends one at each clearing (`charge_invariant`, `clear_spends_one`), attains the bound at three cycles (`cycles_le_three`) and reads below empty at the sixteenth operation (`not_solvent_four`), while every clearing still costs between $(3/16)\log 3$ and $(5/16)\log 3$. `clear_unfunded` records that an erasure reading only the memory cannot be conditioned on the charge, so the unfunded run continues; `not_positive` derives the agent's restricted support. The drawn unit is not identified with the stage's heat. Charge preparation, external refuelling and gate fabrication remain separate	Theorem (conditional)

Conceptual Claim	Lean 4 Implementation	Status
Preparing the prior	A constant preparation channel drives the hardware to a positive law without reading it, which is the hypothesis of `memoryHeat_const`, so the same heat identity prices it. `KL_uniform` (`Phase3_Preparation`) gives $D_{KL}(\mu\|\text{unif}) = \log\|V\| - H(\mu)$ with no positivity of μ, so `memoryHeat_uniform` is zero at every starting law: a uniform target has zero mean log-ratio heat. `FiniteProtocol.withPreparation` runs it as stage zero, `withPreparation_law_succ` identifies the shifted run with the original and `preparation_meanHeat` is its cost. `preparation_parameter_marginal` and `preparation_product` are the two fences: the parameter's marginal is untouched and a constant preparation creates no parameter-hardware correlation. `Examples/Preparation.lean` prepares the sensor-memory agent's own prior with zero mean heat, charges the whole run through the existing bound, prices the agent's standard state at $(1/4)\log 3$ from uniform, evaluates a bias-ε target at $(1/2-\varepsilon)\log[(1-\varepsilon)/\varepsilon]$ and proves this family unbounded at thermal scale one. The witness changes the initial hardware law, rejects preparing parameter uncertainty or correlation, and rejects the sharp target from the positivity premise. Funding, exact preparation and gate implementation remain separate	Theorem (conditional)
Resource bounds on a trajectory	`ResourceTrajectory` (`Phase3_ResourceFoundations`) telescopes the ledger $\beta_{n+1} = \beta_n - w_n^- + w_n^+$ on an *arbitrary* state type: `balance_sum`, `draw_le_resources` and `horizon_bound` use no `Fintype`, no probability law and no channel, and `ContinuingProcess.horizon_le_of_cost` is that bound read on the sequence of expected allowances. `draw_eq_energy_heat_control` identifies the draw from a conserved boundary carrying stored, bath and installation/control energy; `draw_eq_step_heat` closes that identification on a named channel: the bath coordinate's gain is the mean heat of the step the stage runs, and local detailed balance makes the ledger's heat term $k_B T$ times *that step's own* `bathEntropy` rather than any real written as a logarithm. `draw_ge_entropy_reduction` is the resulting Landauer bound: the work drawn, net of stored and control energy, is at least $k_B T$ times the entropy the step removes, under support inclusion rather than full positivity. `transferCharge` is a reversible swap on $\mathbb{N}\times\mathbb{N}$ conserving the total, and `sharp_preparation_extendedKL_top` is why it cannot also *make* a sharp charge: a channel driving every state to one target has no reverse path from that target, so its path divergence is ∞ and no finite cost exists for it, while a channel stopping short of sharp is an ordinary priced step. `SourceDelivery` identifies delivery with a source coordinate's loss on an arbitrary state type; `horizon_le_of_sourced_delivery` adds the source's initial content to the store's in one horizon bound, and `no_finite_source_sustains` shows a run at strictly positive cost with a nonnegative source cannot continue indefinitely, so `refuel`'s sustaining hypothesis is the declaration of a supply outside the boundary. Gate fabrication and control work are charged to an observable and never priced, and no theorem connects these trajectories to a particular agent's channels	Theorem
Entropy balance without full support	`ProbDist.SupportIncluded` replaces strictly positive masses by the condition the expansion uses: every forward path of positive probability has positive reverse probability. `KL_nonneg_of_support`, `entropyProduction_nonneg_of_support` and `entropy_balance_of_support` (`Phase3_SupportedThermodynamics`) hold there, and `reversibleSupport_of_positive` places the stronger premise strictly inside the weaker one. A step whose initial law sits on one atom witnesses the gap (`Examples/AgencyFoundations.lean`). Where a forward path has no reverse, `ProbDist.extendedKL` returns ∞ in $[0,\infty]$ rather than a finite real cost, and agrees with the real divergence where support holds. `extendedKL_eq_klDiv` and `extendedKL_eq_klDiv_of_missing` (`Phase3_MeasureFeedback`) identify it with Mathlib's measure divergence on *both* branches, the finite one under a positive reference law and the infinite one wherever an atom is missing, so the finite convention is the measure-theoretic one restricted rather than a second definition beside it. Construction of a physical reverse protocol remains separate	Theorem

Conceptual Claim	Lean 4 Implementation	Status
Gluing over content variables	`Phase5_ContentCover` takes content variables as the base and makes each region a partial decoder whose domain is the variables it decodes within δ on a declared family (`DecodingSetup.domain`). With the discrete topology the presheaf of values is Mathlib's sheaf of functions, so `decode_glue_unique` is `sheaf_glue_unique` on this cover, and `decode_glue_value` identifies the glued section with the true content at $\delta = 0$. `decode_compatible` derives compatibility at 2δ from decodability; `ContentCover.select_close` and `ContentCover.select_dist` give selection for signed and vector contents, and `select_decode_near_value` puts it within δ of the content. `exists_exact_decoder_of_injOn` and `not_decodes_of_read_eq` bound the decoder class from both sides. `Examples/ContentCover.lean`, §37, attains the 2δ bound and glues exactly at zero offset. The tolerance is worst-case, idealizing a statistical criterion; the variables, tolerance and decoder class are declared; the measure representation and the phase-to-content bounds do not transfer, and nothing relates a region's phase order to what it decodes	Theorem
Contents driven by observation	`LocalContent` (`Phase5_ContentDynamics`) carries fields in a real normed vector space. `update_residual` bounds the next mismatch by $(1-\eta)\varepsilon + \eta\delta$; `run_residual_floor` gives $(1-\eta)^n\varepsilon + \delta$. Zero residual and coverage give a unique patch-selecting extension. `compatible_map` preserves agreement under a common nonexpansive linear map, with the converse rejected in `Examples/VectorContent.lean`, §22. Scalar fields are the $\mathcal{C} = \mathbb{R}$ instance; the mass-profile bridge additionally requires nonnegativity. Observation agreement is an input unless derived from the encoder hypotheses below; neural and probability-sheaf identifications remain separate	Theorem
Coherence bounds vector-content disagreement	`compatible_of_shared_coherence` (`Phase5_ContentDynamics`) uses the pointwise chord estimate `chord_sq_le_of_coherence`, drawn from the phasor-spread identity `sum_normSq_sub_order_parameter`, to give residual $L\sqrt{2N(1-r^2)}$. Site-dependent encoders must be equal as functions on every overlap (`SharedEncoder`) and uniformly Lipschitz. Exact locking gives exact agreement; `run_residual_of_shared_coherence` propagates a uniform residual bound through observation updates. A common scalar encoder is a special case. `Examples/VectorContent.lean`, §23, proves disagreement at perfect locking without the shared-encoder premise. The factor $\sqrt{N}$ is attained up to a constant by one antiphase oscillator, so the global bound is vacuous for large populations below near-perfect order; the bound assumes common content coordinates	Theorem (conditional)
How many pairs a bound can miss, and agreement across a connected cover	Two readings of the same pair sum $\sum_{ij}(1-\cos(\theta_i-\theta_j)) = N^2(1-r^2)$ (`sum_gap_eq`) that decline the step bounding the deviations of one pair by the whole spread. Markov's inequality on it (`card_gap_le`, `card_chord_le`) gives `chord_fraction_le`: the fraction of pairs separated by a chord of c or more is at most $2(1-r^2)/c^2$ at any population size, and the same sum read by rows gives `card_site_gap_le` and `site_fraction_le` for the sites far from a fixed fraction of the population. Separately, `patchNerve` makes two patches adjacent when they share a site, and `chord_le_of_patch_walk` chains a per-patch chord diameter β along a walk of k hops into $(k+1)\beta$ — $k+1$ because both endpoint sites pay a step inside their own patch — with β read off each patch's own resultant by `chord_le_of_patch_walk_coherence` and the content statement by `compatible_of_patch_nerve`, witnessed on a pair no patch contains (`Examples/Phase4.lean` §35). The counts name no pair and no site: the exceptional set they permit is unlocated and may be exactly the sites a cover's overlaps sit on, so a guarantee about a named overlap still needs the uniform bound. The chained bound supplies no walk — connectivity and diameter are properties of the declared cover — and its accumulation is linear, so on the sheet of Section 12.6 it reaches 20 patch diameters and not a hemisphere	Theorem

Conceptual Claim	Lean 4 Implementation	Status
Agreement read off the coupling graph rather than off the resultant	`is_frequency_locked` is the balance equation $\omega_i + \sum_j K_{ij} \sin(\theta_j - \theta_i) = \Omega$, equivalent to rigid rotation solving the dynamics (`is_frequency_locked_iff`) rather than an added postulate. `couplingForm` is the coupling-weighted Dirichlet form, identified with Mathlib's `lapMatrix` form on an unweighted graph (`couplingForm_eq_lapMatrix`), which is what anchors λ_2 to algebraic connectivity; `SpectralGap` carries the Rayleigh inequality as declared data and `SpectralGap.scale` makes the coupling strength visible in it. Symmetrizing (`sum_mul_coupling_sin`), Jordan's inequality termwise (`couplingForm_le_pairing`) and Cauchy–Schwarz give `spread_le_of_frequency_locked`, $4\lambda_2^2\|\theta - \bar{\theta}\|^2 \le \pi^2\|\omega - \Omega\|^2$, read at a named pair by `chord_le_of_frequency_locked` and as a content residual by `compatible_of_frequency_locked`, with no population count, no patch and no cover geometry in it. The quarter-turn hypothesis already puts every chord below $\sqrt{2}$, so the reading constrains the pair only where $\lambda_2 > \pi\|\omega - \Omega\|/2$. Four inputs are declared and none discharged: λ_2 is hardware data in the standing of κ, computed in `Examples/Phase4.lean` §34 for the complete graph alone; existence of a locked configuration is assumed and the critical coupling enters nowhere; the quarter-turn confinement is assumed and is not implied by locking; and the detuning enters in the population's ℓ^2 norm, so heterogeneity is paid for through the data rather than removed. The balance equation is the identical-frequency spatial model's, so reading it at the scalar threshold still needs the reduction of Section 12.4	Theorem (conditional)
Phase-derived task observations	`PhaseSensor.channel` (`Phase3_PhaseSensor`) computes symbol probabilities as fractions of a supplied phase population. `channel_ne_iff_fiberCount_ne` characterizes distinguishability. `Examples/PhaseSensor.lean`, §27, identifies the sensor-fed learner with the existing observational learner and gives an uninformative control with the same $r^2 = 1/4$ at every parameter and action. Only the informative case improves reward. Phase law, sampling and readout are inputs; no phase dynamics or coherence-driven learning advantage follows	Theorem (conditional)
Learning installs coupling	`Examples/LearnedCoupling.lean`, §28, puts the learner's law on the declared actuator. `installedEnergy_learned` gives $U_n = 3/4 - 2^{-n}/4$; `kernel_update_on_executed_step` and `work_learned` compute the kernel change and work on its paths. Cumulative work stays within $(1 + \log 3)/4$. The matched control stays at $U = 1/2$ with zero mean work and is excluded from the coherent scalar branch at $\kappa = 4$, $D = 1$. The learner clears the necessary energy line after one update; evolving phase feedback and physical identification remain open	Theorem (conditional)
Closed phase–register feedback	`Phase3_PhaseFeedback` composes installed phase dynamics, sensing and register update into one joint transition. `Phase3_FeedbackLedger` proves the executed-law installation and finite-store balances and conditional exhaustion. A binary witness uses phase-flip local detailed balance and compares informative, blind, replayed and frozen controls on held-out next-phase prediction. Sensor and reset work prices are declared; their reverse protocols, independent content decoders, physical replenishment and cortical calibration are not supplied	Theorem (conditional)

Conceptual Claim	**Lean 4 Implementation**	**Status**
Reconstruction requires distinguishable codes	`Phase6_Reconstruction`: equal codes force errors summing to at least the states' distance. `Encoding.dist_le_of_contracting_reconstructs` bounds any family reconstructed within ε by one fixed Λ-contraction to diameter $2\varepsilon/(1-\Lambda)$; a broad family needs the input-conditioned map of the next row. `Encoding.card_le_two_pow` gives $M \le 2^b$ for a reconstructed family pairwise more than 2ε apart. `ReflexiveBoundary.dist_le_prediction_residuals` applies the bound to the existing self-prediction map. `Examples/Reconstruction.lean`, §24, gives a contracting positive control and a blinded fixed-point rejection. The count generalizes from an alphabet to a metric: `Encoding.sub_two_mul_le_lipschitz_mul` transfers the state separation through an L-Lipschitz readout, `isSeparated_image_encode` makes the family's codes a δ-separated set whenever $\delta L \le d_{\min} - 2\varepsilon$, for a family separated pairwise by more than $d_{\min}$, and `encard_le_packingNumber_range` bounds the family by `Metric.packingNumber` of the codes the encoder writes; under resonance `encard_le_packingNumber_restrictToAvatar` eliminates the encoding and states the bound of the region itself. Mathlib supplies both notions, so this is an application rather than a construction. This is an additional criterion on a declared family, metric and tolerance; code count is not a physical memory model, the gain L is declared and derived for no physical readout, and an unbounded-gain decoder is constrained by neither form	Theorem
A self that follows a scene	`Phase6_ConditionedReconstruction` conditions the readout on a declared input. With each F_u a Λ-contraction on the family and F L-Lipschitz in the input, `dist_le_of_conditioned_contracting` bounds two members reconstructed within ε by $(2\varepsilon + L\,d(u_s, u_t))/(1-\Lambda)$, replacing the fixed-map diameter; `dist_selfState_le` moves the fixed points by at most $L\,d(u,v)/(1-\Lambda)$ and `dist_selfState_input_le` puts each member within $\varepsilon/(1-\Lambda)$ of its input's fixed point. `card_le_card_codes_fibre` and `dist_le_of_encode_const` make the code criterion bind within an input fibre, which excludes the stored-state readout (`storedReadout`) on any family with separated states per input. `Examples/ConditionedReconstruction.lean`, §36, gives a two-scene family no fixed contraction reconstructs and a whole-line contraction whose self-states attain the bound. The input, its metric, L and Λ are declared; no physical readout, cortical scene variable or derivation of Λ is supplied, and across fibres the input rather than the code carries the distinction	Theorem
Locality limits timely responsiveness	`Network.run_eq_of_agree_on_ball` (`Phase6_Locality`) derives causal-past indistinguishability from local synchronous updates. `not_reconstructs_of_outside_past` uses the reconstruction bound to exclude accurate reports across separated interventions outside that past by the specified round. `Examples/Locality.lean`, §25, checks a delay line, a cut path and coincidental accuracy for one world. Snapshot gluing remains possible; physical paths and deadlines are inputs	Theorem (conditional)
Thermodynamics and control off finite state	`MeasureFeedbackStep` (`Phase3_MeasureFeedback`) carries an initial probability measure and a Markov kernel on *any* measurable space, with entropy production the divergence of the forward path law from the time-reversed one valued in $[0, \infty]$. Nonnegativity is then structural rather than derived, `entropyProduction_eq_top` keeps absolute irreversibility as ∞ with no positivity or finiteness hypothesis, and `path_divergence_splits` is the chain rule: the divergence separates into a term in the two initial laws and a term in the two channels compared on the forward law. `entropyProduction_map_le` is data processing — coarse-graining by any measurable observable cannot increase production — and `initial_divergence_le_entropyProduction` the second law in divergence form. `MeasureControlProblem.exists_optimal` selects a maximizing action over a finite permitted set with no `Fintype` on the state at all, and `exists_optimal_compact` over a compact one, discharging continuity rather than replacing it by finiteness. Witnesses sit on $\mathbb{N}$, where a definite jump has infinite production and a stationary step none, and on $\mathbb{R}$, where the initial law is a Gaussian assigning zero to every singleton — a law no `ProbDist` can carry. The *Shannon-entropy* form of the balance is not generalized: it needs a reference measure, and differential entropy is not entropy	Theorem

Conceptual Claim	Lean 4 Implementation	Status
Finite goal-directed control	`FiniteControlProblem` (`Phase3_AgencyThermodynamics`) evaluates terminal reward and both substep costs on the same `FiniteFeedbackCycle`, with an explicit deterministic policy and shared world and memory channels. `exists_optimal` proves finite feasible maximization. `Examples/AgencyControl.lean` computes every deterministic bit policy's performance, heat and work and proves a unique feasible optimum that strictly improves a declared lamp task. A feasible baseline misses the target; a better-performing policy exceeds the budget; equal final laws can have unequal heat. Objective, reservoir model, initial resource and budget are supplied. This comparison covers fixed policies; adaptation of a policy register is treated separately below	Theorem (conditional)
Finite learning dynamics	`FiniteFeedbackStep.iterate` carries successive laws and telescopes entropy and first-law balances. `Examples/PolicyLearning.lean` proves strictly improving expected lamp reward, geometric convergence to a stationary mixture, positive update heat and bounded cumulative heat from the register's energy loss. Its energy encodes known task values. The limit is suboptimal; individual updates and a different initial preparation can lower reward. Task costs are prospective lamp task expectations. Preparing the register, executing task episodes, implementing the policy readout, cortical identification and the active chain's resource and spatial bridges require separate accounting	Theorem (conditional)
Learning from observed outcomes	`FiniteObservationalLearner` (`Phase3_ObservationalLearning`) holds an environmental parameter fixed and updates a policy register from the observed outcome of the action its readout executed. Neither the update nor the readout takes the parameter as an argument and the reward enters no channel, so nothing supplies the task's values. Act–observe–update composes into one `FiniteFeedbackStep`, carrying reward, heat and entropy on the same law. `Examples/ObservationalLearning.lean` gives every mass of that law at every horizon for a two-action task with an unknown rewarding action, hence strictly improving expected reward $3/4 - \frac{1}{4}(1/2)^n$ from a uniform prior, strict improvement over a frozen register and over an uninformative world under a common allowance, cumulative heat $(\log 3/4)[1 - (1/2)^N]$ and its equality with external work. `zero_work_needs_parameter_independent_heat` proves that a free relaxation would carry no information about the parameter, so the work is irreducible. The limit is below certainty, each action is executed half the time at every horizon, and a reward-decreasing transition keeps positive probability. Preparing the prior, supplying the work, task episodes beyond the declared updates, fabricating the readout, optimal-policy convergence and cortical identification are separate	Theorem (conditional)
Actuated scalar coupling	`ActuatedCoupling` fixes density $g[H(p) - H(p')]b(x)$ for a named feedback step. Regular mesh refinement proves convergence; `Chain.actuated_limit_le_budget` bounds its limit by $g(B/(k_B T)) \int b$ when the total profile is nonnegative. A noisy actuator has positive drive and changing spatial energies; stationary execution, changed gain and changed profile alter the result. The constitutive law is supplied and produces no kernel or microscopic actuator work account	Theorem + instance postulate
Installed kernel's mean-field strength	`KernelArrangement.continuumEnergy_eq_mean_field_coupling` (`Phase9_InstalledCoupling`) identifies the product integral with the iterated mean coupling by Fubini, when the measure and kernel agree. Compactness, finite mass and continuous profiles give integrability. `continuumEnergy_eq_modes` computes the same quantity as $\sum_{\Xi} p(\Xi) \sum_i c_i(\Xi)(\int \phi_i)^2$. Neither identity supplies the physical kernel identification; the threshold additionally requires a normalized substrate	Theorem

Conceptual Claim	Lean 4 Implementation	Status
Installed-energy threshold obstruction	`PricedArrangement` declares positive κ, nonnegative occupancies and $(\int \phi_i)^2 \le \kappa u_i$. `continuumEnergy_le_installedEnergy` proves $K \le \kappa U$ for the actual configuration law. `no_coherence_of_installedEnergy` excludes the coherent scalar branch and growing uniform-state modes at $\kappa U \le 2D$; `stationary_eq_incoherent_of_installedEnergy` identifies the stationary density as uniform. `Chain.installedEnergy_required_by_chain` derives $U > 2D/\kappa$ from E56 and E67. `Examples/InstalledCoupling.lean` witnesses both sides and the boundary, rejecting the converse, negative prices, an underpriced mode, a factor fitted to force exclusion, and signed occupancies. Installed energy is not heat; κ and field identification are supplied, and no heterogeneous-field dynamics or cortical identification is proved	Theorem (conditional)
Microscopic actuation and kernel refinement	`LocalActuator` (`Phase3_LocalActuator`) sets a product-space kernel $\sum_i c_i(\Xi)\phi_i(x)\phi_i(y)$ and a stored energy $\sum_i u_i c_i(\Xi)$ from the same mode occupancies. `kernel_update` and `kernel_update_local` give a transition's change and its spatial locality; `expected_kernel_update` averages those pathwise changes over the executed step, and `installation_first_law` charges the same paths with the stored-energy change plus the channel's heat. `KernelMesh` (`Phase2_KernelMesh`) reconstructs each sampled weight throughout its cell pair, `energy_eq_integral` identifies the cell-mass weighted energy with that reconstruction's product-measure integral, and `energy_tendsto` converges it under shrinking sample error on a compact substrate. `Chain.e45Active_of_localActuator` and `Chain.e45_of_localActuator` discharge the two spatial edges for an arrangement holding, respectively, a named step's final law and a declared predictive system's joint law; neither proof consults its premise. A thermal switch installs one mode for positive external work, changes a pairwise value, changes its coupling energy between the first two partitions of an atomless interval, and rejects a zero-work and a wrong-limit claim. Profiles, prices, readout, reservoir and substrate are declared hardware; fabrication, continuing power and cortical identification are not established	Theorem + instance postulate
One register's resource ledger	`RegisterLedger` (`Phase3_AgencyThermodynamics`) carries one register's update, its operations as elementary steps controlled by that register, local detailed balance for each at that register's temperature, and the identity that its dissipated heat is the total those operations deliver to its reservoir. `opHeat_nonneg_of_compressive` is the path-model second law at that temperature and `opHeat_le_dissipation` derives the allocation; `ledger_bathEntropy` exhibits the total as the register's own bath entropy change. `Examples/RegisterBudget.lean` witnesses it on the one-bit eraser with two operations delivering $(2/5)\log 2$ and $(3/5)\log 2$, and fences both inputs: a compressive operation above the whole dissipation, and a ledger whose other operation draws heat out of the reservoir. The identity and the compressiveness are physical inputs; sequencing the operations into one evolving law, a continuing energy source and cortical identification are separate models	Theorem (conditional)
A register's microscopic bath	`Examples/RegisterBath.lean` realizes each of the four maps of a register bit by a permutation of register and bath acting on a bath prepared pure, and computes heat as the bath's mean energy gain on the executed paths: $\log 2$ for the erasures, 0 for the injective maps, inhabiting Landauer's heat equation with the reachable bath supports. `sum_energyTransfer` and `reported_heat_eq` (`Phase3_ContinuingAgent`) give this ledger for any finite protocol with no positivity. Three results fence a map-indexed budget: the same gate run twice restores the prepared law, with shares $\log 2$ and $-\log 2$ about a total of zero; a drive that holds the register fixed excites its bath at $2\log 2$, so the map does not fix the cost, and the work is $\log 2$, $2\log 2$ and zero for the three runs; and a deterministic gate is not strictly positive, while an involution's log-ratio heat vanishes on every realized path though its bath gains $\log 2$. Positivity, local detailed balance and the ledger identity are therefore inputs rather than consequences of reversibility. The instance is local to its file; preparation, work supply, fabrication and identification of this bath with the ledger's reservoir are separate	Theorem (negative)

Conceptual Claim	Lean 4 Implementation	Status
Agency and feedback thermodynamics	`Phase3_Agency` composes policy, world, observation and memory-update channels. It proves the signed conditional-information identity, conditional nonnegativity under finite information, and passive recovery. `Phase3_FiniteInformation` connects finite laws to the same KL definition and proves finite-state finiteness. `Phase3_AgencyThermodynamics` derives the joint entropy balance from normalized positive path laws; local detailed balance supplies heat. Its `FiniteFeedbackCycle` constructs the swapped intermediate law, identifies the composed channel output and proves total entropy, heat, first-law and budget relations at a common thermal scale and energy observable. Reciprocal and biased two-bit witnesses check nontrivial updates and reject a mismatched intermediate law. `Chain.ActiveBound` carries the named elementary process's entropy, heat and budget into `chain_active`; the cycle bound supplies neither the register allocation nor coupling convergence. Extra registers, protocols, learning and cortical readouts require further models	Theorem (conditional)
Capacity of a finite phase space	The first premise, discharged by theorem rather than by footnote (`Phase1_PhaseSpaceCapacity`). $H(p) \leq \log \lvert X \rvert$ for any distribution on a finite state space, with equality exactly at the uniform law (`shannon_entropy_eq_log_card_iff`); a stream with $\lvert P \rvert^k > \lvert X \rvert$ histories cannot be recorded injectively (`absorb_not_injective`) and its entropy exceeds anything the system can hold; and on a *finite* space an unreachable state is an erasure, hence heat (`finite_phase_space_dissipates`). Finiteness is assumed, not derived, and is fenced — `Nat.succ` refutes the last step on an infinite space (`Examples` §19). The symmetry group and the metric are still consumed by no theorem	Theorem
Pointwise Vacuum Minimization	`pointwise_vacuum_of_global_min` (`Phase1`), potential energy defined as an integral; holds a.e., and a double-well witness proves the a.e. cannot be strengthened	Theorem
Landauer's Principle	Derived from statistical mechanics (`Phase3`); the axiom-free reversibility core becomes a heat bound via the field `StatisticalMechanics.heat_eq`	Theorem + instance postulate
Symmetry Breaking implies Boundaries	Derived from pointwise potential minimization (`Phase1`); witnessed on a double well with a $\mathbb{Z}_2$ symmetry whose minimization hypothesis is proved, not assumed	Theorem
Inevitability of Boundaries (π_0)	The domain-wall case of the Kibble mechanism (`Phase1` §3): on a connected substrate, a continuous field taking values in two different connected components of the vacuum manifold must leave the vacuum somewhere (`exists_notMem_of_connectedComponentIn`). Witnessed on the double well, whose vacuum manifold is computed to be exactly $\{-1, 1\}$ and proved disconnected (`Examples` §11.1); the field class is non-empty and the forced wall is located. Higher homotopy (π_1 vortices, π_2 monopoles) available and not taken	Theorem
Nonpredictive information bounds dissipation	Still, Sivak, Bell and Crooks' bound is the field `PredictiveDissipation.still_bound` (`Phase3_PredictiveThermodynamics`). Everything downstream is derived: mutual information is defined from Mathlib's `klDiv`; the data processing inequality for the chain $X_t \to S_t \to S_{t+1}$ proves the nonpredictive information non-negative (`predictiveInfo_le_mutualInfo`), from which $W_{\text{diss}} \geq 0$ follows rather than being assumed; a dissipation budget bounds wasted memory, and zero dissipation forces every retained bit to be predictive. Witnessed in `Examples` §18 by two systems over one two-bit law, one permitted to dissipate nothing and one forced to dissipate its whole memory, with a counterexample showing the no-back-action hypothesis is load-bearing. Descent of σ alone does not imply $D_{KL} \to 0$. On a finite register with a bipartite environment the postulate is discharged rather than assumed: `PredictiveDissipation.ofLandauer` (`Phase3_LandauerBridge`) builds a predictive structure whose thermal scale and dissipated work are the register's own and derives `still_bound` from `landauer_bound`, given the identification of wasted memory with erased entropy. That identification is the physical content and is not free — a reversible update erases nothing, so it forces the waste to zero (`nonpredictive_eq_zero_of_injective`) — and on the one-bit eraser the bound is attained (`Examples` §18.5)	Theorem + instance postulate
Structural Resonance KL Bound	Gibbs' inequality proved; the bound is the field `StructuralResonance.kl_bound` (`Phase3_KLBound`). The postulate is consistent and yields $\sigma \geq 0$. No limit may be read off it	Theorem + instance postulate

Conceptual Claim	**Lean 4 Implementation**	**Status**
Fixed-phase kernel descent	The bounded drift operator and its Hilbert-space adjoint give the squared-drift gradient. The theorem `structural_resonance_decreases_sigmaContinuum` (`Phase8`) proves descent for the specified flow with phases fixed. An atomless witness is supplied in `Examples`, §20. Neither joint dynamics nor thermodynamic entropy production is established	Theorem
Phase-locking minimizes squared drift	For a *reciprocal* coupling kernel, the phase-locked field has the least squared-drift objective of *any* field (`phase_locked_minimizes_entropy_of_symm`, `Phase8` §2a). The mean-matching condition is derived: the coupling term is antisymmetric under exchanging its two sites, so total drift is conserved across phase fields (`total_drift_eq_of_symm`). Reciprocity is necessary — a one-way coupling on two sites strictly beats its own phase-locked state (`Examples` §16)	Theorem
Electrodynamic Discretization	Lebesgue-Bochner edge weights (`Phase2`) sitting on faces of the complex and partitioning the region's stress-energy; mesh refinement converges, witnessed on $[0, 1)$. The scalar read off the stress-energy tensor is a marked modelling choice: the partition holds for any non-negative choice, and two instances differing only in it have different couplings (`Examples` §6)	Theorem + instance postulate
Coarse-graining of a symmetric coupling	`mesh_refinement_convergence` (`Phase2`) proves convergence of the scalar coupling energy along regular refining triangulations. On the formalized $[0, 1)$ grid the energy is the composite trapezoidal rule, so a C^2 integrand has error at most $\zeta/(12N^2)$. The result produces neither a kernel nor dynamics: `edge_region` supplies no product-structured data on $M \times M$. A kernel limit needs a different discrete architecture, which `KernelMesh` supplies through cell pairs rather than edges; the dynamical mean-field problem stays separate from both	Theorem (scalar energy only)
The field is the endogenous EM field	**Conditions formalized; the identification itself not derived.** A continuum coupling kernel is posited after the scalar convergence result; identifying that kernel with cortical ephaptic coupling is the framework's central empirical commitment (the physical-calibration subsection of the main article). `IsEMFieldCoupling` (`Phase9_EMIdentification.lean`) is the predicate `E56` carries: a jointly continuous kernel, a normalized substrate, a coupling constant K equal to that kernel's double average, no single site contributing more than half of it (`no_site_dominates`, which holds vacuously wherever the substrate has no atoms, the continuum case the identification is about), and a noise D that is the field's own. No object in this development denotes cortex, so whether cortex satisfies them is not a Lean question. Supported as *modulatory and integrative*, not as primary: measured endogenous fields of 1–5 mV/mm shift spike timing by 0.4–2 ms at the 0.4 ms per mV/mm of the Fermi calibration (Fröhlich & McCormick, 2010; Anastassiou & Koch, 2015). Calibration and observational tests are stated in the main article	Modelling assumption + empirical commitment
Frustration as memory capacity	**Outside the regime the chain covers.** `ThermodynamicCover.A_pos` excludes frustrated couplings at the gluing node, and circle-valued offsets from patch phases are coboundaries. Cortical phase singularities supply integer transitions between local real lifts instead: `loopWinding_eq_zero_of_hasGlobalLift` proves non-zero winding obstructs a global lift, with a winding-one witness. No theorem connects winding sectors to content	Theorem (topological obstruction) + empirical candidate

Conceptual Claim	Lean 4 Implementation	Status
Phase synchronization to Unity	Sheaf-theoretic gluing (`Phase5`) over a `ThermodynamicCover`, whose fields are all local: the global section is constructed by `probability_glue_unique`, and the statement that every local section is its restriction, `sync_to_section_eq`, is a theorem about it. With compatible finite-measure representatives on the finite cover, `SpatialMeasure.finite_glue_unique` constructs their unique finite-measure extension and `ThermodynamicCover.existsUnique_measure_representation` identifies its sheaf image with that section; arbitrary sections need not have finite total mass. Witnessed by two patches built from different mass profiles, each section constructed over its own patch, whose glued state is neither of them (`Examples` §14). The cover's equilibrium hypothesis is not an assumption on every instance: `ThermodynamicCover.ofConvergentTrajectory` (`Phase5_EquilibriumBridge`) discharges it from `kuramoto_tendsto_global_minimum`, and `Examples` §17.1 is a three-patch cover whose phase field is the limit of a trajectory started *unsynchronized*. The overlap-agreement hypothesis is untouched by any dynamics and is read pointwise: two patch sections have the same restriction to their overlap exactly when they assign the same mass to every shared site (`restrict_eq_iff_densityOn_eqOn`), and a two-patch cover at one phase with profiles differing at the shared site fails it while meeting every other requirement (`overlap_agreement_fails`, `Examples` §20). Under an ε-tolerance on overlaps the sheaf condition is replaced by a selection rule: normalized subordinate weights select a state within ε of every patch, two such selections differ by at most ε at any cover multiplicity (`select_dist_le`), zero tolerance recovers the unique gluing, and for every positive ε a three-patch family at that tolerance has no global section at all (`apx_no_global_section`, `Examples` §21). The cover is a choice, and the class constrains it in one way only, that synchronized patches agree where they overlap. The presheaf's value on the empty region is a singleton (`subsingleton_section_of_eq_bot`), so agreement across an empty overlap is free (`section_agrees_of_disjoint`), a cover by pairwise disjoint opens is an instance for *every* assignment of local data and phases (`LocalSectionSynchronization.ofDisjointCover`) and the unique global section is still delivered (`glue_unique_of_disjoint`); `Examples` §4.1 is the witness, one patch per site on the three-site substrate at the chain's coupling, legal at every tuple of local data (`shardCover_glued`). `HasNonemptyOverlaps` is the non-degeneracy this identifies, failed at every pair by that cover and met by the two-patch one (`cortexSync_hasNonemptyOverlaps`). It is a predicate and not a field, so no result above assumes it; that the patches' shared quantity is fixed independently of the agreement it establishes is not formalized at all	Theorem + instance postulate
Reflexive Topology (the Self)	Banach fixed point and residual bound $d(s, s_*) \leq \varepsilon/(1-\Lambda)$ for a self-prediction map (`Phase6`); metric, completeness and the Lipschitz law at $\Lambda = e^{-(K-K_c)\tau/2}$ are hypotheses, discharged on the three-site witness. Displacement errors add before division by $1-\Lambda$; the witness attains the bound, and nonexpansiveness alone permits distinct fixed points. Coherence yields $K > K_c$; it does not prove the Lipschitz law. `E89` supplies that law and identifies the glued state with the fixed point	Conditional theorem + modelling postulate (witnessed)

Conceptual Claim	Lean 4 Implementation	Status
When restriction resonance fails	`Resonates` states the predicate on bare maps, and a mismatch between the pairs the encoder confuses and the pairs the reference confuses refutes it (`not_resonates_of_confuses`, `not_resonates_of_splits`). With accuracy it becomes a condition on the region: a boundary reconstructing two relevant states more than 2ε apart that its avatar region reads identically is not resonant (`not_isRestrictionResonance_of_reconstructs`), and a resonant one restricts injectively on any such family (`restrictToAvatar_injOn_of_isRestrictionResonance`). `not_resonates_regionReading_of_outside_past` is the deadline form, from the causal past. The region is itself a choice, and these two obligations do not constrain it: at the whole substrate restriction is the identity (`restrictToAvatar_mk_top`), so `topAvatar` is resonant, reconstructs every declarable family at tolerance zero and restricts injectively on all of them, with nothing encoded in any proper region (`topAvatar_satisfies_obligations`). `IsProperAvatar` is the cheap form of the condition this identifies, a predicate and not a field, and `Examples/Reconstruction.lean` pairs the two regions on one substrate: the proper region provably fails to reconstruct the family it cannot resolve while the improper one reconstructs it exactly, both being resonant (`properness_is_what_costs`). The packing bound in the row above is the quantitative form, and it is empty at the whole substrate, where the encoder's range is everything (`topAvatar_encode_range`). The criteria refute resonance; they construct no resonant avatar, a region of sufficient resolution is not thereby shown to carry one, and properness is required by no result above	Theorem (negative)
Kuramoto Lyapunov descent	Descent chained to the minimum characterization by the rotating frame (`Phase4_RotatingFrame`); identical frequencies	Theorem
Kuramoto dynamics run	The initial value problem is well posed on all of $\mathbb{R}$: the field is globally Lipschitz and globally bounded (`kuramotoField_lipschitz`, `kuramotoField_norm_le`), and through every state at every instant passes exactly one trajectory, defined for all time (`kuramoto_cauchy_problem`). Along *every* trajectory of a zero-frequency system the dynamic potential converges (`dynamic_potential_tendsto`) and the total dissipation $\int_0^\infty \sum_i \dot\theta_i^2$ is finite, equal to the potential's total drop (`dissipation_integral_tendsto`). The motion stops: Barbalat's lemma is proved here, Mathlib lacking it, and applied, so $\sum_i \dot\theta_i^2 \to 0$ along every trajectory (`velocity_sq_tendsto_zero`). Where it stops is proved under an initial-data condition: if the phases start within a quarter turn and the potential excess starts below $a/2$, where $a > 0$ bounds the coupling below, the trajectory converges to a phase-locked configuration that is a *global* minimizer, with $r^2 = 1$ (`kuramoto_tendsto_global_minimum`), by a Polyak–Łojasiewicz estimate rather than a LaSalle principle. Unconditionally this is false — splay and twisted states are equilibria — so the condition cannot be removed. Witnessed on three sites, where no closed form exists (`Examples` §17)	Theorem (conditional)
Windings are stationary and the resultant does not see them	`Phase4_KuramotoDynamics` §6 works over `circulantSystem`: any finite abelian group of sites with coupling $A_{ij} = f(j - i)$ for an even f, of which the ring $\mathbb{Z}/n$ and the periodic sheet $\mathbb{Z}/n \times \mathbb{Z}/n$ are instances. `char_is_kuramoto_trajectory` makes a character state stationary there and `winding_is_kuramoto_trajectory` specializes it to $\theta_k = 2\pi qk/n$. `order_parameter_r_sq_char` puts the global resultant at exactly zero for any winding nontrivial on some separation, with `order_parameter_r_sq_winding` the ring condition $n \nmid q$ and `torus_coupling_does_not_determine_order` the pair on the sheet the sweep integrates; `coupling_does_not_determine_order` collects them into one system with two stationary states, at $r^2 = 1$ and $r^2 = 0$, so the coupling does not determine the observable. Stationarity is unconditional and needs no basin hypothesis. One winding number is sampled, and Section 12.6 develops the sweep	Theorem

Conceptual Claim	Lean 4 Implementation	Status
Which windings a coupling sustains	`hasDerivAt_charJacobian` linearizes the field at a winding into the circulant operator of the twisted kernel `twistedKernel`, $g_d = f_d \cos \Delta\theta_d$, even by `WindingData.cos_psi_neg` and diagonalized by `charJacobian_chi_re`. `charLambda_nonpos` with `charJacobian_stable_of_quarter_turn` excludes a positive rate for a nonnegative coupling whose winding advances by at most a quarter turn across every separation it reaches; `charJacobian_growing_mode` exhibits an unbounded solution of the linearized equation otherwise. The potential results refute rather than scope: `char_not_potential_min` denies a nontrivial winding the global minimum, and `charLineSecond_zero` with `charLinePotential_lt_of_charLambda_pos` and `char_not_local_min_of_charLambda_pos` deny it a local one wherever a mode grows. The criterion bounds rates at the linearization and yields no asymptotic stability, so the persisting states occupy a strictly narrower range and the decay-length boundary of Section 12.6 is measured rather than computed from it	Theorem
Patch order, and the content bound at a winding	`norm_order_parameter_le_mean_patch_order` keeps the patch estimator at or above the global one, for any cover whose patches share a size and whose sites share a multiplicity, and `cos_le_mean_patch_order_winding` puts nearest-neighbour patch order at or above $\cos(2\pi q/n)$. `order_parameter_complex_patch` identifies a patch's resultant with the order parameter of its subtype, `chord_le_of_patch_coherence` and `compatible_of_patch_coherence` re-read the coherence bound at that count, and `chord_le_of_char_patch` puts it at $2\lvert\sin \Delta\theta_g\rvert$ on a winding, where `chord_bound_char_global` against `chord_le_two` shows the global bound vacuous there. The bound attaches to a patch of a winding and not to its winding number: nothing here connects a winding sector to content	Theorem
A winding's degree, read from the state	`ringTransition`, `winding_succ` and `loopWinding_ringTransition` supply the electrode-loop transitions of `Phase5_PhaseLifts` from the state rather than from measured data, summing them around the loop to q. `winding_degree_obstructs` is the pairing — on one state at one coupling the resultant is exactly zero and the degree is q — which by `not_hasGlobalLift_of_loopWinding_ne_zero` obstructs every global real-valued phase lift. Scope is the ring's: `ZMod` n has no interior, so this is a loop's winding and not a spiral, the transitions are those of the affine lifts rather than of an unwrapped measurement, and the integer is identified with no cortical observable	Theorem
An amplitude layer and its defects	`Phase4_KuramotoDynamics` §7 adds the real Stuart–Landau field $\dot z_i = (\Gamma - \lvert z_i\rvert^2) z_i + \sum_j A_{ij}(z_j - z_i)$, the gradient flow of a potential against the real inner product on $\mathbb{C}^V$, which is what makes it an extension of this development rather than a replacement: the convergence results above rest on gradient structure and complex coefficients destroy it. `deriv_phase_of_amplitude_trajectory` gives the exact phase equation at a site of positive amplitude, with no slow-amplitude assumption, and `amplitude_phase_is_kuramoto` specializes it to equal amplitudes: the phase model is this model's phase equation with the amplitude ratios set to one. `charState_is_amplitude_trajectory_iff` makes $z_g = a\chi(g)$ stationary exactly when $a^2 = \Gamma + \lambda_\chi$, `exists_charState_iff` turns that into an existence band, and `phase_model_admits_winding_outside_band` states the disagreement with the phase model, which admits every winding at every coupling. `standing_wave_of_amplitude_band` reads the $q = n/2$ state in the lab frame as one spatial profile times one time factor. `plaquetteCirculation_eq_zero_of_winding` proves every such state defect-free on every plaquette, `norm_charState` supplying the equal modulus it turns on, and `sum_plaquetteCirculation_eq_zero` with `no_isolated_defect` proves a lone defect impossible on a torus for *any* phase field, exact or not. The band is one of existence: no stability result is proved and the persisting states occupy a narrower range. Γ is estimated nowhere, $\lvert z\rvert$ is a normal-form radius and not a measured envelope, the normal form assumes a supercritical Hopf bifurcation where much cortical wave phenomenology is excitable, and relaxational dynamics supplies no propagation speed	Theorem

Conceptual Claim	Lean 4 Implementation	Status
Wiring support and finite-region coupling	`rigid_gap` (`Phase7_Rigidity`) gives a strict resource-matched disadvantage when a better-correlated pair is unwired. `Phase7_FiniteRegion` preserves a matrix's total weight and phase correlation under a positive-mass cell embedding (`kernelIntegral_cellKernel`, `fieldCorrelation_cellKernel`). Point-supported vanishing under an atomless measure is therefore not a consequence of having finitely many components. `no_forced_gap_of_best_wired` gives an optimal wired allocation under its global pairwise maximum hypothesis, including diagonal pairs; the chosen distinct pair must then have correlation one. No energy pricing, dynamics or device exclusion follows	Theorem
Critical Coupling Threshold ($K_c = 2D$)	Self-consistency equation (`Phase8_SelfConsistency`): at or below K_c, threshold included, only $r = 0$ solves it; above K_c *exactly one* $r > 0$ does (`supercritical_fixed_point_existsUnique`), so the non-negative solutions are $\{0\}$ below and $\{0, r\}$ above; that r tends to 0 as $K \downarrow K_c$ and increases strictly with K (`coherent_branch_strictMono`); that it does so at the mean-field rate $r^2/(K - K_c) \to 1/D$, i.e. $\beta = 1/2$ (`Phase8_CriticalExponent`: `coherent_solution_critical_exponent`, witnessed by `coherentBranch_critical_exponent`); a solution is a density reproducing its own order parameter. The von Mises density enters this claim by classification rather than by hypothesis: `Phase8_FokkerPlanck` classifies the stationary densities of the equation and `supercritical_stationary_existsUnique` restates the bifurcation in those terms. Conditional still on no link to a trajectory and on no finite-N limit	Theorem (conditional)
Order and probability-current dissipation	`Phase8_FokkerPlanck`: `current_integral_sq_le` proves $(\int J)^2 \le D\sigma_J$ for continuous current and a continuous positive normalized density. `uniform_current_dissipation` attains equality under constant drive. `supercritical_zero_dissipation` and `no_positive_order_only_bound` give ordered self-consistent equilibria of zero cost and refute universal strictly positive order-only bounds. `cosine_rate_sq_le_dissipation` bounds the continuity operator's cosine moment. `Phase8_OrderSpeed`: `order_magnitude_speed_bound` and `density_order_endpoint_cost_bound` constrain a changing rotation-invariant order magnitude on a smooth path with $0 < r < 1$; no extension through zero order or thermal conversion is claimed. `arbitrary_coupling_speed_zero_cost` and `uniform_time_dependent_solution` refute an unrestricted coupling-speed bound on a genuine uniform solution. No cortical drive or ordered actuator cost is derived	Theorem
Information ceiling of a token interface	`mutualInfo_le_output_entropy` and `mutualInfo_bits_le_log_card` bound the existing KL information for arbitrary finite joint laws, including zero atoms. `token_information_bits_le` gives $k \log_2 \lvert V\rvert$ for the entire sequence, without independent emissions; `no_tokens_no_information` gives zero at $k = 0$. The input is finite and the output is fixed-length; cache state, timing and other channels are outside the bound unless included in that output	Theorem
Spectral error of a linear bottleneck	`Phase6_AttentionRank`: `spectral_tail_le_error` bounds arbitrary rank-constrained linear competitors by the omitted squared eigenvalues; `spectralTruncation_error` and `spectralTruncation_rank` prove attainability. `linear_bottleneck_error` derives the rank budget from a linear encoder/readout's intermediate dimension. `softmax_rank_counterexample` checks rank-one logits yielding rank-two attention. The finite eigenbasis is supplied; no Mercer expansion, neural spectrum, numerical rank tolerance or universal transformer limitation is derived	Theorem
Input-weighted reconstruction floor	`Phase6_InputWeightedRank`, `Phase6_SampledRank` and `Phase6_CovarianceRank` extend the linear floor to diagonal input weights, a sampled input law and a supplied covariance factor with singular bases. A shared synthetic task checks the train-law tail against saved rank-limited decoders, then tests an unseen scene family with causal attention and phase-code controls. Test-family transfer, nonlinear error floors, physical energy and experience are not consequences of these theorems	Theorem (conditional)

Conceptual Claim	Lean 4 Implementation	Status
Stationary densities and linear stability	`Phase8_FokkerPlanck`: a periodic gradient drift carries no stationary current (`stationary_current_zero`), so the positive normalized stationary densities for the cosine drift are exactly von Mises (`stationary_iff_vonMises`). `Phase8_Linearization` expands the current exactly and gives each Fourier mode of the uniform state its rate, $K/2 - D$ at the first harmonic and $-Dn^2$ above, so a growing mode exists exactly above $2D$ (`incoherent_instability_iff`). `Phase8_CoherentStability` proves the coherent branch's gap inequality transverse to its rotation orbit with a positive rate (`coherent_linear_stability`), the orbit direction having zero rate (`rotation_generator_zero`), on a strict moment inequality derived from a Riccati identity rather than assumed. At the coherent state current dissipation is zero while the phase-averaged squared drift is positive (`stationary_current_separation`). No evolution equation is solved, no semigroup or operator closure is built, no nonlinear stability or selection follows, and the stochastic differential equation is not connected to the Fokker–Planck equation	Theorem
Conditional composition	`Chain.lean`: `chain` and `chain_active` each take eight named hypotheses, a `ThermodynamicCover` and a `ReflexiveBoundary`, and conclude `UnifiedSelf` through the shared downstream proof. `E89` supplies the Lipschitz law and identifies the particular glued section with a fixed point; Banach supplies uniqueness. `E34` requires a passive predictive structure with the register's temperature, heat and nonzero waste. `E34Active` supplies physical premises and a heat allocation for the named finite feedback process; its entropy bound is derived, and `e34Active_of_ledger` derives the allocation itself inside one register's resource ledger. `E45` is discharged for named hardware holding a declared predictive system's own joint law. `E45Active` is discharged for a scalar actuation law with supplied gain, profile and regular meshes, whose limit is bounded using the named process's heat budget, and for a microscopic arrangement whose product-space kernel is built from occupancies that step changes. `E78` requires the supplied cover to be reached by relaxation at coupling at least K. `chain_hypotheses_jointly_satisfiable` and `chain_active_hypotheses_jointly_satisfiable` witness the two branches with distinct toy components. The active witnesses reject a zero heat budget, an incompatible energy limit, a compressive operation above the register's whole dissipation and a ledger whose other operation draws heat out of the reservoir. Joint satisfiability establishes no common microscopic mechanism; the actuated witness ties the scalar density to the process through a supplied constitutive law and the microscopic witness ties a kernel to it through declared hardware, while mesh regularity supplies convergence in both. `E12` is equivalent to its conclusion, so the first two links are independent roots. The supplied structures' physical fields and the representational assumptions remain additional inputs	Conditional theorem with eight named hypotheses
What a deadline reaches, and what it costs in wiring	`Phase6_Locality` carries a synchronous network as data — who hears from whom, what is sent, and a local update reading only its own state and the messages it received — so locality is structural and `run_eq_of_agree_on_ball` is derived rather than assumed. Three ways to bound a causal past. `ball_rank_le`: a rank no neighbourhood increases confines the past below it at every depth, and `Examples` §31 instantiates that on a decoder-only forward pass indexed by token position and layer, where `mask_ball_subset_le` keeps the past of position t inside the positions at most t at every number of rounds and `mask_no_guarantee` fires the obstruction at every deadline at once. `card_ball_le`: a graph of fan-in d reaches at most $\sum_{i \le T} d^{\,i}$ sites in T rounds, so `not_reconstructs_of_bounded_degree` produces the outside site by counting rather than by structure, and `card_ball_le_mul_pow` reads the deadline a guarantee needs as $\log_d$ of the site count; `card_le_geomSum_of_reaches` and `le_degree_of_reaches_one` run the count backwards, a one-round guarantee requiring a site that hears from all but one of the others. `not_outside_past_of_isFullSupport` is the other side: where every site hears from every site the obstruction's hypothesis is unsatisfiable at any deadline but zero, so the bound is silent there. The count is of hops and not of time, so none of this is a latency or a speed claim; the graph, the deadline, d and the site count are declared inputs, and whether an execution is a given graph is not a question this development answers; the wiring corollary bounds the communication graph a guarantee requires and says nothing about what hardware is buildable; and a full-support kernel models a physical field only while transit is short against the phase time scale, which is a calibration performed nowhere here	Theorem

Conceptual Claim	Lean 4 Implementation	Status
The channel a self-report is counted in	`card_le_card_codes` caps a family of relevant states pairwise more than 2ε apart and reconstructed within ε by the alphabet of codes, for any encoder and readout whatever. `tokenChannel` names that alphabet for a self-report: the code is the token sequence emitted over k steps, because that is what the next step reads, and `card_le_card_tokens` counts the family at the vocabulary size raised to k. `card_le_card_tokens_one` is where it bites, a single emission distinguishing at most one vocabulary's worth of states of the thing reporting, and `not_reconstructs_of_card_tokens_lt` is the usable contrapositive. What is counted is one declared channel's codes, and nothing follows about what a model computes; over k steps the alphabet is large, so the bound constrains a single-step claim and not an extended one; a cardinality is not a bit count and prices nothing; and a claim resting on activations rather than on text has declared a different encoding, to which the same bound applies with that encoding's alphabet, possibly the machine's whole state	Theorem
Capacity at a declared resolution, in both directions	`encard_le_packingNumber_range` replaces the alphabet of the counting bound by `Metric.packingNumber` of the codes an encoder writes: at readout gain L a family pairwise more than r apart and reconstructed within ε packs the code space at every δ with $\delta L \le r - 2\varepsilon$. `measure_le_mul_packingNumber` is the matching floor — a maximal δ-separated family is a δ-cover, so a region's measure is at most the largest δ-ball's times the packing number — and `measure_le_pow_mul_packingNumber` reads it on a finite-dimensional real space as volume over δ^d, with `measure_le_mul_packingNumber_range` carrying it to the encoder's codes. Ceiling and floor then hold of one quantity, which is what makes `two_pow_lt_packingNumber_of_lt_measure` a comparison: a continuum code space out-resolves a b-bit alphabet exactly when its volume exceeds 2^b resolution cells, and `encard_le_two_pow_of_packingNumber_le` runs it the other way. `Examples` §33 exhibits both signs on the unit interval. The floor is about the code space and not about any mechanism reaching it: nothing says an encoder writes those codes, that a readout separates them, or that the states they would encode exist. L is declared and uncalibrated, so a mechanism free to amplify without bound evacuates the ceiling and no floor on the code space repairs that; the two meet only where δ is a measured noise floor and L a measured gain, and neither is measured here	Theorem
Coarse-graining produces no kernel	**Negative result** (`Chain.lean` §9). `ContinuousNeuralField` carries no conditions, so a constructor from a mesh limit would be free of content (`continuousNeuralField_free`). `mesh_refinement_convergence` converges a *scalar*, and `TriangulatedManifold.edge_region` is a subset of M rather than of $M \times M$, so the discrete side has no product structure to pass to a limit. The vertex-sited candidate does carry the weights (`vertexKernel_apply_embedding`) and still contributes exactly zero continuum coupling energy on an atomless substrate (`vertexKernel_fieldCorrelation_eq_zero`) — the finite-support obstruction, cutting against the framework's own coarse-graining step. The obstruction is to this architecture and not to kernels: `KernelMesh` carries cell pairs of positive mass rather than vertices, and its limit is a kernel integral	Theorem (negative)

7 Phase dynamics and numerical controls

The controls examine finite trajectories, population dependence and signed coupling under declared numerical models. They do not identify a cortical interaction or discharge the formal mean-field limit.

7.1 Finite mean-field trajectory controls

Two finite-N controls probe the gap between the stationary-density results and finite trajectories without closing it. First, 500 seeded replicas of 1,000 identical-frequency oscillators were integrated at fixed coupling. Their steady order was 0.06266 below threshold, 0.16049 at threshold and 0.67735 above threshold. The last value lies -0.00535 from the stationary coherent value 0.68270. Refining the integration step by a factor of 4 at the same coupling, duration and tail sampling moves that steady order to 0.68168, so most of the gap is Euler–Maruyama bias. The last refinement still moves it by 0.00147, so the remaining -0.00102 bounds the finite-N part rather than measuring it.

Early-growth fits are regressed over a window whose bounds follow the trace: it opens at $2/\sqrt{N} = 0.06325$, clear of the finite-size floor, and closes at half the static branch order, which runs from 0.20669 to 0.37050 across the swept couplings. The fits give $\lambda = 0.41505K - 0.78619$, against the infinite-$N$ linearization $\lambda = 0.5K - 1$. That upper bound is set from the stationary branch rather than from the linear regime the fit assumes, so the finite window can itself contribute to the shortfall from 1/2: the step refinement above is a control on stationary order and apportions no bias in this growth estimator. The sweep also integrates at the step whose stationary bias that control measures, and carries finite N; this measurement separates none of the three. This exhibits finite sampled escape toward the coherent branch; it proves no stability or dynamical-selection theorem.

Second, 1,000 ensembles were swept from $N = 10$ to $N = 1000$ at $K = 1$ and $K = 3$ (Fig. S1). In the subcritical control, mean order fell from 0.3726 to 0.0402 and absolute circular pair correlation from 0.0458 to 0.0010. The matched-sample sliced joint–product distance changed only from 0.0468 to 0.0419, exposing its finite-sample floor. Above threshold, mean order stabilized at 0.7208, but the unconditional lab-frame pair correlation remained between 0.459 and 0.525 rather than following the $1/N$ reference. The unpinned collective orientation is a shared latent variable, so unconditional independence and independence conditional on that orientation are different claims. In the co-rotating frame, the $N = 1000$ density at measured concentration $a = 2.1624$ had L^1 error 0.1291 from the von Mises target and Bessel self-consistency residual 0.00191. These finite, single-seed diagnostics prove neither form of propagation of chaos and discharge no Lean obligation.

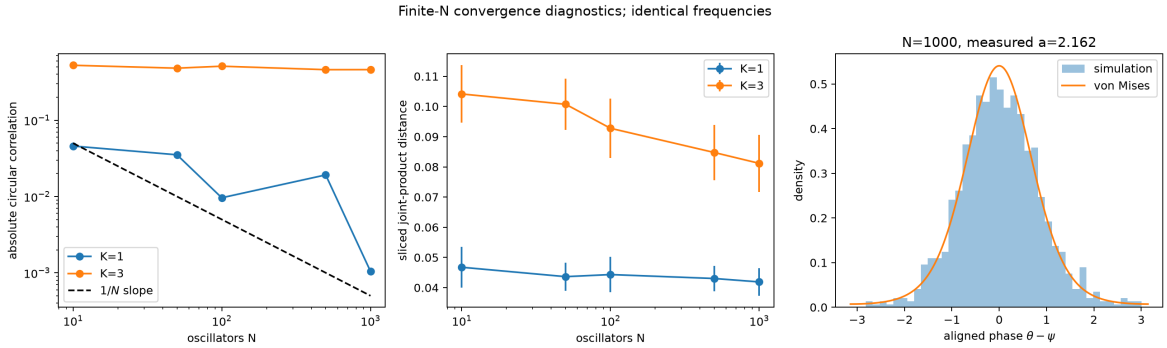


Figure S1: Finite-N convergence diagnostics for identical frequencies. The left panel compares branch-cut-invariant pair correlation with a $1/N$ reference; the center uses matched sliced-Wasserstein joint–product samples; the right compares the co-rotating $N = 1000$ histogram with the von Mises density at measured concentration. The persistent supercritical lab-frame dependence is a negative result for the unconditional protocol, not a theorem about the conditional mean-field limit.

Dale-balanced rescue and calibration limits

Finite simulations separate synaptic balance from the magnitude of fluctuations. The network has no self-edges, an Erdős–Rényi edge probability of 0.2, and excitatory/inhibitory presynaptic fractions of 80/20. Each row's positive and negative weights sum to g and $-g$, respectively. With identical frequencies, $D = 1$, $N = 500$, $g = 1$, step size 0.01 and duration 100, three seeded runs give zero-field mean order in $[0.0508, 0.0574]$, close to the $1/\sqrt{N}$ reference. Controls at $g \in \{0, 1, 2, 4\}$ and $N \in \{250, 500, 1000\}$ show why row balance alone is insufficient to prescribe that reference: signed coupling can amplify finite-size fluctuations. Order decreases with size in both the weak and strong controls while $N\langle r^2\rangle$ remains above the independent-phase value of unity; these finite sizes do not establish an asymptotic scaling law. The $g = 1$ regime is also below the positive-only mean-field threshold, so its low order does not establish a causal role for inhibition in preventing coherence.

A uniform per-pair term ϵ adds mean-field strength $K_{\text{eff}} = N\epsilon$. The operational crossing, defined by second-half mean $r = 0.2$ with all larger sampled couplings above it, occurs at $K_{\text{eff}}/D \in [1.96, 1.99]$ across the three seeds (Fig. S2). Each value interpolates across one step of the ϵ grid. That grid is geometric with ratio 1.136, and the one step of it the crossing falls in, $[1.857, 2.110]$, is subdivided linearly into intervals of width 0.025 and rerun on the same three seeds, which places the crossing in $[1.958, 2.009]$ — the union of the refined intervals the three seeds cross in. That bracket contains $K_c/D = 2$, so refining the grid tenfold locates the crossing without separating it from the mean-field threshold at $N = 500$.

This range describes finite-run variability, not a confidence interval or an infinite-size critical point. Equal-duration half-step controls probe numerical sensitivity near the transition; they do not establish convergence of the integrator.

The field conversion requires an explicit calibration rate γ:

$$E_{\mathrm{req}} = \frac{K_{\mathrm{eff}}}{\gamma N_{\mathrm{cortex}}\,\mathrm{shift}\,f}.$$

Since shift f is dimensionless, γ supplies the inverse-time units of the dynamical coupling. Under the numerical convention $\gamma = 1$ in the simulation's inverse-time unit, the shared Fermi arithmetic gives field ranges $[0.5856, 0.5930]$, $[0.0732, 0.0741]$ and $[0.0217, 0.0220]$ mV/mm at decay lengths 0.1, 0.2 and 0.3 mm. These conditional numbers fall below the 1–5 mV/mm comparison band and scale inversely with γ. The simulation does not measure that calibration, so it cannot independently validate physical-unit consistency or a cortical field threshold. Neither rescue nor these controls discharge a Lean or gluing obligation.

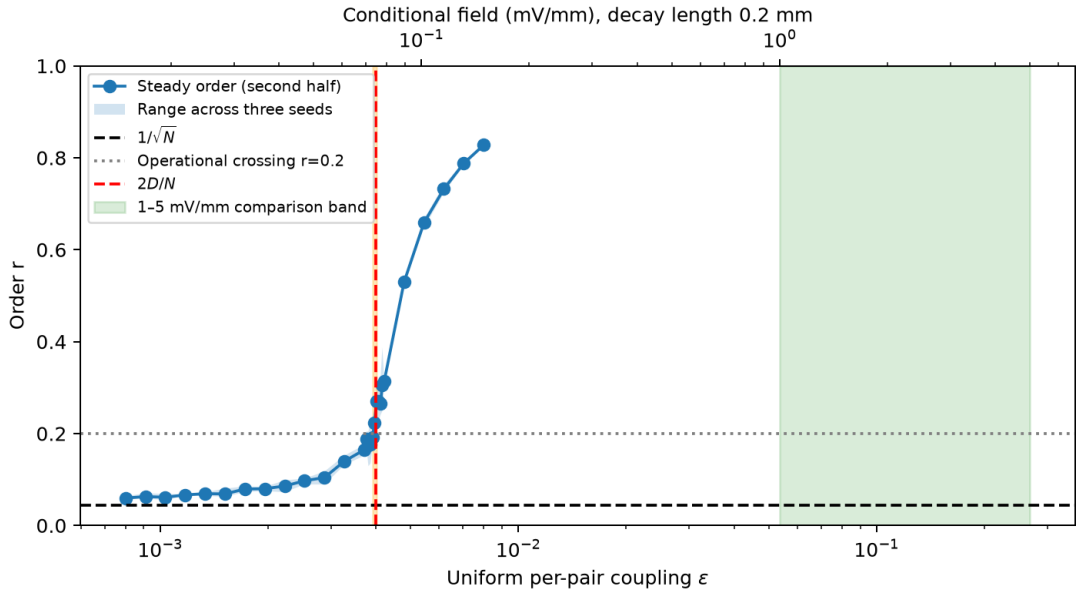


Figure S2: Finite-N rescue in the declared weak-synapse regime. Points are three-seed mean order; blue shading gives the seed range. Orange shading marks the union of sampled crossing brackets. The top axis uses the conditional $\gamma = 1$ field conversion at decay length 0.2 mm; green shading is the 1–5 mV/mm comparison band. The dashed horizontal line is $1/\sqrt{N}$. The balanced strong-synapse control remains a distinct regime and does not satisfy the floor-level baseline gate.

8 Recovery dynamics and mechanism identification

8.1 Assumptions

The stationary calculation assumes identical frequencies in a rotating frame and positive phase diffusion. Composing its square-root branch with a smooth crossing at nonzero speed gives the article's Eq. (11). Applying that deterministic tracking approximation to a trajectory additionally requires adiabatic branch tracking. A deterministic mean-field control starts from the initial stationary density displaced by a fixed amount in its order parameter and crosses on a leg symmetric about K_c. Its terminal quasi-static residual runs from 0.393 to 0.393 over scanned rates from 0.0001 to 1. None meets the declared tolerance under this initialization; slowing the crossing does not remove that residual. This finite scan supplies no general tracking guarantee at onset. The stationary theorem proves no dynamical selection, and finite-size fluctuations require an operational escape criterion above the fluctuation floor.

Away from the crossing, the comparison gives admissible rates as a function of distance from threshold. The declared geometry and phase time-scale ranges used to interpret the control give a ratio between

600 and 60×10^3. At the lower end the stationary reading meets the tracking criterion only down to $|K - K_c| = 0.8$, and at the upper end down to 0.05. These ranges are modelling inputs to the comparison, not a measured cortical coupling trajectory. The physical hypothesis additionally requires a calibrated geometry–coupling relation.

A coupling that fluctuates rather than ramps is read against a statistic of $K(t)$ that the fluctuation rate selects. At correlation time 0.03, substituting the mean coupling reproduces the measured order to 0.035 where substituting the quasi-static average $\mathbb{E}[r_{\mathrm{ss}}(K)]$ is wrong by 0.341; at correlation time 30 with a drive that stays on one side of threshold the two exchange roles, at 0.084 and 0.020 over 8 drives. A slow drive that crosses admits neither. Mean-coupling substitution is refused in 25 of the scanned drives, and in the strongest case — telegraph coupling of mean 2, amplitude 1 and correlation time 3 — the measured order is 0.0924 where r_{ss} of the mean coupling is exactly zero, above the seed 0.02 the run started from. Its spread across replicas is 0.1061, the size of the effect itself. These findings concern order over the simulated horizon and the variability of individual runs; they establish no shifted asymptotic threshold.

8.2 Prespecified spatial onset comparison

Equation (12) is a deliberately local extension of the first-harmonic growth rate $K/2 - D$. Its two candidate couplings share the same measured drug-exposure trajectory $g(t)$ and differ in a spatial covariate. The field score $h_i = N_i \mathcal{T}_i E_i f_i$ uses an anatomical population count, timing sensitivity, pre-onset extracellular amplitude and frequency estimated independently of the onset outcome. The connectivity score c_i is a prespecified weighted structural degree from a separate network measurement; it is a defined comparator, not a claim that every synaptic model predicts hub-first recovery. Normalizing each score across regions leaves one scale a_F or a_C to fit on training sessions. The same independently estimated D_i, observed baseline order $r_i(t_0)$, threshold relative to each region's baseline fluctuations, and recording observation model enter both candidates. A site crossing is predicted only while the first-harmonic small-order approximation applies.

Estimate the two covariate maps from anatomy and pre-onset recordings without using onset outcomes. Fix regions, the pre-onset window, the regional threshold, the recording window and the observation model before scoring. Split sessions within subject into training and held-out sets before fitting either scale; a claim of transfer to new subjects requires a separate subject-held-out evaluation. Both models use the same measured drug-exposure trajectory, baseline order, local diffusion, regional timing-error law and rules for missing channels. Fit the scales and timing-error parameters on training sessions only. The primary score is the sum of held-out regional log predictive scores: an observed crossing contributes the log predicted onset-time density, and a region with no crossing by the end of the window contributes the log predicted survival probability. The shared observation model supplies both probabilities and includes sampling resolution and channel noise. A missing channel has no observed onset and contributes no score term; report its frequency and spatial pattern. Site order, with ties at the recording resolution, is a secondary outcome. Apply the same exclusions and censoring rules to both models.

Predeclare a comparison of mean held-out log score across sessions, with uncertainty computed at the session level so that multiple regions in one session are not treated as independent replications. A region permutation of each covariate is a negative control. Repeat the score over a prespecified range of baseline-scaled thresholds and reference montages to reveal threshold and common-reference sensitivity. Record drug exposure, arousal, motion, reference choice, electrode coverage and task timing as nuisance inputs under the same observation model. Behavioural time-to-response to a standard stimulus is a separate endpoint: estimate its association with regional onset after those inputs are controlled, and report failure to observe a response as censoring at the behavioural recording limit. A missing association would count against the proposed link between the measured onset and recovery, even if one spatial map predicts onset well. Neither endpoint is a direct measurement of coupling or experience.

Before recruitment or acquisition, pilot data must estimate regional phase-noise and timing-error scales, the rate and duration of censored and missing observations, within-subject variability across taper rates, drug-exposure measurement error, and the correlation between field and connectivity scores. These determine a precision or power calculation for a prespecified held-out score difference; no sample size follows from the mean-field theorem. If the covariate maps are nearly collinear, or their predicted onset differences are smaller than temporal and spatial resolution, the design cannot discriminate them. If both models score poorly, neither specified local approximation is supported. A held-out field-model advantage would support the field covariate relative to this weighted-degree comparator, not every possible synaptic model. Interregional coupling, spatially varying phase spread, and the mapping from drug exposure to

effective coupling remain physical inputs. An independent perturbation of the proposed field contribution would be needed for causal identification. Content decoding and an experiential measure pose separate tests.

8.3 Finite-ramp numerical check

Figure S3 shows the ensemble-mean order and each replica's escape across the sweep.

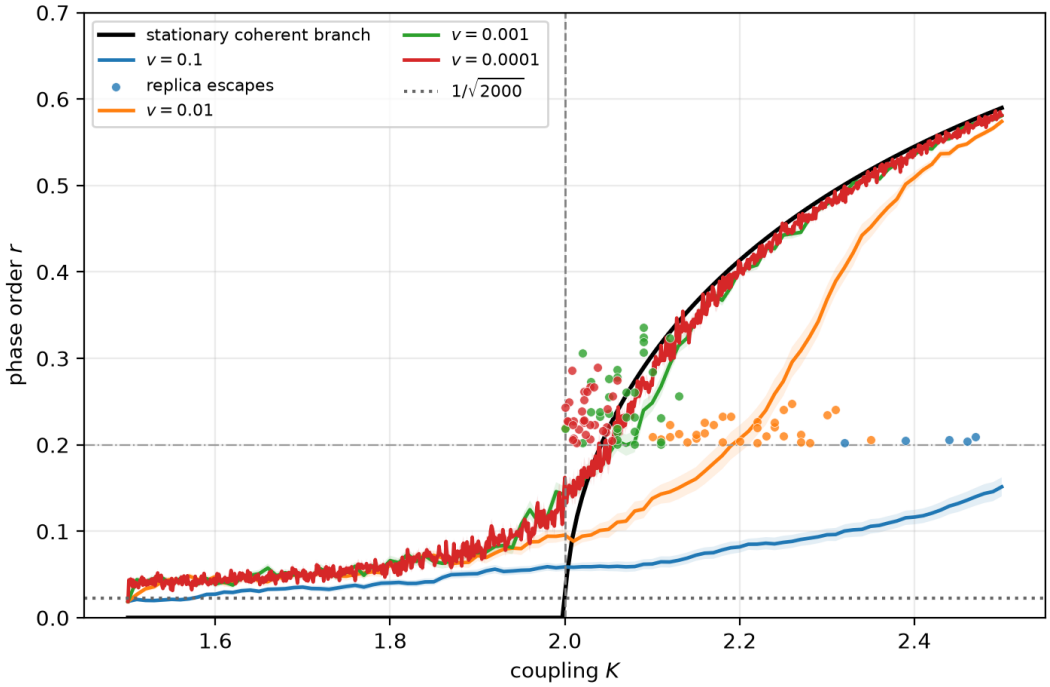


Figure S3: Finite-rate departures from the stationary branch. Curves show ensemble-mean order for four ramp speeds v, with one-standard-error bands. Coloured points mark each replica's first sustained escape, plotted at its recorded order; censored replicas have no point. The dashed vertical line is the static threshold, the dash-dotted horizontal line is the escape criterion, and the dotted horizontal line is the finite-size reference. The replica-bootstrap interval for the delay exponent is reported below.

A finite-N Euler–Maruyama ensemble quantifies the range over which the adiabatic reading is accurate. With identical frequencies, $D = 1$, $N = 2000$ and 32 seeded replicas, escape is defined as the first of three consecutive sampled points at $r \geq 0.2$ after K_c. The sweep runs 8 ramp speeds from $\dot{K} = v = 0.1$ down to $v = 0.0001$, placed so that the converted window at the end of this subsection contains three of them. The 2 fastest legs are right-censored — at the fastest, only 5/32 replicas escape before $K = 2.5$ — so the scaling is fitted on the 6 legs at which every replica escapes. The decade-spaced legs among them have mean delays $\Delta K = 0.1947$, 0.0684 and 0.0253.

The log–log fit over all 6 gives $\Delta K \propto v^{0.444}$, with a 95% slope interval of $[0.394, 0.494]$. Resampling replicas independently within each uncensored speed gives a 95% interval $[0.398, 0.496]$; 0.019 of 10000 bootstrap fits reach or exceed 1/2. The first interval measures fit uncertainty across mean points and the second seed variability conditional on these speeds and this escape criterion. Neither covers integrator error or estimator bias. The reference below applies the same estimator to a run carrying no ensemble. Split at the median speed, the faster half returns 0.492 $[-0.289, 1.273]$ and the slower half 0.447 $[-0.406, 1.299]$: the two point estimates agree, and each half is three legs across one decade, so its own interval is far wider than the shortfall being tested. The split therefore shows that the apparent exponent does not drift across the span, and neither half measures it.

The reference is the same delay measured in the deterministic mean-field limit, seeded at threshold at the finite-N fluctuation floor 0.0224 and stopped at the same criterion. Every one of its 8 speeds escapes, so the censoring above is the ensemble's, and the fit over the 6 legs the ensemble fits returns 0.425: the interval on the measurement contains its reference, and the shortfall from 1/2 survives the removal of finite-N sampling.

The reference shows that the escape criterion contributes to the shortfall. A run entering threshold at the floor and stopped at $r \geq 0.2$ crosses the saturation between those two levels, which the linear escape condition $\frac{1}{2}\int_0^T (K(t) - K_c)\, dt = \ln(r_{\rm esc}/r_{\rm seed})$ does not describe; that condition alone gives $\Delta K = 2\sqrt{v \ln(r_{\rm esc}/r_{\rm seed})}$, a square root in v at any fixed pair of levels. Tightening the criterion shortens the crossing: at $r \geq 0.1$ the same estimator over the same legs returns 0.482 and at $r \geq 0.05$ it returns 0.490, against a seed the criterion cannot go below and at which the measured delay would be zero. An onset delay read off an absolute threshold therefore carries an exponent that depends on how far above the fluctuation floor that threshold sits. This control does not establish that the criterion explains the full stochastic-ensemble shortfall.

Section 8.2 states the square-root hypothesis motivated by the ideal crossing. Testing it through an anaesthetic-rate manipulation requires an independently established relation between that manipulation and the coupling crossing rate.

At fixed $v = 10^{-2}$ the delay is monotonic in N, increasing from 0.0619 through 0.1947 to 0.2341 as N increases from 500 through 2000 to 8000. The corresponding ratios $\Delta K/\sqrt{\ln N}$ are 0.0248, 0.0706 and 0.0781, which do not collapse to a size-independent constant. At $N = 500$, individual replicas reach $r = 0.436$ before K_c, while at $N = 8000$ they reach only 0.106; the absolute criterion $r \geq 0.2$ therefore fires almost immediately at the smallest size. The control demonstrates the protocol hazard warned of in the main text's measurement protocol: a fixed absolute criterion meeting a size-dependent critical fluctuation floor.

The early-foot fit $r = A(K - K_0)^{\beta_{\rm eff}}$ is declared over the order window $0.1 \leq r \leq 0.4$ and is reported on the 4 decade-spaced legs, each of which carries a realized window a reader has to be told. Its exponents are read against the same estimator applied to the exact stationary branch on each leg's own coupling grid, which returns $\beta_{\rm eff} = 0.443, 0.443, 0.443, 0.448$ from fastest to slowest: the window reaches far enough past threshold that the estimator falls short of 1/2 even on a branch whose exponent is exactly 1/2, and that shortfall belongs to the window rather than to the drive. Against that reference the drive still dominates at $v = 10^{-2}$, where $\beta_{\rm eff} = 0.898$ stands well above the adiabatic value; the two slower legs return 0.471 and 0.417, which fall on either side of their references rather than approaching them from one direction.

The legs realize different windows and are comparable only so far as those agree: $K - K_c \in [0.04, 0.31]$, $[0.01, 0.19]$ and $[0.00, 0.20]$ over 28, 19 and 200 samples. At the fastest ramp $v = 10^{-1}$ the ensemble mean never reaches the declared window within the simulated interval: the fit truncates to $r \in [0.100, 0.151]$ at $K - K_c \in [0.32, 0.50]$ over 19 samples, so $\beta_{\rm eff} = 0.964$ is the local log-slope of a strongly delayed trajectory rather than an onset exponent. The shift K_0 is returned at its lower bound on 4 of the 4 legs, so pinning describes the estimator on this family of traces and does not single out the fastest fit. A mapping from these dimensionless rates to human recovery times requires an independently calibrated phase diffusion and coupling trajectory.

The independently estimated (a, r) traces remain close to $I_1(a)/I_0(a)$ even where the trajectory lags the stationary branch. Post-critical RMS deviations are $0.00245, 0.00599, 0.00689, 0.00705$; using the slowest run as the finite-N floor Dev_0, none reaches $2\mathrm{Dev}_0 = 0.01410$, and no leg of the denser delay sweep does either. The specified critical speed v^* is therefore not bracketed by this sweep.

The three delay exponents for the stochastic ensemble at $N = 2000, r \geq 0.20$, $N = 2000, r \geq 0.05$, and $N = 8000, r \geq 0.05$ evaluate to 0.444, 0.637, and 0.720 respectively. Evaluated on the matched six-leg subset $v = 10^{-2}$ to 10^{-4}, the exponents are 0.444, 0.634, and 0.710. On the matched subset, the sequence does not show convergence toward 0.5. Log-fit residuals are largest for the $N = 8000$ low-threshold control, so a common power law across settings is not supported by these fits. Several trajectories exceed the absolute escape level before the static threshold, and the postcritical crossing rule can yield a zero or grid-limited delay. The diagnostic compares log residuals and leave-one-speed-out slopes and checks precritical order, censoring and coupling-grid resolution from saved trajectories. These finite-sweep estimates establish no common asymptotic exponent.

At a Lorentzian half-width of $\Gamma = 1.5$, the completed sweep has too few uncensored speed legs for an exponent interval. The control therefore cannot decide whether heterogeneity changes only the prefactor or the exponent. A lower, predeclared escape criterion or wider coupling window is needed. This explicitly scopes any robustness statement to the mean-field idealisation actually tested. Furthermore, conduction delays and excitatory-inhibitory (E/I) balance remain open controls. The threshold behavior of the Kuramoto model with a Lorentzian frequency distribution is a known result in the literature (Acebrón et al., 2005), and frequency heterogeneity shifts this operational threshold.

This numerical check exhibits finite sampled dynamics only: it proves neither stability nor dynamical selection and does not discharge the adiabatic assumption.

8.4 Alternative models and failure conditions

Exponential relaxation is one comparator, not the entire class of electrical models. Slow latent variables, observation filtering and finite-size effects can also generate delayed or sigmoidal traces. Comparisons should use common preprocessing and observation models, with onset criteria specified before inspecting outcomes. A linear apparent onset counts against the adiabatic branch interpretation only in a regime where its assumptions hold.

The concentration–coherence relation tests the stationary-density model. Its failure over an informative concentration range constrains that model; its success does not identify the coupling medium. The ratio a/r is a consequence of the same density ansatz and is unreliable near the fluctuation floor. A coupling–geometry association requires an independently validated geometry proxy and an independently defined recovery endpoint. Stationary absence of a fold does not rule out finite-rate hysteresis.

8.5 What successive stationary phase distributions cannot separate

The physical controls of the protocol (step 4) need D separately from K because the stationary density does not carry them separately. In the identical-frequency rotating frame that density is von Mises with concentration $a = Kr/D$, so a window of equilibrated phases fixes the ratio K/D and nothing beyond it. Three recovery mechanisms therefore agree on every stationary observable, window by window. Write $s \in [0, 1]$ for a declared window coordinate, $q(s)$ for a schedule rising from 1.5 to an endpoint b, and $r(q)$ for the coherent branch of the stationary self-consistency relation, which is zero for $q \leq 2$. Increasing coupling is $K = q(s)$ at $D = 1$; decreasing diffusion is $K = b$ at $D = b/q(s)$; and a common drive is oscillators with no coupling to each other, each relaxing towards a shared phase ψ under a prescribed $-h(s)\sin(\theta - \psi)$ at $D = 1$ with $h(s) = q(s)\, r(q(s))$. The first two share $K/D = q(s)$; the third has stationary concentration h/D, which is that same number. Their stationary distributions coincide window for window, so no estimator reading only those distributions separates them. The third is one observationally matched drive and not a claim about every drive.

An ideal feasibility check confirms that the coincidence is exact rather than merely close. Nine equally spaced windows, exact densities on a periodic grid of 4096 points, endpoints $b \in \{2.6, 3.2, 4\}$ and orientations $\psi \in \{0, 0.7\}$ give 18 cases. Concentration is estimated by regression of log density on the cosine and sine of phase and coherence independently by circular integration; each family is then fitted for b over the same range against the same predicted concentration curve, with the generating family, K, D and true b withheld from the fit. The three families return the same b on every case, agreeing to within 8.9×10^{-16} at squared errors of at most 5.2×10^{-14}, and 0 of the 18 cases identify a unique generator. No family meets the declared gate of unique recovery on 80% of its cases, and the study stops there: observation noise, filtering and pooling applied to distributions that are already identical cannot separate them.

The observable the protocol is missing is the physical clock. Coupling and diffusion scaled together leave every stationary density untouched and change the rate at which phase moves, so resolved increments carry what equilibrated windows discard: along an ideal continuous path the quadratic variation accumulates at rate $2D$, and conditional increment means constrain the drift. Separating an endogenous coupling from a common drive is not settled by the clock either, since in the mean-field description the two impose the same drift on each oscillator; that separation needs a measured drive or a controlled perturbation. Neither measurement is validated by this pilot, and neither would identify the coupling as electromagnetic. The spatial-onset comparison of step 3 is accordingly a comparison of time courses under a declared observation model, and identifies the mechanism behind a recovery only to the extent that the observations reach past the stationary density.

9 Coupling plasticity controls

The following controls examine whether the specified squared-drift objective acquires its diagnostic environmental structure while preserving coherence. They report transient and complete-interval objective measurements separately. Kernel alignment is the measured quantity; overlap compatibility requires the distinct content specification in Section 10.

9.1 Joint phase and coupling plasticity control

This control tests alignment between a coupling kernel and a specified environmental template. It does not measure overlap compatibility. Its role in the article is to show why coherent adaptation under an objective does not itself supply the representational structure that gluing requires.

A finite driven control co-evolves phases and symmetric nonnegative couplings at fixed total coupling. It uses $N = 100$, $D = 0.1$, $dt = 0.01$, mean row coupling 4, and three nearly equal frequency clusters at 0.5, 1 and 1.5 per unit time. Couplings update every 50 fast steps with learning rate 0.2 for 400 time units, across three seeds. The symmetric-edge derivative of $\frac{1}{2}\sum_i v_i^2$ is $(v_i - v_j)\sin(\theta_j - \theta_i)$; each update clips negative weights, hollows the diagonal and restores the total resource. This constrained update does not guarantee monotonic descent during joint noisy dynamics. Each seed runs four arms on a single phase-noise stream: the gradient step, a symmetric hollow random step rescaled to the Frobenius norm of that arm's own gradient, that arm's own gradient step composed with a random node relabelling drawn afresh at every update, and a frozen kernel. Each direction is computed from the arm's own phases and couplings, so the controls follow the gradient arm's recipe and no control receives its contemporaneous step; the step sizes the recipe produces are comparable rather than identical, since the arms' kernels diverge. At zero learning rate all four arms reproduce one trajectory, which is the contract the shared stream is tested against.

The nonzero grand-mean frequency drives the collective coordinate, since symmetric interactions cancel in its mean drift. That cancellation fixes $\sum_i v_i$ across phase fields, so the objective cannot fall below $N\langle\omega\rangle^2/D = 990$ and the update can only equalize drift across nodes rather than remove it.

The objective is sampled at the update cadence, at the instant each update has just acted, so every objective mean in this subsection is a mean of that immediate post-update readout and not a time mean over the intervals between updates. On it the frozen arm sits at 1380.6063 to 1382.3448; the gradient arm's second-half mean is 1110.5659 to 1116.6277, removing a fraction 0.6773 to 0.6914 of the headroom between the frozen arm and that floor. The random arm removes -0.0173 to -0.0103 of it, so the reduction is specific to the gradient direction and not to displacing a kernel of that step size at fixed resource. What that readout measures is the objective the update leaves behind it at the moment it acts. Whether the objective stays below the frozen arm through the interval that follows is a separate measurement, reported in Section 9.2, and it is the one a claim of sustained minimization would need. The alignment statistics below are read from the kernel, which is constant between updates, so sampling them at the update cadence weights them exactly.

A step size set the same way is not a matched deformation. Fresh isotropic steps cancel across updates, so the random arm ends at a cumulative kernel norm growth of 1.2879 to 1.3056 against the gradient arm's 12.4094 to 14.6958, and hardly enters the channel by which clipping and rescaling concentrate a fixed resource onto fewer edges. The permuted arm enters that channel instead. A node relabelling is an isometry of the matrix and a bijection of the edges, so it preserves the Frobenius norm of the arm's own gradient and the whole multiset of its entries and changes only which edges receive them; the arm therefore takes a step of that gradient's size and of its shape, uncorrelated with the frequency clusters by construction. It reaches a norm growth of 11.4185 to 16.3296, a range bracketing the gradient arm's, and a second-half mean order of 0.9184 to 0.9189, which lies inside the gradient arm's range. Its descent fraction is -0.3319 to -0.1607: a step of a gradient's own magnitude delivered to the wrong edges raises the post-update objective. What the two controls jointly support is therefore a comparison between arms whose step sizes are set by the same rule and whose observed deformation and coherence ranges overlap, not between arms equated by construction.

The learning rate is selected by a sweep from 0.001 to 0.4 on one seed: the post-update objective's second-half mean falls from 1377.9075 at the smallest rate to a minimum at 0.2 and rises again to 1248.2272 at the largest. Rates past the minimum lose that reduction while still meeting the plateau criterion, so the criterion does not by itself select a rate and no convergence in the rate is claimed.

Second-half mean order ranges from 0.9167 to 0.9222, so coherence is maintained throughout. All seeds meet the declared plateau criterion: the last-quarter mean differs from the preceding quarter by at most 5%. The environmental template is the covariance of unit-variance latent signals shared within each frequency cluster, rescaled and hollowed to match the coupling resource; it is a diagnostic template, not an input to plasticity. Alignment is reported by two statistics unchanged under rescaling the kernel. The ratio of mean within-cluster to mean between-cluster coupling starts at 0.9824 to 0.9948 and ends at 0.3507 to 0.6135, against 0.9965 to 0.9997 for the random arm and 0.7290 to 1.2935 for the permuted one. The permuted range is wide and contains the no-preference line. Concentrating a fixed resource onto few edges makes a ratio of block means fluctuate, and with three clusters two thirds of the node pairs are between-cluster, so a dominant edge lands between more often than within and the typical value

sits a little under one. The three permuted values nonetheless lie above the gradient arm's three without overlap, at kernel deformations that overlap almost exactly. Against 2000 node relabellings drawn per run, the final kernel is further from the true template than from a fraction 0.9510 to 1.0000 of them. The frozen arm's initial kernel already sits at 0.6435 to 0.9110, so the gradient arm is read against the frozen arm and not against the null median; the permuted arm, at the gradient arm's deformation, reaches 0.2680 to 0.8410.

A second partition of the same nodes settles what the loss of within-cluster coupling is a loss against. The frequency partition is the structure the environment carries; an interleaved partition of the same three group sizes, built so that each group draws equally from all three clusters, carries no frequency information, and a kernel with only frequency structure is neutral under it up to a correction of order $1/N$. Neither partition enters the update. The gradient arm's interleaved ratio runs from 0.8071 to 3.0223 across the three seeds, on both sides of the no-preference line, where its frequency ratio is 0.3507 to 0.6135 on all three. The interleaved partition is one member of a family, so the calibrated form of the same reading is to score the frequency partition against 2000 relabellings of itself, each keeping the group sizes and losing the correspondence with frequency: at most a fraction 0.0495 of them give a lower ratio than the frequency partition does, on every seed, against 0.0755 to 0.3480 for the frozen arm's untouched kernel. Concentrating a fixed resource onto few edges makes any such ratio fluctuate widely, which is why the interleaved value alone is uninformative and the percentile against the family is not.

Two readouts are reported and not relied on. The initial-minus-final Frobenius distance to the template lies in $[-58.0962, -48.1760]$, but the kernel norm grows by a factor 12.4094 to 14.6958 over the same interval, and a fixed total resource constrains the sum of the entries rather than their norm; most of the distance growth is therefore concentration of mass onto fewer edges rather than loss of alignment, and the distance to the shuffled template grows with it. The off-diagonal correlation with the template ends at -0.0362 to -0.0208, close to its starting value: at the sparsity the update produces, edge-level variance dominates a block-mean contrast that the within-over-between ratio resolves.

Minimizing squared drift under a conserved mean drift equalizes drift across nodes, which rewards coupling between clusters of unlike frequency, whereas the template rewards coupling within clusters of like frequency. The disagreement is a property of this pairing of objective and template. The blind-partition reading is what separates that from the weaker statement it would otherwise be confused with: the update does not degrade whatever partition a reader scores it against, but locates the one partition the environment carries and moves against it. Which structure a different objective would move against, or whether any objective of this family aligns with an environment, is not settled by one pairing and is not claimed. The squared-drift objective includes the drive and fluctuating relative drift; it is not a thermodynamic entropy-production estimate, which would require probability-current information. The four panels of Fig. 4 in the main article are these trajectories: the dissipation function, smoothed over 21 saved samples, against the $N\langle\omega\rangle^2/D$ floor; the order parameter against the $1/\sqrt{N}$ reference; the within-over-between coupling ratio against the no-preference line at unity; and the distances to the cluster template and to its node-label shuffle, which grow together as the kernel norm grows. Only the gradient arm lowers the post-update objective, and only it loses cluster alignment, the permuted arm holding the same deformation without either; these diagnostics provide no evidence of structural resonance or a thermodynamic entropy-production rate.

These finite trajectories establish neither stationarity of the adaptive process nor noiseless phase locking, and do not discharge the E45 modelling assumption or extend the fixed-phase descent theorem. What they exhibit is a reduction at the sampled instant together with a loss of alignment, and not a system that holds the objective below its frozen counterpart while it runs.

9.2 Sustained reduction over complete update intervals

A reduction at the sampled instant and a reduction maintained between updates are different quantities, and the second is the one a claim that a coupling update minimizes the objective needs. A bounded study measures it using a design fixed before execution, with the grid, estimators, selection criteria and stop rule specified below.

The substrate is the one above: $N = 100$, $D = 0.1$, total coupling 400, physical duration 400, and the same initialization, frequencies, projection and objective. Only the learning rate and the physical interval between updates vary, over the product of rates $\{0.01, 0.05, 0.2\}$ and intervals $\{0.05, 0.5\}$ time units, each against a frozen arm on the same initial state and phase-noise stream. Objective and order are recorded at every integration step, on the state that drives that step, and averaged within complete update intervals; the terminal endpoint carries no duration and is excluded. The two assessment windows

are the third and fourth quarters of the run.

The criteria are declared in advance and conjunctive: each window's objective mean at least 5% below the paired frozen arm, each window's mean order at least 0.8, a second-half within-over-between coupling ratio at most 0.8 and at least 0.1 below the frozen arm's, and the final frequency partition at percentile at most 0.05 among 2000 size-preserving relabellings. Selection among the cases meeting the objective and coherence criteria is on the worst quarter's objective ratio alone and is never shown an alignment value; a selected case would then have to meet every criterion on three held-out seeds at two integration steps, at fixed physical duration and cadence.

Table S2: Every case of the declared grid, each against its paired frozen arm. The objective columns are complete-interval means over the third and fourth quarters divided by the frozen arm's; the order column is the smaller of the two windows' mean order; the last two are the second-half within-over-between coupling ratio and the percentile of the frequency partition among 2000 size-preserving relabellings of the final kernel. The criteria are both ratios at most 0.95, order at least 0.8, and the alignment conditions of the text, all together.

Rate	Interval	Q3 ratio	Q4 ratio	Order	Within/between	Percentile
0.01	0.05	0.9974	0.9973	0.955	0.745	0.0090
0.01	0.5	0.9972	0.9969	0.951	0.770	0.0145
0.05	0.05	0.9878	0.9864	0.850	0.462	0.0060
0.05	0.5	0.9917	0.9934	0.902	0.548	0.0765
0.2	0.05	0.9924	1.0135	0.836	0.223	0.0560
0.2	0.5	0.9980	1.0032	0.923	0.582	0.2690

Of the 6 cases of Table S2, 0 meet the objective and coherence criteria together, so the design selects no candidate and the held-out confirmation is not run. The smallest worst-quarter objective ratio falls at rate 0.05 and interval 0.05, whose two windows sit 1.22% and 1.36% below the frozen arm against the 5% the criterion asks for, at mean order 0.850 to 0.879, a within-over-between ratio of 0.462 and a blind-partition percentile of 0.006. Coherence and loss of alignment appear together across this grid; a sustained reduction of the declared size does not accompany them anywhere in it.

The grid contains the rate and cadence of the preceding subsection's run, and in that case the two readouts separate: the post-update second-half mean is 1118.95 where the complete-interval mean is 1383.87 against the frozen arm's 1383.06, a ratio of 0.998 in one window and 1.003 in the other. What the update takes off the objective at the instant it acts, the interval that follows returns.

The stop rule belongs to the design: the grid is not widened and no second candidate is tried once the outcomes are in. The study reports that this bounded family does not supply a sustained reduction, coherence and loss of alignment together. That is not a proof that no declared regime supplies them, and it is a statement about this squared-drift objective rather than about thermodynamic entropy production.

10 What a measurement of overlap compatibility would require

The main article's compatibility requirement needs a neural observable distinct from phase order. This section specifies that measurement: decoded local descriptions, restrictions to shared domains, uncertainty under decoding noise, and validation on examples with known compatibility. An implementation of the specification on constructed data is reported below; no recording is analysed.

10.1 Three components

A restriction map. Two regions with overlapping receptive territories may each support a decoder for content on the shared part. Compatibility is a relation between the resulting objects, so the decoder and its restriction map must be specified before any statistic is computed. The formalization uses finite spatial measures: on its discrete substrate two patch sections have the same restriction exactly when they assign the same mass to every shared site (`restrict_eq_iff_densityOn_eqOn`). A content decoder using probability laws needs additional structure: equality of overlap marginals is necessary, but global extension and uniqueness must be tested or assumed separately rather than imported from the spatial-measure theorem.

A statistic with a null. A distance between the two restricted distributions is not yet a measurement, because decoding error alone makes that distance positive. The statistic needs the distribution of the

distance under compatibility — what two decoders of the same underlying content produce given their own noise — so that incompatibility is a departure from a stated null and not a nonzero number. Overlapping territories supply dependent samples, and the uncertainty has to be computed accordingly.

Validation where the answer is built rather than inferred. The estimator must be exercised on data in which compatible and incompatible overlaps are constructed by design, so that its sensitivity is known before it is applied where the answer is not. The formalization supplies the configuration to target: local states at a single common phase whose densities disagree where the patches meet, which satisfies every other requirement of the gluing theorem and fails compatibility alone (`overlap_agreement_fails`). An estimator that cannot separate that configuration from its compatible counterpart at realistic decoding quality is not yet a measurement of the property the framework needs.

10.2 Three ways such a measurement fails silently

A decoder that shrinks both regional posteriors toward a shared prior reports compatibility whatever the regions encode, having replaced two local descriptions with one global one. A decoder carrying independent per-region bias reports incompatibility whatever the regions encode. These are opposite failures of the same component, neither is visible in the compatibility statistic itself, and both are why the constructed-answer validation above is a requirement rather than a refinement.

The third is circularity. A compatibility statistic computed from phases measures phase agreement, which the counterexample shows can coexist with incompatible local descriptions. E78 connects coherence to a supplied cover's relaxation; the cover's own compatibility field supplies overlap agreement. A statistic for that agreement must use the local descriptions and their restrictions.

10.3 What the implemented estimator separates

The three components above are instantiated on constructed data: per-site distributions restricted to the shared sites, equality of the restricted tuples as the compatibility relation, and a dependence-aware null built from two decoders of the same content. Over 600 trials on an overlap of 8 sites at decoding error 0.2, the statistic separates constructed incompatibility from constructed compatibility at $\mathrm{AUC} = 0.986$, on the 95% interval $[0.980, 0.992]$, against 0.476 on $[0.443, 0.508]$ under the null, a departure of 71 null standard deviations. Both intervals resample trials rather than sites, for the reason the statistic's own spread does: sites within a trial share a latent content field, so a site-level resample would report a width no run supports. The compatible construction's interval contains one half, which is what the constructed answer requires of it. Separation survives to decoding error 0.3 and to a shared territory of 2 sites.

The two silent failures are exhibited rather than assumed. A decoder shrinking both posteriors toward a shared prior drives the statistic to 0.0146 on incompatible content, and independent per-region bias drives it to 0.3791 on compatible content, against 0.2225 for honest decoders under the null: each manufactures the answer the other suppresses, and the statistic alone reports both. Both are rejected by the accompanying per-patch diagnostics, which is what makes those diagnostics a requirement rather than a refinement.

Circularity is separated by measurement rather than by argument. On content built from the phase, the phase carries a share 1.000 of the content entropy against 0.000 when it is not, so the share distinguishes the two models; the phase-derived bound does not, reading 1.852 and 1.804 on the two. On the constructed counterexample — phase-locked patches whose local descriptions disagree — a phase-derived statistic reports compatibility where the content statistic reports incompatibility. No phase-derived quantity substitutes for the restricted local descriptions.

10.4 Scope

A validated observable would make agreement between the chosen local descriptions measurable. Establishing that coherence maintains it would require a dynamical test in addition, and selecting the content model and cover would remain an independent task. A full test of the consciousness identification would also measure the proposed self-encoding and an experience endpoint: compatibility and supercritical coupling alone leave the reflexive clause unevaluated. Absence of report must be distinguished from evidence against an episode. If compatibility cannot be distinguished from incompatibility at achievable decoding quality, the identification is outside the reach of that measurement.

11 Exploratory EEG exercise of the concentration–coherence estimators

This section reports an exploratory check of the parameter-free relation $r = I_1(a)/I_0(a)$ on public propofol-sedation EEG (OpenNeuro ds005620 v1.0.0 (Bajwa et al., 2024, 2025)). The saved summary and the generated macros record every numerical result below. This is an estimator exercise, not a test of awakening or of the electromagnetic-field commitment.

11.1 Data

The recordings are OpenNeuro dataset ds005620, version 1.0.0 (Bajwa et al., 2024), collected for a repeated-awakening study of propofol sedation (Bajwa et al., 2025) under approval of the Regional Committees for Medical Research Ethics South East Norway (REK 2015/1520) and released under a CC0 licence. The dataset contains 21 participants, one of whom its authors mark as excluded. Each recording was made on Brain Products hardware at 5 kHz; in every analysed file the first 62 channels are the same scalp EEG electrodes, followed by two EOG channels and, in six of the eight files, one EMG channel. This exercise uses 8 participants (`1010`, `1016`, `1022`, `1033`, `1045`, `1054`, `1060`, `1067`), none of them excluded. They form an exploratory convenience subset: no selection rule is claimed for them, and the cross-subject summary below is not an estimate for the dataset as a whole. For each participant the analysed block is the first resting recording under sedation (`task-sed_acq-rest_run-1`); the narrowband analysis uses participant `1016`'s block of the same kind. No awake recording enters a reported result.

11.2 Preprocessing

For each recording block, the first 26 seconds were downloaded from the public S3 bucket: 20 seconds of analysed data with 3 seconds of padding on each side. Each channel was detrended, bandpass filtered with a fourth-order Butterworth filter using second-order sections (`sosfiltfilt`), and converted to instantaneous phase with the Hilbert transform. The padding is discarded after filtering in every montage. The primary analysis uses 61 adjacent bipolar pairs formed from the 62 scalp EEG channels; EOG and EMG channels are excluded.

11.3 (a, r) estimation

Each recording was divided into non-overlapping 100 ms bins (500 samples at 5 kHz). The estimand is the pooled spatiotemporal phase distribution over 61 bipolar pairs and 100 ms, not an instantaneous spatial distribution. Pooling increases the sample count for log-density estimation; statistical adequacy still requires the bias and dependence calibration below. For each bin t:

- **Order parameter r:** $r_b = \left| \frac{1}{NM} \sum_{j=1}^{N} \sum_{m=1}^{M} e^{i\theta_j(t_{bm})} \right|$, where $N = 61$ bipolar pairs and $M = 500$ samples per bin. This is the resultant length of the pooled sample, not an average of instantaneous spatial resultant lengths.
- **Concentration a:** The slope of the log-density of the same phase distribution regressed on $\cos\theta$, as specified in the main text. Phases are centred on their circular mean, binned into 40 bins on $[-\pi, \pi]$, and log of the binned density (with a 0.5 pseudocount) is regressed on $\cos\theta$ by *unweighted* ordinary least squares. Under the von Mises model $\log \rho(\theta) = a\cos\theta - \log(2\pi I_0(a))$, so the slope is a.

The estimator must not be a maximum-likelihood one. The von Mises family is a one-parameter exponential family whose sufficient statistic is $\sum_j \cos\theta_j$ — that is, r itself — so any ML estimate of the concentration solves its score equation $I_1(\hat{a})/I_0(\hat{a}) = r$ and returns $\hat{a} = A^{-1}(r)$ identically. The pair $(\hat{a}, r)$ would then lie on $r = I_1(a)/I_0(a)$ for any input whatsoever, including phases drawn from a distribution that is not von Mises, and the collapse would be an algebraic identity rather than a test. The implemented regression is unweighted, so it uses the log-density shape rather than defining concentration by inversion of the first trigonometric moment. Other weighting rules are different estimators and are not assumed equivalent to maximum likelihood.

The pseudocount biases the estimator, so its null residual must be calibrated at the sample size and dependence structure of the intended measurement. The exploratory bins contain 30500 pooled phase values. At 100 independent samples, 500 von Mises draws with true $a = 2.0$ return mean $\hat{a} = 1.346$

and mean residual +0.146 with standard deviation 0.028. A large residual at that count can therefore be estimator bias rather than evidence against the density relation. Independence is itself optimistic for EEG. Section 11.5 evaluates the site count and dependence needed for discrimination; the pooled exploratory count does not by itself validate a spatial estimator.

The estimator also separates synthetic non-von-Mises alternatives from correctly sampled high-count von Mises data. This confirms that the choice of unweighted log-density regression avoids the maximum-likelihood tautology; it does not calibrate a low-count or dependent neural measurement.

11.4 Constructed signals through the same pipeline

Signals of known coherence were passed through the same bipolar montage, filter, Hilbert transform, trimming and 100-ms binning (`eeg_pipeline_check.py`). A perfectly phase-locked 10-Hz rotation, with amplitudes falling along the montage so that every bipolar pair is locked as well, has instantaneous spatial resultant one on both raw channels and bipolar pairs. The pooled estimand returns $r = 0.0001$ and $a = 0.0004$, with a residual from the Bessel relation indistinguishable from zero: within 100 ms the common phase sweeps a full cycle, and pooling reads that rotation as dispersion. Centring each sample on its instantaneous circular mean before pooling recovers $r = 1$.

A rotating field with fixed von Mises spatial offsets ($\kappa = 4$) and additive white noise has raw-channel instantaneous resultant 0.89. Its bipolar pairs have instantaneous resultant 0.06, because the phase of a difference of two offset sinusoids depends on their amplitudes and offsets as well as on the common phase, and the pooled estimand reads 0.007. The montage and the pooling therefore each remove spatial coherence, and a pooled residual near zero is what any nearly uniform distribution produces: at small a both r and $I_1(a)/I_0(a)$ are close to $a/2$.

The same constructed field tests the quadratic-variation reading. For the latent phase of Eq. (1), squared increments accumulate at rate $2D$ at every step. The detrended squared increments of the band-limited analytic phase accumulate 37 times faster per unit time at 20-ms steps than at 0.2-ms steps, because the filtered phase is smooth below its correlation time. The rate $2D$ is recovered from a resolved trace only under an observation model for filtering and sampling.

11.5 Dependence-aware spatial design

The spatial design uses resultant length for r and unweighted regression of log density on centered cosine for a. At 100 sites, 40 bins clear the declared bias floor at no scanned concentration or dependence. Reducing to 24 bins or fewer permits discrimination only from $a = 1.5$ upward under independence. At the declared design concentration $a = 0.542$, near the low end of the scanned grid, the scanned designs instead require 3000 independent sites, rising to 10000 at half clustering and 31000 when sites share a mean in groups of 10. These are results for the specified estimator, concentration grid, dependence model and bias-floor criterion. A recording design must calibrate those choices in its own observation regime; its minimum site count cannot be inferred from the number of fitted parameters.

11.6 Bootstrap confidence intervals

95% confidence intervals on the mean (a, r) per block were obtained by bootstrap resampling of the time bins (2000 resamples, percentile method). This procedure does not preserve dependence between neighbouring bins; its intervals are exploratory rather than a calibrated time-series uncertainty estimate.

11.7 Narrowband analysis

The pipeline was run separately on theta (4–8 Hz), alpha (8–12 Hz), beta (13–30 Hz), gamma (30–40 Hz), and broad (4–40 Hz) bands.

11.8 Results

Residuals are evaluated per 100-ms bin and then averaged, rather than computed from block-mean a and block-mean r; because I_1/I_0 is nonlinear, the latter is not the mean residual. Figure S4 shows that this particular pooled estimator is sensitive to very short windows. It does not validate an instantaneous spatial estimator.

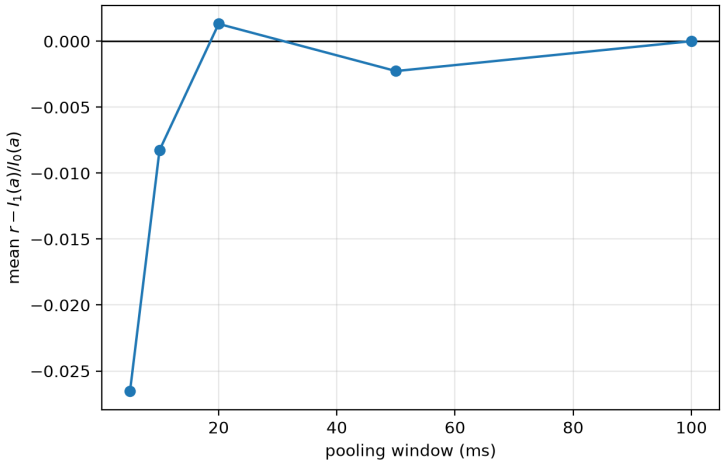


Figure S4: Sensitivity of the mean collapse residual to the spatiotemporal pooling window. This assesses the pooled estimator and does not turn 62-channel EEG into an instantaneous spatial sample.

- **Cross-subject bipolar summary.** Across 8 sedated subjects, the mean concentration is 0.104 (range 0.087–0.142) and mean order is 0.052. The mean residual is -0.0002, with RMSE 0.0027 and maximum absolute residual 0.0364.
- **Narrowband.** The theta, alpha, beta, gamma and broad-band mean residuals are -0.0018, +0.0001, +0.0001, +0.0000 and +0.0000, respectively, for sub-1016, sed run-1, bipolar.

Reference choice remains a material analysis decision. The summary fixes the primary montage and padding rule, but does not establish a mechanism for differences between montages.

11.9 Interpretation

The exploratory bipolar data lie near the relation at their observed range and pooling choice, and Section 11.4 shows why that is uninformative: the same pipeline puts a perfectly locked signal on the relation at $r \approx 0$. The pooled estimand measures neither the spatial concentration of the model nor its coherence. The exercise establishes estimator behaviour and a reference choice that matters; it is not evidence for the stationary density, for an electromagnetic coupling medium or for the awakening prediction.

The observed pooled range is also too narrow to discriminate the Bessel ratio from its linear approximation at the trace's scatter. The curve departs from that approximation by at most 0.0002 over the observed concentration range; a discriminating awakening experiment must span concentrations around 1.5, where the separation is 0.154, and must estimate concentration from phases centred on their instantaneous mean, with its estimator calibrated at the actual spatial sample count.

An informative dataset therefore needs an estimand that removes the common rotation, a montage whose effect on spatial phase is modelled, sufficient concentration range or precision, calibrated estimators and an observation model for dependence and source mixing. Intracranial recordings (ECoG, sEEG or LFP) or narrowband burst analysis may reach a more informative concentration range, but that range must be measured. The design counts of Section 11.5 are calculations for the specified estimator, not sample requirements transferable to another recording. Public mouse calcium imaging provides thousands of spatially located neurons across waking, sleep and anesthesia (Oomoto et al., 2026), but its 7.65 Hz frame rate is below the phase-estimation band used here. Site count, state coverage and bandwidth are separate requirements.

The result does not test the empirical commitment that the endogenous electromagnetic field is the coupling medium, nor the stipulation that the global section is experience. The coupling-gated awakening prediction remains untested here; these recordings concern propofol sedation rather than recovery from sleep.

12 The cortical field identification

The framework identifies its posited continuum coupling kernel with the endogenous electromagnetic field of mammalian cortex. This is a claim about biology, not a consequence of any theorem above it.

12.1 Supporting evidence

Neurons do not only signal through synapses: transmembrane currents generate extracellular fields that shift membrane potentials of neighbouring cells, a mechanism known as ephaptic coupling. Endogenous fields of 1–5 mV/mm entrain spike timing (Fröhlich & McCormick, 2010; Anastassiou & Koch, 2015), and stable electric fields carry information that survives representational drift in the underlying spiking population (Pinotsis & Miller, 2022, 2023a,b). The field is spatially continuous and superposes contributions from large populations — a medium in which spatially separated activity is combined without further mechanism. This is the integrative role the CEMI field theory assigns it (McFadden, 2002, 2020; Pockett, 2012), physically supported by astrocytic syncytia joined by gap junctions (Verkhratsky & Nedergaard, 2018).

These observations support a modulatory interaction, but do not determine its integrated strength relative to phase diffusion. Continuity does not itself establish compatible local states or conscious unity. The model requires independent calibration of the coupling kernel, its spatial aggregation and its rate scale.

12.2 The field-strength objection

The longest-standing objection is that the fields are too weak to matter, pressed from within the field-theoretic camp by Pockett herself (Pockett, 2012, 2002). Vöröslakos and colleagues found that applied current is shunted through scalp, delivering fields at or below 1 V/m in cortex (Vöröslakos et al., 2018). That is a bound on *applied* rather than endogenous fields, but it constrains how much causal work a field of that magnitude can do, and is the reason we do not claim the field drives spiking.

12.3 Why the rate calibration is not derivable

The conversion $K_{\text{eff}} = \gamma N \mathcal{T} E f$ carries its whole physical content in γ, and no rearrangement of the remaining factors supplies it. The frustration control above records the dimensional half of this: s is a time shift per unit field, so sE is a time, and multiplying by an oscillation frequency leaves a pure number — the fractional phase advance a cell receives per cycle from the population field. Multiplying by a population size leaves it a pure number. A coupling in the sense of the main text's stochastic model is an inverse time, and γ is the entire gap between the two.

The substantive half is that γ is a hypothesis about persistence, not a unit conversion to be fixed by convention. It states how a per-cycle phase advance accumulates into a drift rate, which depends on how long a field-driven timing shift survives. Three candidates are available and they are not interchangeable: $\gamma = f$ takes each cycle's advance to be retained in full and independently; $\gamma = 1/\tau$ for a phase relaxation time takes it to decay as the medium relaxes; $\gamma = D$ takes it to accumulate at the rate phase itself diffuses. These need not agree, in cortex or anywhere else, and the field measurement chooses between them only if the persistence of induced timing shifts is measured alongside it.

The third candidate is a physical hypothesis, not an algebraic error. Under $\gamma = D$ the threshold test $K_{\text{eff}}/D > 2$ reduces to $NsEf > 2$; that cancellation is legitimate if an independent experiment establishes that persistence rate. It is unsafe to choose $\gamma = D$ merely because it cancels the unmeasured diffusion, or to recover any γ by requiring the model to cross its own threshold or fit the Bessel relation. Identification requires coupling, diffusion and the persistence rule to be estimated in one cortical regime by routes independent of the field arithmetic, which is item 1 of the main text's protocol.

12.4 Calibrating the scalar spatial reduction

A spatially decaying kernel enters the scalar model through its row sum at zero diagonal, and a finite sheet measures what that reduction costs: over decay lengths from 0.1 to 0.3 mm, spanning 87 to 566 effective neighbours, the measured threshold lies at 0.750–0.753 times $2D$, within 0.003 of the 0.750 a uniform control returns on the same sheet and fit window. Both sit below $2D$ there — the fit window alone brackets them down to 0.658 — so the agreement is between the reduction and its control rather than a confirmation of the threshold's value. Within this control, frequency heterogeneity shifts the

operational threshold: a quenched spread of 1.5 rad s^{-1} moves it to 1.53 times $2D$, while the same sheet reproduces the stationary branch to 0.0048 over 19 supercritical couplings.

12.5 Finite-sheet spatial-decay control

A finite stochastic control asks whether fixing total coupling while shortening its spatial reach fragments the phase field. On a periodic 128×128 sheet spanning 2.0 mm, an exponential kernel was row-normalized to total coupling $K_0 = 8$ at every decay length; the dynamics used thermodynamic diffusion 0.5 rad^2/s, quenched Gaussian frequency spread 1.5 rad/s, and 20,000 Euler–Maruyama steps per length. With steady $r > 0.2$ as the declared operational coherence criterion, the crossing falls between the sampled lengths 0.0140 and 0.0147 mm, and interpolating within that pair puts the smallest length above which every larger sampled length remained coherent at 0.0140 mm. On this grid that is 0.90 lattice spacings, and the shortest sampled length from which coherence is sustained is 0.94 spacings, so the sheet places the boundary below its own resolution.

Rerunning the transition band and the lengths above it on a 256×256 sheet of the same extent — half the spacing, so a given physical length spans twice as many cells — leaves the boundary at 0.01413 mm, 1.01 times its coarse value and 1.81 spacings on the finer grid. Both crossings fall between the same pair of sampled lengths, so that ratio measures the interpolation inside one grid cell rather than a displacement the sweep resolves; what the grid supports is that the boundary lies in that cell on both sheets. Two threshold interpolants agreeing is not convergence of the dynamics they interpolate between: at 0.01467 mm, the shortest sampled length from which both sheets stay coherent, steady order is 0.6176 on the coarse sheet against 0.2687 on the finer one, and the refined crossing is still only 1.81 cells wide. The transition-local trajectories are metastable and non-monotone on both sheets, so that number is an operational boundary rather than a critical-point estimate with statistical precision.

At the empirical controls 0.1, 0.2 and 0.3 mm (Fig. S5), steady r was respectively 0.9424, 0.9427 and 0.9427, and winding-defect density was at most 2.4×10^{-6}. From 0.1 mm out to the sweep's longest range 8.0 mm, which is 4.0 times the sheet's extent, steady order spans only 0.00038, so the empirical band sits inside a saturated plateau rather than near a boundary. Thus this finite, single-seed parameter sweep does not show fragmentation in the empirical decay-length band. It neither measures cortex nor proves a transition, and it leaves seed dependence and the near-boundary metastability unresolved.

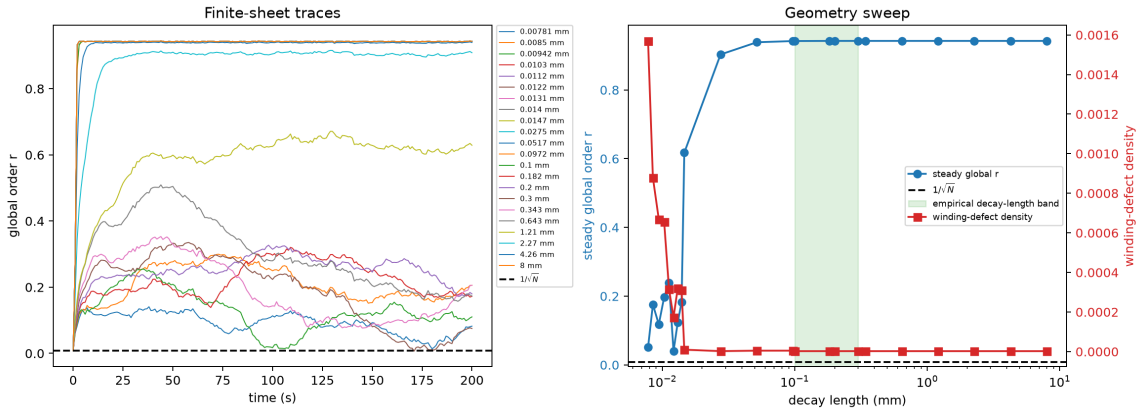


Figure S5: Finite-sheet heuristic evidence under a row-normalized spatial kernel. The left panel shows global order traces and the $1/\sqrt{N}$ finite-size floor; the right compares steady order and wrapped plaquette winding-defect density across physical decay lengths. The shaded band is 0.1–0.3 mm. This sampled simulation establishes no general transition or biological identification.

12.6 What the order parameter reports for a winding state

The coherence observable is the global resultant $r = |N^{-1}\sum_i e^{i\theta_i}|$, invariant under a global phase shift and otherwise indifferent to where a phase sits on the sheet. A field winding once across a periodic sheet has $r = 0$ while every patch of it stays locked, so the quantity in which the threshold $K_c = 2D$ is stated does not by itself separate a winding field from incoherence. What it misses is the winding and not any

motion: at identical frequencies with a symmetric kernel such a field has zero drift and stands still, and a travelling wave is the same winding carried by a drift this case does not supply.

Such a field is not a perturbation of this model. On a torus the twisted state $\theta(x) = 2\pi qx/L$ is stationary for any isotropic symmetric kernel at identical frequencies: $K(x-y)\sin(\theta_y - \theta_x)$ is odd in the separation against a kernel even in it, so the drift cancels pair by pair and vanishes exactly. It satisfies both positivity and the identical-frequency hypothesis, and `velocity_tendsto_zero`, which assumes only the latter, holds of it trivially — its velocity is zero throughout. What a twist violates is the quarter-turn initial spread $|\theta_i - \theta_j| \le \pi/2$ that `kuramoto_tendsto_global_minimum` additionally requires before placing the limit at the potential's global minimum, where $r = 1$. A twist spans the whole circle and lies outside that basin hypothesis. The relation is therefore the one positivity already illustrates: *that* the motion stops is settled without the hypothesis, and *where* it stops is settled only inside a declared basin.

All four statements are theorems, and of a site set general enough to be the sheet rather than only the ring: any finite abelian group under `circulantSystem`, the coupling $A_{ij} = f(j-i)$ for any f with $f(-d) = f(d)$. `char_is_kuramoto_trajectory` proves a winding stationary there, the drift cancelling by pairing each separation with its negative, and `winding_is_kuramoto_trajectory` specializes it to $\theta_k = 2\pi qk/n$ on $\mathbb{Z}/n$. `order_parameter_r_sq_char` puts the global resultant at exactly zero for any winding that is nontrivial on some separation, which on the ring is the condition $n \nmid q$ that `order_parameter_r_sq_winding` states. `torus_coupling_does_not_determine_order` assembles the pair on $\mathbb{Z}/n \times \mathbb{Z}/n$, which is the geometry integrated below. `norm_order_parameter_le_mean_patch_order` proves the patch estimator used below never reports less than the global one, for any cover whose patches share a size and whose sites share a multiplicity — the sliding box is the case where the two counts agree — and `cos_le_mean_patch_order_winding` puts nearest-neighbour patch order at or above $\cos(2\pi q/n)$, which approaches one as the same winding is spread over more sites. `coupling_does_not_determine_order` collects them into one system with two stationary states, at $r^2 = 1$ and $r^2 = 0$.

Which windings a coupling sustains is a separate question from which ones it admits, and it has a criterion of the same shape. Write $\theta_i = \bar{\theta}_i + su_i$ and differentiate the field at $s = 0$: the constant term is the drift that cancels above, and what is left is circulant again with the *twisted* kernel $g_d = f_d \cos \Delta\theta_d$, each separation reweighted by the cosine of the phase the winding advances across it (`hasDerivAt_charJacobian`). That kernel is even, f being even and $\cos \circ \Delta\theta$ even by `WindingData.cos_psi_neg`, so every cancellation of this section applies to it unchanged: the characters diagonalize it, and the rate of the mode $\operatorname{Re}\chi'$ is $\lambda_{\chi'}(g) = \sum_d g_d(\operatorname{Re}\chi'(d) - 1)$, the same eigenvalue functional evaluated at the twisted kernel (`charJacobian_chi_re`). The two directions follow termwise. If $f_d \ge 0$ and the winding advances by at most a quarter turn across every separation carrying coupling, then $g_d \ge 0$ and no rate is positive, since $\operatorname{Re}\chi'(d) \le 1$ (`charJacobian_stable_of_quarter_turn`). If some character has a positive rate, the field $e^{\lambda t}\operatorname{Re}\chi'$ solves the linearized equation exactly and its value at the identity site is $e^{\lambda t}$, which leaves every bound (`charJacobian_growing_mode`). This is the shape `incoherent_instability_iff` already carries for the uniform state, at a winding instead.

That criterion names the mechanism the sweep below varies. A kernel reaching only short separations sees phase advances well inside a quarter turn and $g_d > 0$ throughout; lengthening the decay length brings in separations across which the winding has advanced past $\pi/2$, there $g_d < 0$, and a mode crosses zero. What the criterion does not supply is a trajectory. It is a statement about rates at the linearization, and upgrading it to asymptotic stability needs a Łojasiewicz estimate around a critical point that is not a minimum, where the estimate available here — `lojasiewicz_estimate` through `phase_tendsto` — is built around the global minimum. The decay-length boundary reported below is therefore measured and not computed from this criterion, and the states that persist occupy a strictly narrower range than the criterion admits.

A winding is also not where the potential is lowest, and that is a refutation rather than a gap. `potential_min_iff_phase_locked` makes a global minimizer of the dynamic potential phase-locked, `phase_locked_implies_r_sq_eq_one` puts its resultant at one, and `order_parameter_r_sq_char` puts a winding's at zero; so a nontrivial winding is not a global minimizer (`char_not_potential_min`). The quarter-turn hypothesis of `kuramoto_tendsto_global_minimum` does not merely decline to cover this state; it must exclude it, because the conclusion is false of it. The local statement uses the twisted kernel again: the second variation of the potential along $\bar{\theta} + su$ is that kernel's Dirichlet form, which is minus the linearization's quadratic form (`charLineSecond_zero`), so a mode of positive rate is a direction of strictly negative curvature and the potential falls along it at every small enough displacement (`charLinePotential_lt_of_charLambda_pos`). A winding carrying a growing mode is therefore not a

local minimum of the potential on the state space either (`char_not_local_min_of_charLambda_pos`), the line being a continuous curve through it.

A winding is where the patchwise reading of the content bound, Eq. (12), separates furthest from the global one. The global form degrades to $\sqrt{2N}$ here (`chord_bound_char_global`), which constrains nothing at all: every chord is at most 2 (`chord_le_two`), so the estimate is above the trivial bound for every population of two sites or more. Across a nearest-neighbour patch the same inequality gives $2\,|\sin\Delta\theta_g|$, with $\Delta\theta_g$ the phase the winding advances across the patch (`chord_le_of_char_patch`); on the ring $\Delta\theta_g = 2\pi q/n$, so the bound tends to zero as the same winding number is spread over more sites. A population whose resultant is exactly zero therefore still bounds content disagreement inside each patch, by a quantity that shrinks with the resolution of the array. This is the compatibility reading of a low global resultant at sustained patch order.

The same sheet, kernel and integrator as the preceding control measure how far that basin reaches, at identical frequencies so the zero-detuning hypothesis holds. Over 17 decay lengths the twist survives at 9 of them; between the longest length retaining it, 0.1750 mm, and the shortest losing it, 0.2000 mm, the interpolated boundary is 0.1871 mm, or 12.0 lattice spacings. That boundary falls inside the 0.1–0.3 mm band which, for the uniform initial condition of the preceding control, is interior to a plateau. The two are consistent: the plateau records where a uniform state stays ordered, the boundary where a winding one survives, and a sweep from one initial condition constrains neither claim about the other.

What the observable reports there is the point. At 0.1 mm the sheet reaches $r = 0.9657$ from a uniform start and $r = 0.0083$ from a twisted one — both stationary, at the same coupling, diffusion and kernel. Patch-local order in the twisted state is 0.9015 and winding-defect density is at most 0.000000 at every sampled length, so the low resultant records a wave rather than fragmentation. At fixed K the observable thus takes two values spanning nearly its whole range, and K alone does not determine r. Each length runs twice, once twisted and once uniform; the uniform control's gap between patch-local and global order is at most 0.0001 wherever the twist is lost, so the largest gap the twisted runs show, 0.8944, is not an artefact of the patch estimator.

Repeating the sweep under the quenched spread of 1.5 $\mathrm{rad\,s^{-1}}$ the preceding control carries leaves the twist standing at 7 of the 17 lengths, with the boundary at 0.1421 mm. Disorder therefore shortens the range that sustains a wave without closing it, and both boundaries stay inside the empirical band. At 0.1 mm the disordered sheet reaches $r = 0.9418$ from a uniform start against $r = 0.0348$ from a twisted one, at patch-local order 0.8696 and winding-defect density at most 3.5×10^{-4} over the lengths that retain the twist: the spread nucleates defects without dissolving the wave. The uniform branch of this sweep also reproduces the preceding control's steady order at the three empirical lengths from a different initial condition, which is a check on the two implementations rather than a further result. Across 3 frozen frequency realizations at the 9/3 lengths that retain it, retention holds in 9 of 9 runs, at global order between 0.0092 and 0.0570, so the outcome belongs to the model rather than to one sample of the disorder.

Four limits. Below the coupled regime the comparison carries no information: at 0.0125 mm the uniform control itself reaches only 0.3160, so neither initial condition is ordered and the difference between them is not a statement about waves. A twist imposed as an initial condition measures the stability of a basin, not the probability that dynamics enters one. The sweep fixes one winding number, so nothing here describes the spectrum of wave states a cortical sheet might occupy. And nothing identifies a cortical decay length or establishes a transition. The content bound above attaches to a patch of a winding and not to the winding number: which sector a field occupies is connected to no content by any theorem here.

Fig. 3 of the article plots this sweep. Its left panel shows global order against patch-local order for a subset of decay lengths; its right panel compares the twisted and uniform initial conditions across lengths, with the 0.1–0.3 mm band shaded and the interpolated retention boundary marked. The sheet and kernel are those of Fig. S5 above, at identical frequencies.

12.7 An amplitude field, and what it reaches

Every oscillator of the phase model sits on the unit circle, so the state carries a phase and no amplitude, and a field with an amplitude node—$u = \cos(kx)\cos(\omega t)$—has no representation in the dynamics integrated above. Amplitude enters elsewhere in this work only at the measurement layer, where the loop calculation of `Phase5_PhaseLifts.lean` reads winding off locally unwrapped phase maps. The real

Stuart–Landau field supplies the missing degree of freedom on the same coupling:

$$\dot{z}_i = (\Gamma - |z_i|^2)z_i + \sum_j A_{ij}(z_j - z_i), \tag{18}$$

with $z_i \in \mathbb{C}$. This is the gradient flow of $F(z) = \sum_i(|z_i|^4/4 - \Gamma|z_i|^2/2) + \frac{1}{4}\sum_i\sum_j A_{ij}|z_i - z_j|^2$ against the real inner product on $\mathbb{C}^V \cong \mathbb{R}^{2V}$, which is what makes it an extension of the development rather than a replacement: every convergence result above rests on gradient structure. Complex coefficients on the cubic term or on the coupling—the genuine complex Ginzburg–Landau equation—destroy that structure, so the real case is the one that inherits the architecture. `Phase4_KuramotoDynamics.lean` §7 carries the results, on the ring and the circulant kernel of the preceding section.

Four statements follow. `deriv_phase_of_amplitude_trajectory` proves that at a site of positive amplitude the phase of a solution of Eq. (18) obeys $\dot{\theta}_i = \sum_j A_{ij}(r_j/r_i)\sin(\theta_j - \theta_i)$ exactly, with no slow-amplitude assumption: the identity is the imaginary part of $\dot{z}_i\overline{z_i}$ divided by r_i^2. `amplitude_phase_is_kuramoto` specializes it to equal amplitudes, where the ratios are one and the phase is a Kuramoto trajectory of the same coupling. The phase model is therefore this model's phase equation with the amplitude ratios set to one, and the two layers are one object.

`charState_is_amplitude_trajectory_iff` proves that $z_g = a\chi(g)$ is stationary exactly when $a^2 = \Gamma + \lambda_\chi$, where $\lambda_\chi = \sum_d f_d(\mathrm{Re}\,\chi(d) - 1)$ is real by the evenness of the kernel. The site set is any finite abelian group and χ any unit character of it carrying a real phase lift, which is the structure every cancellation in this development actually uses: the ring $\mathbb{Z}/n$ and the periodic sheet $\mathbb{Z}/n \times \mathbb{Z}/n$ are both instances, the second built from the first by `WindingData.prod`. That second one is the geometry the sweep integrates — `travelling_wave.py` imposes $\theta_{rc} = 2\pi qc$/side on a periodic `(side, side)` array under a toroidal kernel, which is this state with one of its two winding numbers set to zero — so the theorems and the simulation now stand on the same object. On the ring `windingState_is_amplitude_trajectory_iff` is the specialization, with $\lambda_q = \sum_d f_d(\cos(2\pi qd/n) - 1)$. The consequence is a band: `exists_windingState_iff` proves a q-fold winding of positive amplitude exists if and only if $\Gamma + \lambda_q > 0$. The phase model admits every winding number at every coupling—`winding_is_kuramoto_trajectory` constrains neither q nor f—so outside the band the two layers disagree about which states exist, which `phase_model_admits_winding_outside_band` states as one theorem.

Read against the sweep, that disagreement is about the state space and not about what a recording would show. Evenness is the theorem's only hypothesis on f, but the kernels this work uses are non-negative, and there $\lambda_q \le 0$ with $|\lambda_q|$ growing as coupling mass moves to longer separations: a broad enough kernel takes $\Gamma + \lambda_q$ below zero and the winding ceases to exist. The preceding sweep reaches the same boundary from the other side, the twist it imposes on the phase model remaining a trajectory at every decay length and ceasing to survive integration past one. Existence and stability fail in the same direction, so the added degree of freedom sharpens the account of the regime rather than predicting a new observation in it. Two things it does not supply. The band is one of existence: no stability result is proved here, and for fields of this class the windings near its edge are unstable in the standard way, so the states that persist occupy a strictly narrower range than the inequality allows. And Γ is the distance from the bifurcation, estimated nowhere in this work, so the band is structural—it predicts no numeral and discriminates no experiment.

Inside the band at $q = n/2$ the phasor is $\xi = -1$, so the state is real and alternates in sign along the ring. Eq. (18) carries no frequency term, being written in a frame rotating at the common carrier Ω that the normal form removes; restoring it gives the lab-frame reading $u_k(t) = \mathrm{Re}(z_k e^{i\Omega t})$, which is $a(-1)^k\cos(\Omega t)$: one spatial profile times one common time factor, every site in antiphase with its neighbours. `standing_wave_of_amplitude_band` assembles this, and it is $\cos(kx)\cos(\omega t)$ at $k = \pi$ in the framework's own integrated dynamics. The phase model has no such object, every site there having unit modulus.

These are windings and not spirals, and on the sheet that is a theorem rather than a limitation of the geometry. The defect (Townsend et al., 2015) and (Xu et al., 2023) measure is a point at which the amplitude vanishes and around which the phase advances by a whole turn; `spatial_kernel.defect_winding` reads it as the wrapped phase circulation around a periodic 2×2 plaquette. `plaquetteCirculation_eq_zero_of_winding` proves that circulation is exactly zero on every plaquette of every state above, for any 2π-periodic odd angle convention: the phase advance across a separation does not depend on where the separation is taken, so the four sides of a plaquette cancel in pairs. `norm_charState` is the reason, putting the modulus at $|a|$ at every site. The extra dimension does not change this. What a torus adds is winding around its two non-contractible loops; a spiral is winding around a contractible one, and that forces an amplitude zero which no state of constant modulus has.

The negative result is sharper still and holds of every field, exact or not. `sum_plaquetteCirculation_eq_zero` proves that on a torus the circulations of all plaquettes sum to zero, each edge being traversed once in each direction, so that `no_isolated_defect` gives the statement in the form that matters: if every plaquette but one carries no defect then neither does that one. A single spiral is not a state of this geometry at all, whatever dynamics produced it — defects arrive in cancelling pairs. This is a constraint on what the layer could ever host, not a gap waiting on more work, and it is why the electrode loop of `Phase5_PhaseLifts.lean` and not the interior it encircles remains what this development speaks about. The sweep's near-zero measured defect density is consistent with the first theorem and is not evidence for anything beyond it: the runs carry noise, diffusion and a quenched frequency spread, and are therefore not exact states.

Three further limits bound the layer, of which the article states the two a reader of the neuroscience meets first. Relaxational dynamics has no propagation speed: real coefficients give no nonlinear dispersion and the drift is exactly zero, so nothing here travels, while a cortical wave measurement always reports a speed. Supplying one needs heterogeneous frequencies or complex coefficients, and both cost the gradient structure. Second, $|z|$ is a normal-form radius and a measured envelope is not; identifying them is a physical assumption of the same kind as the Landauer identification held as a named hypothesis in `Phase3_LandauerBridge.lean`, and it is made nowhere in this work. Third, the normal form is valid near a supercritical Hopf bifurcation, where a resting state has just lost stability; much cortical wave phenomenology is excitable-medium behaviour instead, and the two produce similar-looking maps from different mechanisms. The row of Table S1 recording that no theorem connects winding sectors to content stands unchanged throughout: making a defect computable from a dynamical state rather than postulated from data is not making it mean anything.

Two things do not come for free and are not claimed. A cosine profile is not a stationary state of Eq. (18): the cubic term does not preserve a single Fourier mode, so only the circularly polarized modes are exact and the only real one among them is the alternating mode above. A standing wave with an interior amplitude node is a fixed point of a nonlinear system with no closed form, reached by bifurcation from $\Gamma = -\lambda_q$.

12.8 Falsification conditions

The physical identification would be challenged if independently calibrated field-mediated coupling were insufficient for the proposed regime, or if a selective perturbation of that interaction left the proposed neural and behavioural endpoints unchanged. Such a perturbation requires controls for direct effects on transmembrane currents and other coupling pathways; selectivity is not assumed experimentally available. Recovery tests must distinguish the coupling-gated account from other slow mechanisms under the same observation model. Coherence alone is not equated with report or experience, so dissociations must be evaluated against independently specified endpoints and the additional compatibility and reflexivity hypotheses.

These conditions bear on a coupling field and not specifically on an electromagnetic one. Of the four protocol steps of the main article (Section 9.3), only the spatial-onset comparison of step 3 discriminates, because its field score is built from pre-onset extracellular field amplitude and is where a field enters the protocol. The other steps bear on the coupling hypothesis at large, and extracellular ion diffusion, gap junctions modelled as a density, astrocytic calcium or a mechanical medium satisfy or fail them identically: they manipulate the emergence rate, fix an independent behavioural endpoint, and match diffusion and the observation model, and none of them mentions a field. A null result on the spatial-onset comparison therefore constrains the electromagnetic identification; a null result on the others constrains the coupling-gated account whatever carries the coupling.

13 Coupling support and continuum observables

Two mathematical comparisons distinguish a restriction on wiring support from a restriction on the measure used to evaluate a kernel.

Theorem S4 (Resource-matched coupling reallocation): If a fixed wiring support excludes a pair whose phase correlation exceeds the maximum over its wired pairs, a coupling using that pair can achieve greater total correlation at the same resource budget. This is a conditional support comparison, not a general inequality between biological and silicon hardware.

Coupling reallocation (`Phase7`): The lemma `exists_better_coupling_allocation`, wrapped as the theorem `continuous_beats_rigid_topology`, proves the reallocation step: given a coupling matrix

satisfying `is_strictly_suboptimal` (positive weight on a less phase-correlated pair while a more correlated pair exists), there is another valid coupling with identical total resources and strictly greater total phase correlation.

Lean 4 Implementation (`Phase7_Rigidity.lean`): Two results specify the scope of the comparison. *(i) Rigidity as a support constraint.* `RealizableIn S R A` says A is a valid coupling meeting resource budget R and vanishing on every pair the architecture has no wire for — weights tunable, wiring fixed. `totalCorrelation_le_of_realizable` proves no realizable coupling exceeds $R \cdot M$ when M bounds phase correlation over the wired pairs, and `exists_realizable_pair` attains it. `rigid_is_strictly_suboptimal` then *derives* `is_strictly_suboptimal` for every realizable coupling, whenever some pair (k, l) the wiring misses has correlation exceeding M: the hypothesis Theorem S4 would otherwise have to assume is a consequence of the wiring pattern. `rigid_gap` gives the quantitative form, a gain of at least $R\,(\cos(\theta_l - \theta_k) - M) > 0$ at equal budget. `Examples` §8 witnesses this on three sites wired only to the maximally anti-correlated pair, where the gap is exactly 2, and runs `exists_better_coupling_allocation` on it with no assumed hypothesis. *(ii) Continuity as a measure-theoretic distinction.* `fieldCorrelation` integrates a kernel against $\mu \otimes \mu$ as in the spatial coupling model. A finitely-sited architecture is a `Finset X`; `fieldCorrelation_sited_eq_zero` proves that on an atomless substrate such a kernel contributes exactly zero, for every phase field and every choice of weights, while `fieldCorrelation_patchKernel` gives $c\,\mu(U)^2$ for a patch. `sited_architecture_below_field_optimum` combines them; `Examples` §8 instantiates it on $\mathbb{R}$ with a thousand units of weight at the origin against the unit interval.

13.1 Finite regions preserve the discrete observable

`Phase7_FiniteRegion.lean` embeds a finite matrix A in a `KernelMesh` whose cells C_i have positive masses $m_i = \mu(C_i)$ under a finite measure. Put

$$K(x, y) = \frac{A_{ij}}{m_i m_j}, \qquad \theta(x) = \theta_i \quad (x \in C_i,\ y \in C_j).$$

Integration of a function of cell indices is its mass-weighted sum (`integral_locate_pair`), hence

$$\iint K(x, y)\, d\mu(y)\, d\mu(x) = \sum_{i,j} A_{ij}, \tag{19}$$

$$\iint K(x, y) \cos(\theta(y) - \theta(x))\, d\mu(y)\, d\mu(x) = \sum_{i,j} A_{ij} \cos(\theta_j - \theta_i). \tag{20}$$

These are `kernelIntegral_cellKernel` and `fieldCorrelation_cellKernel`. Symmetry and nonnegativity are preserved, and a zero matrix diagonal makes each within-cell block zero. On an atomless substrate, a reconstruction with nonzero correlation cannot be supported on finitely many points (`cellKernel_not_sitedOn`).

`Examples/FiniteRegion.lean`, §26, uses Lebesgue measure on $[0, 4)$, cells of masses one and three, and a matrix with total weight twelve. Both the embedded weight and correlation equal twelve exactly. Thus finite point support and finitely many positive-mass regions have different consequences for the same integral.

The same module supplies a control for the wiring gap: `no_forced_gap_of_best_wired` constructs an optimal allocation when the support contains both directions of a distinct pair whose correlation bounds every pair, including diagonal pairs. This hypothesis forces the chosen pair's correlation to be one. The three-site witness wires the perfectly correlated pair and attains the resource bound, so no disadvantage is forced by fixed wiring itself.

The matched resource in the cell embedding is total coupling weight; pricing it as energy requires hardware data. The embedding is piecewise constant, with no continuity, dynamical approximation or refinement limit established here. Point-supported vanishing is about measure-zero support, and the strict wiring gap requires the specified missing pair. Neither result excludes digital hardware from approximating field dynamics or supporting consciousness.

14 Symmetry breaking and the emergence of boundaries

A boundary is one of the physical structures the broader composition assumes. Under the stated vacuum and connectedness hypotheses, a field reaching two vacuum components must leave the vacuum some-

where. Relating that locus to a finite information-processing register is a separate physical premise; the result supplies neither a general condition for agency nor a derivation of the cortical coupling model.

Let $\mathcal{M}$ be a pseudo-Riemannian manifold representing spacetime, and let $\phi : \mathcal{M} \to \mathcal{V}$ be a physical field mapping to a vacuum manifold $\mathcal{V}$ (the space of energy minima). Let G be a continuous symmetry group acting on $\mathcal{V}$.

Theorem S5 (Domain-wall obstruction): A continuous field on a connected substrate whose image reaches two different connected components of the vacuum manifold must leave that manifold somewhere. This is the π_0 case of the topological-defect argument. A separate extension obstruction states that non-null-homotopic boundary data admit no continuous extension over the specified contractible interior.

Lean 4 Implementation (`Phase1_Primitives.lean`): We formalize this by defining the set of degenerate energy minima `DynamicalVacuum` and proving `boundary_defect_forces_interior_vacuum_break`: if a boundary field configuration ϕ_∂ is not null-homotopic, then it admits no continuous extension over a contractible interior. This is an obstruction to extending the specified boundary data. Note the direction: this result establishes that a defect *obstructs extension*, not that a defect must exist. The companion lemma `contractible_interior_forces_trivial_boundary` proves the contrapositive.

A direction that produces a defect is proved, at the level of π_0. `exists_notMem_of_no_common_preconnected` states it in the form a witness can discharge — if no preconnected subset of the vacuum manifold M contains both $\phi(x_1)$ and $\phi(x_2)$, and the substrate is connected, then ϕ leaves M somewhere — and `exists_notMem_of_connectedComponentIn` restates it with $\phi(x_2) \notin$ the connected component of M containing $\phi(x_1)$. The proof uses the connectedness of continuous images and `IsPreconnected.subset_connectedComponentIn`, and nothing else; no hypothesis mentions energy, a potential or a symmetry. This is the domain-wall case of Kibble's classification and we take only that case: $\pi_1 \neq 0$ forcing vortices and $\pi_2 \neq 0$ forcing monopoles are recorded as available and not taken, since the manuscript's argument uses domain walls alone.

Separately, `spontaneous_symmetry_breaking` shows that a global energy minimizer lies in the vacuum manifold almost everywhere. The proof introduces no additional axioms and uses `pointwise_vacuum_of_global_min`.

`Examples` §11 witnesses that result on the double well $V(v) = (v^2 - 1)^2$ over a two-point spacetime carrying a Dirac probability measure. Two features of the witness are deliberate. Its minimization hypothesis is *proved* rather than assumed — everywhere else in the development the existence of a global minimizer is a hypothesis, since no compactness or direct-method argument is formalized. And the almost-everywhere in the conclusion is shown necessary: the minimizing field sits at the vacuum on the support and at 0 off it, so the pointwise version of the statement is false of the very field the theorem applies to. The same section instantiates the $\mathbb{Z}_2$ action $v \mapsto \pm v$ and discharges the energy invariance, giving the first instance of `SymmetryInvariantAction`; the vacuum manifold $\{-1, 1\}$ is proved degenerate and the two minima exchanged by the symmetry, so a minimizer and its image have equal energy and differ.

`Examples` §11.1 then supplies the witness for the π_0 theorem on the same well. The vacuum manifold is *computed* there, `wellVacuum_eq` proving `DynamicalVacuum` $V = \{-1, 1\}$ exactly rather than merely containing both minima — which matters, because a larger vacuum set might have been connected — and `wellVacuum_separated` proves the disconnection in the form the theorem consumes, by the observation that a preconnected subset of $\mathbb{R}$ spanning -1 and 1 contains 0, which is not a minimum. `wellV_domain_wall` is the resulting statement about *every* continuous field on *any* connected substrate reaching both minima. Non-emptiness of that class is checked separately rather than assumed: the clipped identity on $\mathbb{R}$ satisfies the hypotheses, and `kink_zero_notMem` locates its forced wall at the origin. These results concern vacuum membership and domain walls; they construct no conserved quantity.

15 The thermodynamic constraint: Landauer's principle

For a finite effective register, the number of distinguishable histories that can be retained is bounded. This does not imply that every update erases. The capacity theorem states that a distribution on a finite state space X carries at most $\log |X|$ nats, with equality only at the uniform law (`shannon_entropy_eq_log_card_iff`), and a perturbation stream of length k over an alphabet P with $|P|^k > |X|$ therefore cannot be recorded injectively (`absorb_not_injective`), whatever the update map is and even when every individual step of it is reversible.

Let the global phase space $\Omega = \Omega_{sys} \times \Omega_{bath}$ undergo globally reversible Hamiltonian dynamics, modeled as an injective mapping $f : \Omega \to \Omega$.

Theorem S6 (Thermodynamic Erasure): Under the specified finite-register and reversible system–bath assumptions, the heat identity gives a lower bound proportional to the entropy erased by the register. Erasing one unbiased bit gives $\Delta Q \geq k_B T \ln 2$. Non-injectivity alone does not specify a universal one-bit entropy loss; the bound uses the erased entropy of the particular update.

The same heat, in three vocabularies (`Phase3_LandauerBridge.lean`). The erasure bound charges a register for destroying information, while the predictive bound charges a structure for keeping information it cannot use. `PredictiveDissipation.ofLandauer` builds the predictive structure of a given register: its thermal scale is that register's temperature, its dissipated work is that register's Landauer heat, and Still's bound is derived from `landauer_bound` rather than assumed a second time. The single hypothesis is the identification $(I_{\text{mem}} - I_{\text{pred}}) \leq S(\text{id}) - S(t)$ — the memory wasted is no larger than the entropy erased — and it is a restriction, not bookkeeping: on an injective update the right-hand side is zero and the inequality forces the waste to vanish (`nonpredictive_eq_zero_of_injective`), which is the Norton–Shenker distinction in informational form. The measure-theoretic statement of Landauer (`Phase3_MeasureThermodynamics.lean`) is specialized in the same file rather than left standing beside the combinatorial one: at counting measure on a finite register, `continuous_entropy` is `boltzmann_entropy`, volume contraction is non-injectivity, and `landauer_heat_bound` is T times the entropy erased (`landauer_heat_bound_count_eq`). The witness is the one-bit eraser, where a register destroying $\log 2$ pays for a memory wasting $\log 2$, with equality (`Examples` §18.5).

Lean 4 Implementation (`Phase3_CombinatorialThermodynamics.lean`): We explicitly decouple logical erasure from physical erasure. Lean verifies that any contraction in the *physical* state space ($\neg$Function.Injective(t_{sys})) under globally reversible dynamics mathematically necessitates heat dissipation. The combinatorial core, `landauer_from_reversibility` and `second_law`, is fully derived. Finiteness of the phase space is what converts the physically checkable hypothesis into the one the bound needs: on a finite state space a map that fails to reach every state is not injective (`is_erasure_of_not_surjective`), so an unreachable state suffices for $\Delta Q > 0$ (`finite_phase_space_dissipates`). The hypothesis is fenced: `Nat.succ` misses 0 and is injective, so the implication is false on an infinite phase space, and `Examples` §19 records this beside the witnesses. The Norton–Shenker distinction survives the identification — retaining the incoming perturbation is injective and free (`keepRecord_injective`), refreshing the register that held it is not (`absorbStep_is_erasure`). Converting it into a statement about heat requires the postulate $\Delta Q = T\Delta S_{bath}$, which is carried as the class field `StatisticalMechanics.heat_eq`; `Examples` §1 exhibits a one-bit erasure system with an explicit bath that discharges it, so the postulate is demonstrably satisfiable. A register that satisfies the stated erasure and thermal hypotheses dissipates; identifying it with a biological boundary is an additional modelling step.

16 Predictive memory and feedback

For passive memory X_t driven by an environmental signal S_t, the predictive bound is

$$W_{\text{diss}} \geq k_B T \left[I(X_t; S_t) - I(X_t; S_{t+1}) \right]. \tag{21}$$

The Markov structure $X_t \to S_t \to S_{t+1}$ makes the bracket nonnegative (Still et al., 2012). Zero retained information also satisfies it, so a task and a learning rule require additional specification. Relating this bound to Landauer heat requires a common thermal scale and budget, and memory waste bounded by erased entropy.

An acting agent changes the signal it will receive. Specify policy, world, observation and memory-update channels. For their finite laws,

$$\begin{aligned} I(X_t; S_t) - I(X_t; S_{t+1}) = I(X_t; S_t \mid S_{t+1}) \\ - I(X_t; S_{t+1} \mid S_t). \end{aligned} \tag{22}$$

The second conditional term vanishes in the passive model. With feedback, the signed difference can be negative as action creates correlations with the future. A causal reading requires an environmental state sufficient to describe the relevant dynamics.

For a controlled substep, the active thermodynamic bound instead uses the joint entropy. Forward and reverse path laws give $\Sigma = D_{KL}(P_F \| P_R) = \Delta H(X, S) + \langle Q \rangle / (k_B T) \geq 0$ under local detailed balance, with thermal scale $k_B T > 0$ and heat positive into the reservoir (Horowitz and Esposito, 2014; Ito and Sagawa, 2013). Every forward path of positive probability must have a positive-probability reverse for this finite identity. A missing reverse gives infinite divergence. The support conditions, and

the composition of actuation with memory updates on their shared intermediate law, are developed in Section 16.2.

A bound relating environmental and internal distributions must be a physical assumption about specified processes, not a consequence of Gibbs' inequality alone. The class field `StructuralResonance.kl_bound` states $D_{KL}(P\|Q) \le \Delta t\,\sigma$ for an instance. Even under this assumption, descent of σ does not imply $D_{KL} \to 0$ unless the upper bound vanishes. A positive limiting upper bound permits a positive divergence. No convergence of internal geometry to environmental structure follows.

Theorem S7 (Nonpredictive information bounds dissipation): Let X_t be the system's state and S_t an environmental signal evolving by a dynamics of its own on which the system does not act. Then

$$I(X_t; S_{t+1}) \le I(X_t; S_t), \qquad W_{\mathrm{diss}} \ge k_B T\,[\,I(X_t; S_t) - I(X_t; S_{t+1})\,], \tag{23}$$

the second being the bound of (Still et al., 2012). Consequently $W_{\mathrm{diss}} \ge 0$, and any bound on the dissipated work is a bound on the nonpredictive information $I(X_t; S_t) - I(X_t; S_{t+1})$.

The first inequality is the data processing inequality for the Markov chain $X_t \to S_t \to S_{t+1}$, proved in `Phase3_PredictiveThermodynamics.lean` as `predictiveInfo_le_mutualInfo`; mutual information is defined there from Mathlib's `klDiv` as the divergence of the joint law from the product of its marginals, since Mathlib carries no mutual information of its own. The second inequality is the class field `PredictiveDissipation.still_bound`, an irreducible physical postulate for the same reason as every other one here — its `axiom` form is refutable, and `Examples` §18.4 refutes it. Everything downstream is derived: non-negativity of the dissipated work, the bound of the nonpredictive information by the dissipation budget, and the statement that a quasi-static drive leaves every retained bit predictive. `Examples` §18 inhabits the class twice over one correlated two-bit law — a frozen signal, permitted to dissipate nothing, and a scrambled one, forced to dissipate its entire memory — and fences the theorem with a Markov kernel on the pair that writes the system's state into the signal and thereby *raises* mutual information, so the no-back-action hypothesis is shown to carry weight rather than to decorate. The equality of (Kawai et al., 2007), $\langle W_{\mathrm{diss}} \rangle = k_B T\, D_{KL}(\text{forward}\,||\,\text{reverse})$, is carried as a predicate on an instance and yields `nonpredictive_le_arrow`: the irreversibility of the dynamics is at least the memory that fails to earn its keep.

The information constraints distinguish erased entropy from retained nonpredictive information (Rosas et al., 2026). Estimating either side from data requires a separate observation model. A partially observed signaling model derives finite-time estimates of information gathering and dissipation from counting statistics (Nicoletti et al., 2026); a corresponding cortical estimator is not supplied here.

16.1 Agency and the scope of the passive bound

The stochastic channels of `Agency` (`Phase3_Agency.lean`) specify one perception–action cycle:

$$\begin{aligned} A_t &\sim \pi(\cdot \mid X_t), & S_{t+1} &\sim K(\cdot \mid S_t, A_t), \\ O_{t+1} &\sim G(\cdot \mid S_{t+1}), & X_{t+1} &\sim \mathcal{U}(\cdot \mid X_t, A_t, O_{t+1}). \end{aligned}$$

The sampled action is retained through actuation and supplied to the memory update. The history used for predictive information records (X_t, S_t, S_{t+1}), keeping memory at its pre-action state. Its present marginal is the input law and its future marginal is the post-action world paired with that pre-action memory. The environmental state must be sufficient for the declared Markov dynamics; an observation alone need not satisfy this condition.

`Feedback.information_balance` proves Eq. (22) from the existing KL-based mutual information. Conditional information uses the chain-rule expression $I(X;(S,S')) - I(X;S')$ or $I(X;(S,S')) - I(X;S)$, respectively. Data processing proves both nonnegative under finite joint information, and `mutualInfo_finite` discharges that condition on finite discrete spaces, including laws with zero atoms. Real subtraction preserves the sign of their difference. The passive constructor retains the present signal while evolving it through its own kernel; `passive_feedback_zero` and `passive_signedWaste` recover the passive information statements exactly. The witness `writingAgent` selects its own bit as the action, writes it into an initially independent world bit, senses the result and updates its memory. Its signed difference is strictly negative, and its composed cycle equals the explicit write kernel. Conditional dependence identifies causal influence only with an adequate causal model.

16.2 A finite thermodynamic model with feedback

For a finite controller–world law $p(x, s)$ and a normalized channel $K_x(s' \mid s)$, one elementary update holds x fixed. The final law and the forward and reverse path laws are

$$\begin{aligned} p'(x, s') &= \sum_s p(x, s) K_x(s' \mid s), \\ P_F(x, s, s') &= p(x, s) K_x(s' \mid s), \\ P_R(x, s, s') &= p'(x, s') K_x(s \mid s'). \end{aligned}$$

Both path laws are normalized. The reverse uses the same autonomous transition channel, started at the forward final distribution. `FiniteFeedbackStep.entropy_balance` assumes strictly positive initial masses and transition probabilities. `entropy_balance_of_support` assumes only that every forward path of positive probability has positive reverse probability, which is what the expansion uses and which strict positivity implies (`reversibleSupport_of_positive`). Time-dependent protocols and unidirectional transitions require a corresponding treatment of reversal and support.

`FiniteFeedbackStep.entropy_balance` (`Phase3_AgencyThermodynamics.lean`) expands the path log-ratio to obtain

$$\Sigma = D_{KL}(P_F \| P_R) = H(p') - H(p) + \sum_{x,s,s'} P_F(x, s, s') \log \frac{K_x(s' \mid s)}{K_x(s \mid s')} \geq 0. \tag{24}$$

Gibbs' inequality proves nonnegativity; the balance is derived from these distributions. `ProbDist.KL_eq_klDiv` identifies the finite divergence with Mathlib's measure-theoretic divergence for a positive reference law. The information identity uses the same forward law and its proven initial and final marginals.

Support inclusion is the weaker of the two conditions and the one the identity needs. Gibbs' inequality holds under it (`KL_nonneg_of_support`), so zero masses in the initial law or in a transition are permitted and $\Sigma \geq 0$ survives them; this matters because an agent's executed channels are sparse, and a step whose initial law sits on a single atom satisfies the support condition while failing positivity. What the weaker condition does not do is make every channel priceable. A forward path with no reverse — deterministic erasure is the standard case — has infinite divergence, and `ProbDist.extendedKL` records it as such in $[0, \infty]$ rather than returning the finite number the totalized real logarithm would supply. Where support does hold, that extended value is the real divergence above. `ProbDist.extendedKL_eq_klDiv` and `extendedKL_eq_klDiv_of_missing` (`Phase3_MeasureFeedback.lean`) identify the extended value with Mathlib's measure divergence on both branches — the finite one under a positive reference law and the infinite one wherever a forward atom has no reverse — so the finite convention here is the measure-theoretic divergence restricted to finite spaces, not a second definition standing beside it.

Local detailed balance is the physical identification $q(x, s, s') = k_B T \log[K_x(s' \mid s)/K_x(s \mid s')]$, with heat positive into the reservoir (Horowitz and Esposito, 2014; Ito and Sagawa, 2013). It gives $\langle Q \rangle \geq k_B T[H(p) - H(p')]$. The first-law hypothesis $w = \mathcal{E}(x, s') - \mathcal{E}(x, s) + q$ separately fixes work. These are bounds on the joint controller–world system. The passive signed information difference alone is insufficient to supply a general work bound. `Chain.active_entropy_budget` gives the scalar budget consequence of `ActiveBound`, the fourth node of the active composition.

16.2.1 Composing actuation and memory updates

`FiniteFeedbackCycle`, in the same module, carries an actuation step and a memory channel $\mathcal{U}_t(y \mid x)$. Its sensing step starts at $p_1^{\text{swap}}(t, x) = p_1(x, t)$, where

$$p_1(x, t) = \sum_s p_0(x, s) K_x(t \mid s), \qquad p_2(y, t) = \sum_x p_1(x, t) \mathcal{U}_t(y \mid x). \tag{25}$$

The second initial law is constructed from the first final law. The theorem `cycle_law` identifies p_2 with the output of the composed `Agency.cycle` channels, with action $a = x$ and observation $o = t$. Strictly positive $p_0, K, \mathcal{U}$ give positive substep path laws. At a common thermal scale $k_B T > 0$, local detailed balance for each channel yields

$$\Sigma_A + \Sigma_M = H(p_2) - H(p_0) + \frac{\langle Q_A \rangle + \langle Q_M \rangle}{k_B T} \geq 0. \tag{26}$$

The intermediate entropy cancels because exchanging coordinates preserves joint entropy. No intermediate or final law is assumed stationary. `entropy_budget` therefore gives $H(p_0) - H(p_2) \le B/(k_B T)$ whenever $\langle Q_A \rangle + \langle Q_M \rangle \le B$.

With one energy observable $\mathcal{E}(x, s)$, the separate first-law hypotheses are $w_A(x, s, t) = \mathcal{E}(x, t) - \mathcal{E}(x, s) + q_A(x, s, t)$ and $w_M(t, x, y) = \mathcal{E}(y, t) - \mathcal{E}(x, t) + q_M(t, x, y)$. `mean_first_law` cancels the same intermediate expected energy and proves

$$\langle W_A \rangle + \langle W_M \rangle = \mathbb{E}_{p_2}\mathcal{E} - \mathbb{E}_{p_0}\mathcal{E} + \langle Q_A \rangle + \langle Q_M \rangle. \tag{27}$$

These results cover the two autonomous channel substeps with action and observation aliases. Extra physical registers, time-dependent control and learning require specified dynamics and costs. These balances supply no upper budget; the register ledger below allocates one.

16.2.2 A reciprocal two-bit witness

In units $k_B T = 1$, let the interaction energy be zero for $x = s$ and $\log 3$ for $x \neq s$. Each substep redraws one bit, agreeing with the fixed other bit with probability $3/4$. Actuation updates s using action $a = x$; sensing updates x using observation $o = s$. The action and observation are aliases of existing states and require no separate storage register. Their physical content is the specified pair interaction and thermal updates. The alignment preference is supplied by that interaction.

Starting from independent uniform bits, actuation produces probability $3/8$ on each agreeing pair and $1/8$ on each disagreeing pair. Both marginals remain uniform. The subsequent sensing step preserves this joint law while allowing individual memory transitions. In `Examples/AgencyThermodynamics.lean`, the exact actuation heat is $(\log 3)/4$ and entropy production is $\log 2 - (\log 3)/2 > 0$; the sensing step has zero mean heat and entropy production. Each substep satisfies detailed balance and the first law with zero external work. The positive actuation heat comes from a decrease in mean interaction energy. The example accounts for relaxation of an initially nonequilibrium law; a maintained supply of useful work or a learning process would require additional dynamics and resources.

`Examples/AgencyCycle.lean` specializes the general composition, total heat budget and first law to these same channels and laws. A biased initial memory law gives an asymmetric intermediate law and a sensing step that changes the joint distribution, with positive actuation heat. The regression `mismatched_intermediate_rejected` excludes a second step initialized without the required coordinate swap; the symmetric equilibrium law alone cannot detect that error. The reciprocal cycle also rejects a zero total heat budget.

17 Finite policy selection and learning

The finite models of this section separate policy selection, learning from a phase-derived sensor and repeated operation. A binary closed-loop witness lets a register install the phase energy landscape, lets the phase evolve under a bath-supported channel, senses the new phase and updates that same register. A fixed decoder trained on one family predicts the next phase on a held-out family with accuracy 0.783 under an informative sensor, versus 0.552 under a matched blind sensor. Their marginal phase order is the same. The executed law also charges installation, sensing, reset and preparation against a finite work store; phase-flip bath heat is recorded separately. The symbol task and mode prices are supplied toy inputs, so this does not identify a cortical mechanism or a heat cost for the full sensing and reset protocol.

For a complete cycle with a common energy observable $\mathcal{E}$ and thermal scale $k_B T$, positive forward and reverse path support, and local detailed balance at each stage, the conditional joint-law budget is

$$k_B T[H(p_0) - H(p_N)] \le Q_{<N} = W_{<N} + \langle \mathcal{E} \rangle_0 - \langle \mathcal{E} \rangle_N \le B + \langle \mathcal{E} \rangle_0 - \langle \mathcal{E} \rangle_N, \tag{28}$$

when expected cumulative work stays within the supplied allowance B at every prefix. This bounds expectations. The numerical witness specifies a bath for the phase transition; its sensing, reset and preparation prices do not establish reverse protocols or heat for the complete cycle. Separate constructions track a store on each trajectory and specify when a mean allowance represents an available resource; an expected budget alone permits overdrawn paths.

For the declared process, a finite source with positive net cost funds finitely many cycles. A common physical ledger for the sensing and reset channels, and a cortical calibration of every mode price, remain separate obligations.

These models ask three different questions. A finite optimization selects a policy for a supplied task and budget. A thermal register adapts toward a reward encoded in its energy. An observation-driven register learns correlation with an unknown task parameter. Their preparations, clocks and energy assignments are specified separately; policy comparison does not itself execute a learning process. Table S1 maps each construction to its Lean declarations, and the modules named below carry the witness arithmetic.

17.1 Policy selection under a heat budget

`FiniteControlProblem` supplies a common initial law, world channel, memory channel and terminal reward. A deterministic policy $a = \pi(x)$ induces a complete update whose final law `cycle_law` identifies with the channels carrying that policy. Action is a function of existing memory, observation is the world state, and no stochastic action path is discarded before heat is calculated. The reservoir model assigns each substep the transition log-ratio of its own channel at thermal scale $k_B T$, which is the local-detailed-balance modelling input; performance and total mean heat are read from the same process, including its policy-dependent intermediate law. For a finite permitted set Π with a feasible member, `exists_optimal` proves

$$\exists \pi_* \in \Pi : \quad Q(\pi_*) \le B, \qquad J(\pi) \le J(\pi_*) \quad \text{for every } \pi \in \Pi \text{ with } Q(\pi) \le B. \tag{29}$$

Neither strict improvement nor a task objective follows from existence alone.

`Examples/AgencyControl.lean` supplies a lamp task with all four deterministic bit policies permitted, and proves their exact expectations from the finite path laws. At budget $B = (\log 3)/4$, `unique_optimum` identifies copying as the unique feasible maximizer: it doubles performance over a baseline of equal heat and equal mean work, the constant-true policy performs better and is infeasible, and copying and inversion reach identical final joint laws at unequal heat, so even the full endpoint law cannot replace path accounting. The reward, permitted policies, initial resource and budget are declared inputs. Selection compares fixed autonomous policies for one update and accounts for neither installing nor switching nor learning them.

17.2 Adaptation toward an encoded objective

A physical bit b indexes the two feasible lamp policies, and at thermal scale one its register energy $E_L(b) = 4 \log 3 \, [J(\pi_1) - J(\pi_b)]$ encodes their supplied task values against the channel

$$T = \begin{pmatrix} 5/8 & 3/8 \\ 1/8 & 7/8 \end{pmatrix}, \tag{30}$$

rows and columns ordered $0, 1$. Every transition has outward heat $E_L(b) - E_L(c) = \log[T_{bc}/T_{cb}]$ and zero work; no observation or reward estimator is part of this model. `FiniteFeedbackStep.iterate` initializes each update from the preceding final law, and its telescoped entropy balance and first law preserve strict support. In `Examples/PolicyLearning.lean` the occupancy converges geometrically to $\Pr(b = 1) = 3/4$, so `performance_improves` gives strict expected improvement at every finite update while `performance_limit` places the limit at $7/16$, below the better policy's own value. Cumulative heat is bounded by $(\log 3)/2$ and comes from the prepared register's relaxation, so this finite total supplies no continuing power source and does not bound the cost of repeatedly performing the task. Improvement depends on the objective, channel and preparation: `reward_decreasing_path_positive` retains a copy-to-baseline transition, and `overprepared_reward_decreases` lowers expected reward from an initial occupancy of $7/8$.

17.3 Learning an unknown task from observations

Which of two actions is rewarded is an environmental bit w the learner never reads. Its register bit is read out as the executed action, the world returns success with probability $3/4$ on a match and $1/4$ otherwise, and the update is win-stay, lose-resample. The fences are signatures rather than prose: in `FiniteObservationalLearner` the update has type $O \to R \to \mathrm{ProbDist}\, R$ and the readout $R \to A$, so no channel can read the parameter and no precomputed policy value can be decoded into the action, and the reward enters `performance` and no channel. Acting, observing and updating compose into one `FiniteFeedbackStep` on (W, R) at fixed parameter, so the existing path-law entropy balance, local

detailed balance and first law apply to exactly the process the reward is read from. Its composite channel is

$$T_w = \begin{pmatrix} 5/8 & 3/8 \\ 1/8 & 7/8 \end{pmatrix} \quad (w = 1), \tag{31}$$

rows and columns ordered $r = 0, 1$, with its mirror image at $w = 0$ and every entry positive. From the uniform prior `law_apply` gives every mass at every horizon, so the executed action's expected reward is $J_n = 3/4 - (1/4)(1/2)^n$: `performance_improves` is strict improvement from $J_0 = 1/2$, `performance_limit` places the limit at $3/4$ and `performance_below_certainty` keeps every term below one. `action_marginal_unchanged` executes each action half the time at every horizon: what the learner acquires is correlation with the parameter, not a preference between actions.

Two baselines share the task, prior, readout and allowance. A frozen register and a world returning success with probability $1/2$ whatever is done both stay at $1/2$ for every n, the second spending nothing, and `improves_on_baselines` states the strict improvement over both within the common allowance. The claim therefore rests on the observations carrying task information rather than on the update rule or the heat, and `observation_informative` exhibits that information as conditional on the action while `outcome_marginal_uninformative` shows the outcome's marginal carries nothing about w at the prior. Heat is the transition log-ratio of the same composite channel at thermal scale one, evaluated on the actual evolving laws:

$$\langle Q_n \rangle = \frac{\log 3}{8} \left(\frac{1}{2}\right)^n, \qquad \sum_{n=0}^{N-1} \langle Q_n \rangle = \frac{\log 3}{4} \left[1 - \left(\frac{1}{2}\right)^N\right] \leq \frac{\log 3}{4}. \tag{32}$$

The register's two states are energetically degenerate, so it stores nothing and `cumulative_work` makes that whole sum external work. The work is irreducible rather than a modelling convenience: `zero_work_needs_parameter_independent_heat` proves that a register energy giving zero work on every path forces the heats at the two parameter values to agree, and `no_zero_work_energy` exhibits their disagreement, so an energy landscape that paid for this learning would be one that already encoded the answer. Expected improvement is not pathwise improvement, the register leaving the rewarding action with probability $1/8$. Preparing the prior, supplying the work that drives the updates, executing task episodes beyond the declared updates, convergence to an optimal policy and any cortical identification are outside this model.

17.4 A sensor channel computed from phases

In `Phase3_PhaseSensor.lean`, a `PhaseSensor` declares a phase configuration $\theta(w, a, v)$ for environmental parameter w, executed action a and oscillator v, together with a finite-symbol readout $h : \mathbb{R} \to O$. Sampling one of $N > 0$ oscillators uniformly gives

$$\Pr(o \mid w, a) = \frac{|\{v : h(\theta(w, a, v)) = o\}|}{N}. \tag{33}$$

The channel probabilities are derived from those data. The phase law, sampling rule and readout remain supplied; a general real-angle readout here is not required to satisfy the periodic Lipschitz condition of the content encoders. Two parameters give different channels at an action exactly when some symbol's fibre counts differ (`channel_ne_iff_fiberCount_ne`).

`FiniteObservationalLearner.ofPhaseSensor` feeds this channel into the existing learner without giving its update or action readout access to w. In `Examples/PhaseSensor.lean`, §27, four oscillators take phases 0 or π, read as success or failure. A matching action gives three success symbols and one failure; a mismatching action reverses those counts. The resulting learner is proved equal to the learner in Section 17.3, so its entire law and reward formula apply. The control always has three success symbols and one failure, independently of parameter and action, and keeps expected reward $1/2$. Both populations have $r^2 = 1/4$ for every parameter and action; the informative learner strictly outperforms the control after any positive number of updates. A constant readout also removes the advantage while preserving the phase law.

These witnesses distinguish information carried by a phase readout from coherence. They derive no phase configuration from Kuramoto dynamics and no learning advantage from increased coherence.

18 Repeated operation, memory and energy supply

An executed sequence carries each final law into the next operation. Its mean thermodynamic balance, solvency on individual trajectories and physical source are separate specifications, and the constructions below identify which resource is supplied and which operation spends it. Sensor memory adds recording and clearing to the protocol; preparation adds a stage before its declared initial law; the energy-boundary results state the conditions needed to connect a drawn unit to physical work. Each paragraph names the module that carries it, and Table S1 maps the declarations.

A continuing agent and its expected work allowance. `FiniteProtocol` in `Phase3_ContinuingAgent.lean` allows a different channel at every stage, carrying each actual final law into the next initial law, and `sum_entropy_balance`, `sum_first_law` and `cumulative_entropy_budget` telescope the finite path balances over it without stationarity. `ContinuingProcess` adds a common energy, stage heat observables, a thermal scale and an initial work allowance B; its predicate `Sustains` N requires $B - W_{<k} \geq 0$ at every $k \leq N$ for work read as an expectation on the executed paths, which is the article's Eq. (28). When each stage costs at least $c > 0$, `horizon_le_of_cost` gives $Nc \leq B$. These bounds include the system's mean energy drop and do not identify the allowance with a register's erasure heat.

`ContinuingAgent` implements a three-stage cycle on (w, r, e): the parameter w is fixed, the register r selects the action and the environment flag e retains what that action did, with acting, learning and reset each driving one coordinate while the others drift. The learner's readout, observation and update do not take w as an argument; the actuator responds to it, as a physical task must, and reward is used only to score the law. On the bit-valued witness of `Examples/ContinuingAgent.lean` the proved mass recursion gives expected task performance rising geometrically from below to 65/127 and staying short of one (`performance_error`, `performance_tendsto`), while `first_cycle_work` and `totalWork_cycles_lower` price the cycles: the learning stage costs $(\log(5/3))/16$ every cycle, reset has strictly positive mean heat, and the resulting per-cycle floor exceeds any declared allowance at some cycle (`work_unbounded`). The sequence that converges is one no finite store funds to its limit, so reaching that limit is a statement about replenishment rather than about learning. Two controls delimit it: uniform observations leave chance reward and zero learning-stage heat, and replacing reset by idle drift still learns while changing the workspace law and subsequent behaviour. The store limits expected work at each prefix and not the battery on every realization, which is what the next construction supplies.

A store carried on each trajectory. A store reading β with draw w^- and supply w^+ obeys the following ledger and funding condition on reachable transitions $s \to t$:

$$\beta(t) = \beta(s) - w^-(s,t) + w^+(s,t), \qquad w^-(s,t) \leq \beta(s) + w^+(s,t) \tag{34}$$

`PathwiseStore`, in the same module and witnessed in `Examples/PathwiseStore.lean`, makes the reading a *coordinate of the state*, which is what makes the account pathwise: `Reachable` is positive mass under the protocol's own law, `Ledgered` is Eq. (34) on every executed transition out of a reachable state, `Solvent` N says no reachable state carries a negative reading, and `solvent_of_funded` is the conclusion. No positivity hypothesis appears, so deterministic gates and zero masses are admitted, as they must be for the register–bath model of Section 19.3. Averaging the ledger over the executed paths gives `totalDraw_le_of_solvent`: pathwise solvency implies the expected bound that `Sustains` asserts. Two results fence what a supply buys — a supply covering each executed draw sustains every horizon, and a net draw of at least $c > 0$ on every executed transition before N gives $Nc \leq \beta_0$ — and two witnesses separate the readings. The *gambler* satisfies the mean predicate with equality while running overdrawn at probability 1/4 (`sustained_not_solvent`), so the refinement runs one way only. The *charger* spends a unit and is given it back, proving solvency at every horizon while drawing without bound, satisfies neither fence, and loses its ledger entirely when the supply is removed.

Finite sources and uncertain delivery. If usable source energy R pays for the supplied transfer, internal replenishment cancels from the combined ledger, $R' = R - w^+$ and $(\beta + R)' = (\beta + R) - w^-$. In `Examples/FiniteSupply.lean` the construction `withSource` adds R to the store's reading and sets external supply to zero, so `totalDraw_le_initial_resources` bounds the drawn work by the initial combined reading and `horizon_le_of_finite_source` applies the pathwise horizon bound to it. Both horizon bounds restrict their cost premise to the finite run: a finite state space cannot sustain a uniformly positive net energy loss forever. The witness has source, buffer, load and selector bits with

reversible, energy-conserving gates, so conservation discharges the buffer's ledger and a unit reaches the load with probability 1/2; `forced_not_ledgered` rejects a ledger claiming delivery on every path, and `charger_no_finite_source` rules out any nonnegative finite source for the indefinitely replenished charger. Signed supply permits energy returned to the source; the transfer identity is a model obligation, and preparing the source, controlling the gates, external refuelling and identifying this hardware with the learning agent's actuator are outside the model.

Recording observations and clearing sensor memory. For a memory marginal μ and a strictly positive target law ν, the constant clearing channel has mean log-ratio heat

$$\langle Q_M \rangle = k_B T[H(\mu) - H(\nu)] + k_B T\, D_{KL}(\mu\|\nu), \tag{35}$$

identified with reservoir heat by local detailed balance, the KL term being the additional cost of that fixed channel on the supplied marginal. `Phase3_SensorMemory.lean` proves it as `memoryHeat_const`, with no positivity of μ: a state the memory never holds contributes to neither side. `MemoryAgent` runs act, record, learn and clear four-periodically, so `protocol` is an ordinary `FiniteProtocol` and the telescoped balances cover it unchanged, and the fences are again signatures — `record` and `erase` do not take the parameter, `update` takes neither parameter nor world, and the reward enters no channel. `memoryLaw_clear` is the one stage at which the memory is autonomous and `memoryLaw_record` the stage at which it is not, which is the content: the correlation recording writes is what clearing destroys. `clear_stageHeat_eq` splits the clear stage's heat into the register's, the memory's and the world's shares, and `clear_memoryHeat_const` is Eq. (35) for the agent's own run. On the bit-valued witness of `Examples/SensorMemory.lean` a marginal argument that uses no joint law prices every clearing between $(3/16)\log 3$ and $(5/16)\log 3$. Four controls fence the reading: a measurement ignoring the world leaves the same marginal and pays identical heat (`blind_same_cost`), so the cost is set by the marginal and not by what the memory is about; replacing erasure by symmetric drift costs nothing and carries the measurement forward; and `over_ordered_absorbs` delivers a negative share, a case the inequality does not see.

Funding repeated clearing from a carried source. A nonnegative source starting no higher than b, losing at least $c > 0$ at each clearing and receiving no return at other stages, can fund only m complete cycles satisfying

$$mc \le b, \tag{36}$$

the stage-dependent cost permitting free operations between clearings. `PathwiseStore.balance_le_of_stage_cost` generalizes the reading bound from a constant to a stage function, and `horizon_le_of_periodic_cost` is Eq. (36) for a run solvent at pm whose other stages return no more than they draw. `net_draw_eq_zero_of_positive` is the obstruction: on a `Positive` protocol the ledger holds between every pair of states in both directions, so the reading is constant and every net draw vanishes. In `Examples/FundedMemory.lean` the witness reuses the measurement, update, erasure and drift of the sensor agent, so what the source funds is what Eq. (35) prices; an induction on the reachable states proves the charge after n operations exactly, and `cycles_le_three` is Eq. (36) attained, with the reading below empty at the sixteenth operation. Two negatives are recorded rather than repaired. `clear_unfunded` proves that from the sixteenth operation every executed clearing draws nothing while the floor still prices it: `erase` reads only the memory, so no channel of this agent can be conditioned on the charge, and the run past three cycles is unfunded rather than halted. The drawn unit is a declared reading of the state and is not identified with the stage's log-ratio heat.

Preparing the initial law. A preparation drives the hardware to a strictly positive law without reading what the hardware held, which is the hypothesis of `memoryHeat_const`, so Eq. (35) prices it. `KL_uniform` gives $D_{KL}(\mu\|\mathrm{unif}) = \log|V| - H(\mu)$ with no positivity of μ, so the two terms cancel and `memoryHeat_uniform` is zero at every starting law and thermal scale. `FiniteProtocol.withPreparation` prepends the preparation as stage zero, so the telescoped balances and store bounds cover it, and `preparation_parameter_marginal` holds for every preparation channel: it writes the system coordinate and leaves the parameter's marginal as it found it, so the agent's uncertainty about the world is an input. In `Phase3_Preparation.lean` and its examples, preparing the sensor agent's uniform prior is free while preparing its standard state costs $(1/4)\log 3$; a bit law biased to ε costs $(1/2 - \varepsilon)\log[(1-\varepsilon)/\varepsilon]$, and `sharp_prep_unbounded` proves that family has no finite upper bound, its sharp endpoint rejected from the identity's positive-target premise rather than assigned an infinite

value. `known_parameter_not_preparable` and `correlated_not_preparable` bound what a state-blind preparation reaches: it cannot make a certain parameter uncertain, and it cannot produce a correlated two-bit prior.

Relating the drawn resource to physical energy. `Phase3_ResourceFoundations.lean` discharges three premises of the resource ledger on named objects rather than declaring them. The first is what the drawn unit *is*: `draw_eq_step_heat` identifies the bath coordinate's gain with the mean heat of the step the stage actually runs, which local detailed balance turns into $k_B T$ times that step's own bath entropy, and `draw_ge_entropy_reduction` is the consequence, under support inclusion rather than full positivity, so it reaches the sparse laws an agent runs on. Control remains on the right-hand side and is subtracted, so a large control budget weakens the bound rather than being constrained by it: fabrication is unpriced, and it is visible in the statement where it is unpriced. The second is the sharp charge. `transferCharge` moves one reversibly and creates nothing; what it cannot do is make one, and `sharp_preparation_not_reversibleSupport` says why — a channel driving every state to one target has no reverse path, so the path divergence is ∞, recorded in $[0, \infty]$ and not converted to a real. The approximate charge is priceable and the exact one is not, and both halves are theorems. The third is where delivery comes from. `SourceDelivery` identifies delivery with a source coordinate's loss on an arbitrary state type and `horizon_le_of_sourced_delivery` adds the source's initial content to the store's inside one horizon bound, so refuelling buys a longer run of the same shape; `no_finite_source_sustains` is the other half, a ledger with delivery from a nonnegative source, a store never below empty and a strictly positive cost at every operation being an impossible object. Sustained operation is not a property a finite agent has.

Trajectory identities and measure-valued feedback. The resource results above hold on an arbitrary state type, but they are trajectory identities: real readings indexed by operation number, with no probability law anywhere, so nothing in them discharges an integrability or existence assumption. Carrying the path law itself off finite spaces needs a different vehicle, and in `Phase3_MeasureFeedback.lean` it is Mathlib's divergence valued in $[0, \infty]$. `MeasureFeedbackStep` carries an initial probability measure and a Markov kernel on any measurable space, the forward path law being their composition and the entropy production the divergence from the declared reverse channel read in forward time order. Three properties the finite development imposes by hand then hold by construction: nonnegativity, because the codomain has no negative element; ∞ for a forward law not absolutely continuous with respect to the reversed one, with no positivity or finiteness hypothesis; and integrability as a stated side condition of the real-valued form rather than an implicit property of a finite sum. `path_divergence_splits` is the chain rule and the general replacement for Eq. (26); `entropyProduction_map_le` is data processing, so coarse-graining by any measurable observable cannot increase production; and `initial_divergence_le_entropyProduction` is the second law in divergence form. Control follows: `MeasureControlProblem.exists_optimal` selects a maximizing action over a finite permitted set with no finiteness on the state at all, and `exists_optimal_compact` over a compact one, with continuity discharged on witnesses that sit on $\mathbb{N}$ and on $\mathbb{R}$, the latter carrying a Gaussian initial law that assigns zero to every singleton. The *Shannon-entropy* form of the balance does not carry across: it needs a reference measure and densities against it, and differential entropy relative to Lebesgue measure is not entropy, not being even nonnegative. The divergence form is the statement that survives, and the entropy identity remains a finite-state result.

19 Thermodynamic bridges and coupling installation

The composition needs identifications beyond a heat inequality. A register ledger allocates heat among operations; its bath model tests what microscopic reversibility supplies. A scalar actuation law gives one spatial construction; the microscopic modes of Section 2 supply the product-space kernel whose mean is bounded by installed energy. The final subsection connects a learner to that installation on its own executed law. These connections require the stated identifications of reservoirs, stores and hardware.

19.1 Active and passive composition

Two thermodynamic statements enter, with different standing. The first is proved for the phase model: for a smooth positive periodic density with probability current J and diffusion D, the current functional

$\sigma_J = \int J^2/(D\rho)\, d\theta$ bounds how fast order can change, $\int_0^\tau \sigma_J\, dt \geq [\arcsin r(\tau) - \arcsin r(0)]^2/(D\tau)$ while $0 < r < 1$, and stationary order has no positive floor. The second is conditional: the installed-energy condition of Section 2.1 is necessary for its declared actuator, and restricts a physical candidate only once mode prices, coupling reduction and diffusion are measured. Neither is a heat or power statement until the current is identified with a physical flux, a reservoir at a stated temperature is specified and local detailed balance holds for the steps concerned; only then does σ_J convert to entropy production and the finite register budgets to Landauer heat (Landauer, 1961; Norton, 2005; Shenker, 2000). No composition edge derives coherence, compatibility or reconstruction from a heat bound. Finite-register capacity obeys $H(p) \leq \log|X|$, and a separate boundary result concerns a field reaching disconnected vacuum components (Kibble, 1976); Table S3 states the connecting assumptions. Section 1 gives the current-cost derivation; Sections 16–18 give the predictive-memory and feedback balances, learning and repeated-operation budgets, and the sensitivity of the metabolic comparison to its assumptions.

The passive and feedback composition theorems each accept eight connecting implications (Table S3). In the active branch, a named substep at the register's thermal scale satisfies

$$k_B T > 0, \qquad k_B T[H(p) - H(p')] \leq \langle Q \rangle \leq B, \tag{37}$$

with B the register's heat. The upper allocation follows from a ledger whose operations deliver that total heat to one reservoir and whose other operations do not increase joint entropy. Those ledger assumptions are physical inputs. Neither Landauer's lower bound nor microscopic reversibility supplies them. Finite two-bit witnesses exhibit positive forward and reverse elementary paths, local detailed balance at one thermal scale, and interaction energies of order $k_B T$; these conditions alone do not quantify proximity to equilibrium. Cortex has not been shown to satisfy these premises. The illustrative cortical metabolic expenditure over one region and one phase-diffusion time is 2.4×10^{14} times kT at body temperature, but identifying it with the register heat budget B requires a common physical ledger. Reading the register at the scale where its premises are physical, such as a single channel, requires E45Active to carry that step's budget to the convergence of a continuum kernel. The active branch is an alternative route to the same node rather than a stronger one, and no result of Section 2.1 passes through it: the installed-energy condition is derived from E56 and E67.

Table S3: Connecting hypotheses of the formal composition, eight per branch, and the installed-energy consequence. The branches share E12, E23 and E56–E89. Exact overlap compatibility and the encoding–readout structure are additional inputs. Table S1 maps the node-level results.

Connection	Status	Required assumption
E12 *capacity and boundary*	Physical premise	A finite-capacity system realizes a field leaving its vacuum manifold. Capacity and boundary are independent roots.
E23 *boundary to register*	Bridge assumption	The boundary is associated with a finite register whose update leaves states unreachable.
E34 / E34Active *register to dissipation*	Bridge assumption	Passive: predictive memory at the register's thermal scale and heat budget, with nonzero waste. Active: positive paths, local detailed balance and allocation of register heat to the named substep; the allocation uses the ledger identity and other operations' compressiveness. Both are exhibited at thermal-scale heats on finite state spaces with positive reverse paths, which is not the regime a cortical register occupies.
E45 / E45Active *couplings to continuum*	Modelling assumption	Convergent spatial coupling energies. Declared mode hardware supplies a kernel limit for either branch. A learning witness computes installation and work from its own executed law. Convergence follows from mesh refinement rather than from the budget, which selects no regime at macroscopic scale. Refinement and physical identifications remain inputs.
E56 *field identification*	Physical commitment	A specified physical field realizes the limiting coupling and satisfies the kernel conditions.
E67 *above threshold*	Physical commitment	That field operates above the synchronization threshold.

Table S3 (continued)

Connection	Status	Required assumption
E78 *coherence to cover*	Bridge assumption	Coherence relaxes the supplied cover at coupling at least the named mean-field value. Shared phase encoders bound vector-content mismatch; the exact cover and timely responsiveness remain separate requirements. Patches and oscillators carry one index, so one oscillator carries exactly one content patch: a modelling assumption rather than a consequence, and not an innocent one, since field-mediated coupling is spatial while two territories describing one quantity need not be neighbours. Where the cover's nerve is connected the same edge carries through the sites its patches share, at a residual linear in that nerve's diameter; where the configuration is frequency-locked it carries through the coupling graph's algebraic connectivity λ_2 instead of through any patch count, which asks the cover for nothing and is declared hardware data in the standing of κ.
E89 *contraction to fixed point*	Modelling assumption	A map mediated by local self-encoding fixes the glued section and has the stipulated Lipschitz factor $\Lambda < 1$. Reconstruction over a declared family imposes an additional code-separation criterion.
Installed energy	Conditional theorem	Under Section 2.1's occupancy and mode-price conditions, E56 and E67 require $U_{\mathrm{inst}} > 2D/\kappa$. The learning witness clears this necessary bound; its matched control does not. Neither physical premise follows.

Both branches also take a thermodynamic cover carrying exact overlap compatibility and a reflexive boundary specifying encoding and readout. They conclude that the glued section is the unique fixed point of self-prediction. Compatibility and accuracy at that section enter as premises; uniqueness follows. Identifying encoding with restriction needs the separate resonance condition of Eq. (7).

Each of the eight edges is supplied rather than derived from the one above it. E12 is equivalent to its conclusion, so capacity and boundary are independent roots rather than a step from one to the other; both E45 discharges take their convergence from refinement data and not from the thermodynamic bound; E78 states that coherence relaxes the supplied cover, and the coherence-to-content estimate of Eq. (5) bounds disagreement only for encoders already shared on overlaps. The inferential content of the composition is its last step, where contraction and a fixed glued section give uniqueness. What composing them establishes is that the premises are consistent, named in one place and machine-checked together, and not that coherence follows from installed energy or self-representation from coherence.

Joint witnesses combine a register, boundary, refining mesh and small spatial cover; active witnesses add controlled thermal updates and a priced installation. They establish satisfiability under their stated identifications. Mesh convergence follows from refinement data, and the witnesses establish no common cortical mechanism. Table S1 maps the claims to their Lean declarations and records the counterexamples. The development declares no additional axioms beyond its libraries' foundational axioms.

Formalization: `Chain.lean`.

The theorem `chain` uses the passive predictive node; `chain_active` uses Eq. (37) for a named `FiniteFeedbackStep` M, heat observable q, thermal scale k_BT and budget B. The process's controller has the register's state type. `E34Active` requires positivity, local detailed balance at that register's temperature and $\langle Q\rangle_M \le B$, with B the register's heat. The entropy lower bound is derived by `activeBoun d_of_e34Active`. A lower bound on erasure heat supplies no upper bound for an actuator and equates no operations; the upper-budget allocation comes from a resource model instead. `RegisterLedger` carries one register's update, its operations as elementary steps controlled by that register, local detailed balance for each at that register's own temperature, and the accounting identity $Q_{\mathrm{reg}} = \sum_i \langle Q\rangle_i$ for the heat its operations deliver to its reservoir. `opHeat_nonneg_of_compressive` applies the left inequality of Eq. (37) to an operation with $H(p') \le H(p)$, so its share is nonnegative; `opHeat_le_dissipation` then bounds an operation whose companions all compress by Q_{reg}, and `e34Active_of_ledger` turns that bound into the edge. Under Landauer's heat equation, `ledger_bathEntropy` exhibits the same total as the register's own bath entropy change. The identity and the compressiveness are the model's physical inputs. Equation (26) accounts for both finite update substeps at a common thermal scale; it supplies neither this allocation nor the coupling-convergence implication of E45Active, and the ledger's index labels operations rather than sequencing them into one evolving law.

`E45Active` consumes `ActiveBound` for these same process, heat observable, thermal scale and budget, and requires convergence of the specified coupling-energy sequence to its named limit. Neither an unspecified agent nor a bound for a different process discharges this edge. The scalar constitutive model below discharges this edge for its own process-dependent density, and the microscopic model below discharges it for a kernel on the product space built from mode occupancies the step actually changes. `E45` is discharged the same way for named hardware holding a declared `PredictiveDissipation`'s own joint law. Both proofs discard their premise, because the refinement data and not the bound supplies convergence. Both branches then use `chain_from_coarseGrains`, sharing E56–E89 and preserving the particular cover and its fixed section. Each composition has eight named edges; local-state compatibility, contraction and the fixed-state identification remain supplied physical content.

`chain_active_hypotheses_jointly_satisfiable` combines the noisy actuator with the one-bit register at thermal scale one, the changing grid and the specified field/cover. Its heat fits within the register's $\log 2$ budget, while its passive signed information difference is negative. `thermalAgency_zero_budget_rejected` excludes a zero budget for that same heat observable. `thermalAgency_wrong_limit_rejected` exhibits a valid active bound for which the constant-zero energy sequence cannot converge to the named limit three. The positive mesh witness supplies convergence independently; joint satisfiability establishes no common mechanism or thermodynamic derivation of convergence.

19.2 Allocating one register's heat budget

Formalization: `Examples/RegisterBudget.lean`.

The witness is that same one-bit register, at temperature one with Landauer heat $\log 2$, and its operations act on a lamp bit through one family: the environment relaxes toward agreement with the held register bit at odds $a : 1 - a$, from a law whose agreement probability is w, with interaction energy $\log[a/(1-a)]$ for disagreement and mean heat $\log[a/(1-a)]\,(a-w)$. The controlled operation takes $a = 2/3$, $w = 4/15$ and delivers $(2/5)\log 2$; the other takes $a = 4/5$, $w = 1/2$ and delivers $(3/5)\log 2$. Their sum is the register's $\log 2$, so `registerBudget_e34Active` discharges the bridge with no numerical comparison between separate models, and `chain_active_budget_jointly_satisfiable` composes the whole branch through it. Both inputs are load-bearing. `act_not_compressive` shows the controlled operation increases the joint entropy, so its own second law bounds it neither below at zero nor above; `compressive_exceeds_budget` exhibits an entropy-preserving operation of the same family delivering $(7/3)\log 2$, above the register's whole dissipation; and `leaky_exceeds_budget` exhibits a second ledger for the same register, satisfying positivity, local detailed balance and the identity, whose other operation draws heat out of the reservoir and whose controlled operation is allocated $(7/6)\log 2$. `leaky_other_not_compressive` identifies that operation as the one the derivation excludes.

19.3 Bath-energy transfer and reported heat

For a finite protocol with declared bath energy E_B and reported transition heat q_k, telescoping gives

$$\sum_{k<N} \langle \Delta E_B \rangle_k = \langle E_B \rangle_N - \langle E_B \rangle_0, \qquad \sum_{k<N} \langle q_k \rangle = \langle E_B \rangle_N - \langle E_B \rangle_0 + \sum_{k<N} \langle q_k - \Delta E_B \rangle, \tag{38}$$

The correction records any difference between the reported heat and actual bath-energy transfer.

Formalization: `Examples/RegisterBath.lean`.

`FiniteProtocol.energyTransfer` is the gain of the subsystem whose energy is declared, on an actual transition. `sum_energyTransfer` and `reported_heat_eq` (`Phase3_ContinuingAgent.lean`) give Eq. (38) for any finite protocol: the first is a first-law statement that admits deterministic gates and zero masses and asserts no entropy or temperature identification, and the second exhibits the correction by which a reported ledger differs from the modelled bath's gain.

The witness carries one register bit and one bath bit. `lift` sends each of the four maps of the register to a permutation of the pair (`lift_injective`) that performs the requested map on a bath prepared `false` (`lift_realizes`); an injective map acts on the register alone and an erasure swaps the register's bit into the bath. The prepared law is a uniform register and a pure bath, the bath's levels are 0 and $2\log 2$, the register's are degenerate (`bathEnergy_register_degenerate`), and the thermal scale is one. `operationHeat` is the bath's mean energy gain on the executed paths and `operationHeat_formula` evaluates it as $\log 2$ for the two erasures and 0 for the two injective maps. A `StatisticalMechanics Bool` instance discharges Landauer's heat equation for these values, with the reachable bath supports

as its final supports; the instance is local to the file, so §1's bath, which charges $\log 2$ for every map including the identity, and its consumers are unchanged.

Three results fence a map-indexed budget. *The eraser is not reusable.* `erase_law_one` clears the register and leaves the bath uniform, and `erase_law_two` proves that a second execution of the same gate restores the prepared law, so the register is uniform again. The shares are $\log 2$ (`erase_first_heat`) and $-\log 2$ (`reused_bath_heat`) about a total of zero (`two_stage_heat`), which is the generic bath ledger applied to this protocol rather than a separate calculation, so a share is neither nonnegative nor bounded by the total (`shares_exceed_total`).

The register's map does not fix the cost. `driveGate` holds the register at its value and excites its bath, so it performs the identity on a prepared bath at heat $2\log 2$ rather than 0 (`drivenIdleHeat_eq`, `same_update_different_heat`). The work is accounted with the bath: $\log 2$ for one erasure (`erase_work`), $2\log 2$ for the drive (`driven_work`) and zero for two erasures (`reuse_work`), the register's levels being degenerate.

Reversibility supplies neither ledger premise. A deterministic gate sends one state to one state, so its channel charges nothing to any other (`gate_channel_zero`) and the reduced step is not strictly positive (`gate_not_positive`); and a swap is an involution, so its forward and reverse transition masses agree on every realized path and its log-ratio heat vanishes at every thermal scale (`erase_stageHeat_zero`) while its bath gains $\log 2$ (`erase_stageHeat_not_bath_heat`), the whole of which the reported ledger then omits (`erase_reported_correction`).

This constructs no `RegisterLedger` instance and weakens none of its results: positivity, local detailed balance and the accounting identity remain the physical inputs `RegisterBudget` supplies, and this model shows that a reversible microscopic gate does not supply them. The values follow from the declared energy gap and preparation. Preparing the bath, supplying the work, fabricating the gate and identifying this bath with the reservoir the ledger's operations exchange heat with remain separate.

19.4 A scalar constitutive actuation law

For entropy reduction $d = H(p) - H(p')$, supplied gain $g \geq 0$ and spatial profile b, set $f_A(x) = g\,d\,b(x)$. Regular refinement gives

$$\mathcal{K}_n(f_A) \longrightarrow L_A = g\,d\int_S b\,d\mu, \qquad L_A \leq g\,\frac{B}{k_B T}\int_S b\,d\mu \tag{39}$$

The upper bound additionally needs $\int_S b\,d\mu \geq 0$ and the named process's budget in Eq. (37).

Formalization: `Phase3_ActuatedCoupling.lean`.

`ActuatedCoupling` carries a named feedback step, temperature and heat, a nonnegative conversion gain g, a spatial profile b, and regular meshes of a finite-measure region whose cell diameters tend to zero. Its signed drive $d = H(p) - H(p')$ is computed from the step's actual initial and final joint laws. The constitutive choice $f_A(x) = g\,d\,b(x)$ fixes the density consumed by `discreteEnergy`; it is not an independently supplied energy sequence. Uniform continuity and integrability of b give Eq. (39): `actuated_coarseGrains` applies mesh refinement to this density, and `e45Active_of_actuatedCoupling` discharges the corresponding edge. The budget plays no role in convergence. `actuated_limit_le_budget` uses that same process's `ActiveBound` to bound the limit, provided the total profile is nonnegative.

The noisy two-bit actuator has positive drive $d = (3/4)\log 3 - \log 2$. The witness in `Examples/ActuatedCoupling.lean` uses the nonconstant profile $b(x) = 1 + |x - 1/2|$ on $[0, 1)$. At unit gain the first two spatial energies are $3d/2$ and $5d/4$, so refinement changes the energy. The stationary version of the same channel has zero drive and actuates zero density. Doubling the gain strictly increases the limit at unchanged heat, and replacing the profile by one changes the spatial energy. Calibrating g to give limit three lets `chain_active_actuated_jointly_satisfiable` compose with the existing field and cover witnesses; a zero target is rejected for that same energy sequence. The example retains the thermal actuator's numerical budget comparison, separately from the register-ledger witness. All numerical identities here are exact Lean statements.

The constitutive gain and profile are physical inputs. The drive is an ensemble entropy difference, and its identification with a spatial amplitude is the supplied mechanism law. This model neither derives that law from local actions nor accounts for the work of installing it. The spatial mesh index is independent of learning time. The result is a scalar density and energy limit; the product-space kernel and cortical identification remain separate inputs.

19.5 Learning, installation and work on one executed law

`Examples/LearnedCoupling.lean`, §28, identifies the observational learner's pair (w, r) with the microscopic actuator's configuration. The mode is occupied when the two bits agree; its price is one, profile mass is two and conversion factor is $\kappa = 4$. The arrangement after n updates uses the learner's actual joint law, giving

$$U_n = \frac{3}{4} - \frac{1}{4}\,2^{-n}, \qquad K_n = 4U_n = 3 - 2^{-n}. \tag{40}$$

`installedEnergy_learned` and `continuumEnergy_learned` prove these exact identities. `kernel_update_on_executed_step` identifies the ensemble-kernel change with the occupancy change averaged over that same step's forward paths.

The first law prices the installation on those paths:

$$\langle W_n \rangle = U_{n+1} - U_n + \langle Q_n \rangle = \frac{1 + \log 3}{8}\,2^{-n}, \tag{41}$$

$$W_{<N} = \frac{1 + \log 3}{4}(1 - 2^{-N}) \le B, \qquad B = \frac{1 + \log 3}{4}. \tag{42}$$

Heat is the composite-channel heat of Eq. (32); installed energy adds the mode's storage change to the energetically degenerate register account. These are formal witness identities, not simulation estimates. Work is positive at every update, so identically free installation is rejected.

The uninformative sensor control of Section 17.4 uses the same hardware, prior, update rule and allowance. Its joint law stays uniform, $U_n = 1/2$ and $K_n = 2$; transition-log heat vanishes on each path and mean installation work is zero. At $D = 1$, the scalar field carrying its kernel has no coherent stationary branch at any horizon (`uninformed_no_coherence`). The learner also starts on the exclusion line, and its first update makes $\kappa U_n > 2D$ (`learned_above_line`). Substituting the control's law for the executed final law changes the kernel and is explicitly rejected.

The energy bound is necessary; exceeding it alone proves no coherence. The exact K_n above permits a separate stationary scalar comparison, but neither comparison establishes an evolving coherent field. The sensor phases are supplied configurations, and no dynamics feeds the installed kernel back into them. The allowance bounds expectations, with preparation, source implementation and cortical identification still additional specifications.

19.6 A finite phase–register loop with a work store

`Phase3_PhaseFeedback.lean` defines a joint transition retaining a phase p and register r: the installed action $a(r)$ selects a phase channel, a sensor reads the new phase, and the resulting observation updates the same register. Marginalizing the new register recovers the installed phase channel averaged over the current joint law. The update order retains the correlations that a product of separate phase and learner laws would discard. `Phase3_FeedbackLedger.lean` proves the finite-horizon recurrence and telescopes expected installation energy and a store account on these executed transitions. Its sensing, reset, preparation and mode prices are declared work inputs, and its exhaustion result assumes a positive expected net-cost floor without replenishment.

The numerical witness uses $p, r \in \{0, 1\}$, installed action $a = r$, signed coupling $K_a = h(2a - 1)$ to the phase spin $2p - 1$, mode price $u_a = a\Delta u$, and energy $E_a(p) = -K_a(2p - 1)$. Changing the register therefore changes the installed coupling and the next phase transition's rates. A phase flip occurs with probability $\gamma \exp[-\beta(E_a(1 - p) - E_a(p))/2]$, with the remaining probability on no flip. The chosen parameters keep every phase-transition probability positive. The flip-rate ratio is $P_a(p, 1 - p)/P_a(1 - p, p) = e^{-\beta[E_a(1-p) - E_a(p)]}$, so this phase substep has local detailed balance with the named bath at inverse temperature β. The sensor reports the new phase correctly with probability q and the register copies that report. The exact joint transition is the product of these conditional channels in that order. Deterministic copying need not give reverse support for the whole sensor/register operation, so local detailed balance is asserted for phase flips alone. No continuous-angle probability current exists in this binary phase space.

On each transition, installation work is the change in $E_a(p') + u_a$ at the new phase, while phase-flip heat to the bath is $E_a(p) - E_a(p')$. Their expected difference telescopes to the change in installed energy. Sensing and reset carry additional positive declared work prices; preparation is charged once. The work store subtracts these prices and adds a separately named external charger at each step. A conservative pathwise stop rule uses the maximum possible stage price, so it halts before any path can overdraw the

finite store; continuing operation requires the charger. No thermal lower bound is inferred for sensing, reset or charger channels without their own reverse protocols and reservoir model.

One fixed binary readout is fitted on independently sampled training episodes at a declared bath field and then frozen. Held-out episodes use two different field families and new random seeds. In one held-out family its next-phase accuracy is 0.783 with the informative sensor versus 0.552 with a blind sensor. A replay control follows the informed register marginal from an external schedule but removes its phase–register correlation, and an unchanging actuator is scored under the same initial independent prior and resource prices. The informed, blind and replay controls have the same marginal phase order in this symmetric example; their accuracy difference rests on the retained joint law, not on that order statistic. The benchmark saves binary phase/register symbol agreement, reconstruction error, installed energy, cumulative declared work, phase-flip bath heat and store readings separately. Symbol agreement is a toy shared-content proxy, without independent semantic decoders. The register, sensor, bath, charger and mode prices have no cortical calibration, so this closed-loop witness does not establish a cortical process.

20 Interpretive scope

The formal endpoint is the section glued on a specified compatible cover, fixed by a specified reflexive map. The interpretation developed in the main article identifies the section with conscious unity and its accurate self-representation with a minimal self. Uniqueness is relative to the cover and map; identifying these with meaningful neural descriptions and a neural self-encoding requires independent evidence. Restriction-resonance supplies the relation between the encoding and the state it represents.

The full commitment proposes that unity and reflexivity occur together in conscious episodes. Gluing by itself does not require reflexivity, so a dissociation could challenge that requirement while leaving the unity clause assessable. The physical recovery tests concern coupling and phase dynamics. Testing the interpretive clauses also requires validated measures of compatibility and self-encoding, together with an independent experience endpoint.

In the cup example, the proposed self-encoding concerns the perceiver's position and condition within the scene, without requiring an autobiographical narrative. Holding the external scene fixed while altering an independently identified encoding of the hand or viewpoint should change the reconstructed self-relative relation as the specified readout predicts. That manipulation tests the reconstruction model; whether the represented perspective is experienced remains a separate question.

21 Relationship to other theories

Integrated information and global-workspace theories supply different descriptions of integration and access (Tononi et al., 2016; Baars, 1997; Dehaene et al., 2014). The present model specifies a candidate physical coupling and a consistency condition on local states. Gluing does not derive an exclusion principle, and continuity alone does not force compatible local descriptions. A digital architecture is not excluded by the sheaf construction. These mathematical differences do not settle which systems have conscious experience.

Biological naturalism holds that the properties grounding experience belong to self-sustaining organisms and that simulating them does not secure them (Seth, 2026a,b). The conditions in this work name no biological mechanism. A digital candidate would have to exhibit physically realized coupling in a measured noise regime, agreement about independently identified shared quantities and internal reconstruction at a declared fidelity; satisfying those conditions would still leave the experiential interpretation open. A related position locates the difference in computation across inseparable scales and continuous dynamics of the physical medium (Milinkovic & Aru, 2026). A mathematical description of a field does not by itself realize such dynamics in a computing substrate.

The neuroscientific case against human-like consciousness in text-only language models concerns embodied sensory interaction, thalamocortical organization and biological regulation that links perception and action to survival (Aru et al., 2023). Their absence supports skepticism conditional on those mechanisms' causal relevance. None is required by abstract gluing or reconstruction. Reports of workspace-like organization in some models instead motivate a direct test of the proposed reconstruction condition (Gurnee et al., 2026): declare the state family and fidelity, fix the encoding region and its causal past, and test what the region distinguishes before interpreting an emitted report.

Two results bear on attempts to infer a generative mechanism from a steady state, including readings of the free-energy principle (Friston, 2010). Increasing coupling, decreasing diffusion and a matched

common drive share stationary phase densities, so the density alone cannot select their mechanism. Separately, local descriptions at positive overlap tolerance need not have an exact global extension; a decomposition alone supplies no single consistent state. These are explicit instances of identifiability concerns discussed in (Aguilera et al., 2022; Biehl et al., 2021; Bruineberg et al., 2022).

The comparisons have specific scope. The squared-drift functional is not a variational free energy, and the plasticity control concerns only its stated objective. At coherent stationarity, phase-space current dissipation vanishes while phase-averaged squared drift is positive; the substrate objective is a further functional on a different domain. These distinctions supply no general verdict about cortical inference or the free-energy principle.

22 Formal soundness: why a global axiom is the wrong shape for a physical postulate

Formalization is only as trustworthy as the axioms it rests on, and an axiom that is inconsistent silently makes every theorem downstream of it vacuous without changing whether the remaining statements type-check. The failure mode below is easy to reproduce: a declaration can constrain a symbol it does not itself bind, making the resulting class inconsistent.

The failure has one shape, and naming it is more useful than cataloguing its instances: *an axiom that constrains a symbol it does not itself bind.* Such a declaration is not a postulate about one system. It is a constraint on every model of the class the symbol belongs to, including models built to violate it, and a class that has such a model is thereby inconsistent. Three concrete cases, each of which proves `False` on its own:

1. `spontaneous_symmetry_breaking_pointwise_min` would take the potential-energy functional as an *implicit* argument occurring only in its hypothesis `PotentialEnergy phi = V v0`. Instantiating that functional with the constant map $v \mapsto V(v_0)$ discharges the hypothesis by `rfl` and yields the conclusion for an arbitrary field. With $V(v) = v^2$, $v_0 = 0$ and $\phi \equiv 1$ this proves $1 = 0$.

2. `landauer_heat_eq` would equate the free class field `heat_dissipation` with a bath-entropy difference *for every instance* of the class. An instance with zero heat dissipation and a bath that gains one bit gives $0 = \log 2$.

3. `kl_bound_axiom` would assert $\sigma \geq D_{KL}(P\|Q)/\Delta t$ for *all* P and Q, while σ is a single real fixed by a class field. Since D_{KL} is unbounded above, a point mass against a uniform distribution refutes it.

Two further declarations of the same shape are worth recording because their refutation is different in kind. `phase_invariant_periodic` and `sync_to_section_eq`, as global assertions pinning free fields of `LocalSectionSynchronization`, are refutable as soon as the probability presheaf admits two distinct global sections; the refutation needs no arithmetic at all, only that the sheaf is not trivial. Neither is a postulate in the current development: the periodicity condition is an obligation of the constructor that builds a cover from an invariant measure, and `sync_to_section_eq` is a theorem about the global section the sheaf condition produces.

The rule that prevents it is uniform: *a physical postulate that mentions a class field must be a field of that class, never a standalone axiom quantified over all instances.* A postulate is then an obligation each model discharges, not a global claim about every model. Under that rule the development declares no axioms at all.

22.1 What the witnesses establish, and what they do not

Explicit witnesses establish that the class assumptions can be satisfied. The `Examples` witnesses therefore cover every class that carries a physical postulate: a one-bit erasure system with a genuine bath that discharges `heat_eq` and reproduces $\Delta Q = k_B T \ln 2$; a perfectly resonant system discharging `kl_bound`; and a scalar field on a one-point spacetime discharging the action principles. The gluing structures are the most demanding case and are witnessed by a `ThermodynamicCover` over a three-site substrate carrying the discrete topology and its Borel σ-algebra, covered by two overlapping open patches whose intersection is a single shared site.

What the witnesses establish is that the classes are inhabitable, so the theorems about them are not vacuous. What they do not establish is that a cortex is such an instance. In the cover above the phase

field is already constant — every patch locked at $\theta = 0$ — which is precisely what makes the equilibrium obligation dischargeable. In the model, `ThermodynamicCover.ofConvergentTrajectory` reaches that configuration from initial data inside the quarter-turn region; whether a physical system reaches it is not shown here.

References

Pinotsis, D. A., & Miller, E. K. (2022). Beyond dimension reduction: Stable electric fields emerge from and allow representational drift. *NeuroImage*, 253, 119058.

Tononi, G., Boly, M., Massimini, M., & Koch, C. (2016). Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience*, 17(7), 450-461.

Singer, W., & Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience*, 18, 555-586.

Fries, P. (2015). Rhythms for cognition: communication through coherence. *Neuron*, 88(1), 220-235.

Shadlen, M. N., & Movshon, J. A. (1999). Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron*, 24(1), 67-77.

Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V. A. F. (2015). No-report paradigms: extracting the true neural correlates of consciousness. *Trends in Cognitive Sciences*, 19(12), 757-770.

Metzinger, T. (2003). *Being No One: The Self-Model Theory of Subjectivity*. Cambridge, MA: MIT Press.

Seth, A. K. (2013). Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences*, 17(11), 565-573.

Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., et al. (2023). Consciousness in artificial intelligence: insights from the science of consciousness. *arXiv preprint* arXiv:2308.08708.

Bobadilla-Suarez, S., Ahlheim, C., Mehrotra, A., Panos, A., & Love, B. C. (2020). Measures of neural similarity. *Computational Brain & Behavior*, 3, 369-383.

DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, 11(8), 333-341.

Kriegeskorte, N., & Diedrichsen, J. (2019). Peeling the onion of brain representations. *Annual Review of Neuroscience*, 42, 407-432.

Gallagher, S. (2000). Philosophical conceptions of the self: implications for cognitive science. *Trends in Cognitive Sciences*, 4(1), 14-21.

Baars, B. J. (1997). In the Theater of Consciousness: The Workspace of the Mind. *Oxford University Press*.

Dehaene, S., Charles, L., King, J. R., & Marti, S. (2014). Toward a computational theory of conscious processing. *Current Opinion in Neurobiology*, 25, 76-84.

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

Milinkovic, B., & Aru, J. (2026). On biological and artificial consciousness: A case for biological computationalism. *Neuroscience & Biobehavioral Reviews*, 181, 106524. doi:10.1016/j.neubiorev.2025.106524.

Aru, J., Larkum, M. E., & Shine, J. M. (2023). The feasibility of artificial consciousness through the lens of neuroscience. *Trends in Neurosciences*, 46(12), 1008–1017. doi:10.1016/j.tins.2023.09.009.

Seth, A. K. (2026). Conscious artificial intelligence and biological naturalism. *Behavioral and Brain Sciences*, 49, e315. doi:10.1017/S0140525X25000032.

Seth, A. K. (2026). The stuff matters: Consciousness, computation, and biology. *Behavioral and Brain Sciences*, 49, e366. doi:10.1017/S0140525X26106906.

Gurnee, W., et al. (2026). Verbalizable Representations Form a Global Workspace in Language Models. *Transformer Circuits Thread*. arXiv:2607.15495.

Aguilera, M., Millidge, B., Tschantz, A., & Buckley, C. L. (2022). How particular is the physics of the free energy principle?. *Physics of Life Reviews*, 40, 24-50.

Biehl, M., Pollock, F. A., & Kanai, R. (2021). A technical critique of some parts of the free energy principle. *Entropy*, 23(3), 293.

Bruineberg, J., Dolega, K., Dewhurst, J., & Baltieri, M. (2022). The Emperor's New Markov Blankets. *Behavioral and Brain Sciences*, 45, e183.

Pinotsis, D. A., & Miller, E. K. (2023). In vivo ephaptic coupling allows memory network formation. *Cerebral Cortex*, 33(17), 9877-9895.

Pinotsis, D. A., Fridman, G., & Miller, E. K. (2023). Cytoelectric coupling: Electric fields sculpt neural activity and "tune" the brain's infrastructure. *Progress in Neurobiology*, 226, 102465.

McFadden, J. (2002). The conscious electromagnetic information (cemi) field theory: the hard problem made easy?. *Journal of Consciousness Studies*, 9(8), 45-60.

Pockett, S. (2012). The electromagnetic field theory of consciousness: a testable hypothesis about the characteristics of conscious as opposed to non-conscious fields. *Journal of Consciousness Studies*, 19(11-12), 191-223.

McFadden, J. (2020). Integrating information in the brain's EM field: the cemi field theory of consciousness. *Neuroscience of Consciousness*, 2020(1), niaa016.

Fröhlich, F., & McCormick, D. A. (2010). Endogenous electric fields may guide neocortical network activity. *Neuron*, 67(1), 129-143.

Anastassiou, C. A., & Koch, C. (2015). Ephaptic coupling to endogenous electric field activity: why bother?. *Current Opinion in Neurobiology*, 31, 95-103.

Pockett, S. (2002). Difficulties with the electromagnetic field theory of consciousness. *Journal of Consciousness Studies*, 9(4), 51-56.

Vöröslakos, M., et al. (2018). Direct effects of transcranial electric stimulation on brain circuits in rats and humans. *Nature Communications*, 9, 483.

Verkhratsky, A., & Nedergaard, M. (2018). Physiology of Astroglia. *Physiological Reviews*, 98(1), 239-389.

Mountcastle, V. B. (1997). The columnar organization of the neocortex. *Brain*, 120(4), 701-722. doi:10.1093/brain/120.4.701.

Horton, J. C., & Adams, D. L. (2005). The cortical column: a structure without a function. *Philosophical Transactions of the Royal Society B*, 360(1456), 837-862. doi:10.1098/rstb.2005.1623.

Rakic, P. (2008). Confusing cortical columns. *Proceedings of the National Academy of Sciences*, 105(34), 12099-12100. doi:10.1073/pnas.0807271105.

Bastos, A. M., Donoghue, J. A., Brincat, S. L., Mahnke, M., Yanar, J., Correa, J., ... & Miller, E. K. (2021). Neural effects of propofol-induced unconsciousness and its reversal using thalamic stimulation. *eLife*, 10, e60824.

Bajwa, I. J., Nilsen, A. S., Skukies, R., Aamodt, A., Ernst, G., Storm, J. F., & Juel, B. E. (2025). A repeated awakening study exploring the capacity of complexity measures to capture dreaming during propofol sedation. *Scientific Reports*, 15, 32746. doi:10.1038/s41598-025-12695-z.

Bajwa, I. J., Nilsen, A. S., Skukies, R., Aamodt, A., Ernst, G., Storm, J. F., & Juel, B. E. (2024). A repeated awakening study exploring the capacity of complexity measures to capture dreaming during propofol sedation (Version 1.0.0) [Data set]. *OpenNeuro*. doi:10.18112/openneuro.ds005620.v1.0.0.

Xie, L., Kang, H., Xu, Q., Chen, M. J., Liao, Y., Thiyagarajan, M., ... & Nedergaard, M. (2013). Sleep drives metabolite clearance from the adult brain. *Science*, 342(6156), 373-377.

Syková, E., & Nicholson, C. (2008). Diffusion in brain extracellular space. *Physiological Reviews*, 88(4), 1277-1340.

Barbour, B. (2017). Analysis of claims that the brain extracellular impedance is high and non-resistive. *Biophysical Journal*, 113(7), 1636-1638.

Logothetis, N. K., Kayser, C., & Oeltermann, A. (2007). In vivo measurement of cortical impedance spectrum in monkeys: implications for signal propagation. *Neuron*, 55(5), 809–823. doi:10.1016/j.neuron.2007.07.027.

Gabriel, S., Lau, R. W., & Gabriel, C. (1996). The dielectric properties of biological tissues: II. Measurements in the frequency range 10 Hz to 20 GHz. *Physics in Medicine and Biology*, 41(11), 2251–2269. doi:10.1088/0031-9155/41/11/002.

Norton, J. D. (2005). Eaters of the lotus: Landauer's principle and the return of Maxwell's demon. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 36(2), 375-411.

Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5(3), 183-191.

Kibble, T. W. B. (1976). Topology of cosmic domains and strings. *Journal of Physics A: Mathematical and General*, 9(8), 1387-1398.

Sakaguchi, H. (1988). Cooperative phenomena in coupled oscillator systems under external fields. *Progress of Theoretical Physics*, 79(1), 39-46.

Strogatz, S. H., & Mirollo, R. E. (1991). Stability of incoherence in a population of coupled oscillators. *Journal of Statistical Physics*, 63, 613–635.

Acebrón, J. A., Bonilla, L. L., Pérez Vicente, C. J., Ritort, F., & Spigler, R. (2005). The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics*, 77, 137–185. doi:10.1103/RevModPhys.77.137.

Ott, E., & Antonsen, T. M. (2008). Low dimensional behavior of large systems of globally coupled oscillators. *Chaos*, 18, 037113. doi:10.1063/1.2930766.

Dörfler, F., & Bullo, F. (2014). Synchronization in complex networks of phase oscillators: A survey. *Automatica*, 50, 1539–1564. doi:10.1016/j.automatica.2014.04.012.

Breakspear, M. (2017). Dynamic models of large-scale brain activity. *Nature Neuroscience*, 20, 340–352. doi:10.1038/nn.4497.

Still, S., Sivak, D. A., Bell, A. J., & Crooks, G. E. (2012). Thermodynamics of prediction. *Physical Review Letters*, 109(12), 120604.

Kawai, R., Parrondo, J. M. R., & Van den Broeck, C. (2007). Dissipation: The phase-space perspective. *Physical Review Letters*, 98(8), 080602.

Horowitz, J. M., & Esposito, M. (2014). Thermodynamics with continuous information flow. *Physical Review X*, 4, 031015. doi:10.1103/PhysRevX.4.031015.

Ito, S., & Sagawa, T. (2013). Information thermodynamics on causal networks. *Physical Review Letters*, 111, 180603. doi:10.1103/PhysRevLett.111.180603.

Corberi, F., dello Russo, S., Messuti, G., Scarpetta, S., & Smaldone, L. (2026). Thermodynamic learning. *arXiv preprint*. arXiv:2609.04732.

Banach, S. (1922). Sur les opérations dans les ensembles abstraits et leur application aux équations intégrales. *Fundamenta Mathematicae*, 3(1), 133-181.

Shenker, O. R. (2000). Logic and entropy. *PhilSci-Archive preprint*. philsci-archive.pitt.edu/115.

Rosas, F. E., Mediano, P. A. M., Luppi, A. I., Boyd, A. B., Berta, M., Jensen, H. J., Gastpar, M., & Carroll, S. M. (2026). Four faces of information in natural systems. *Nature Reviews Physics*. doi:10.1038/s42254-026-00971-4.

Sznitman, A.-S. (1991). Topics in propagation of chaos. In *École d'Été de Probabilités de Saint-Flour XIX—1989*, Lecture Notes in Mathematics, vol. 1464, pp. 165–251. Springer. doi:10.1007/BFb0085169.

Dai Pra, P., & den Hollander, F. (1996). McKean–Vlasov limit for interacting random processes in random media. *Journal of Statistical Physics*, 84, 735–772. doi:10.1007/BF02179656.

Bertini, L., Giacomin, G., & Pakdaman, K. (2010). Dynamical aspects of mean field plane rotators and the Kuramoto model. *Journal of Statistical Physics*, 138, 270–290. doi:10.1007/s10955-009-9908-9.

Berglund, N., & Gentz, B. (2002). Pathwise description of dynamic pitchfork bifurcations with additive noise. *Probability Theory and Related Fields*, 122, 341–388. doi:10.1007/s004400100174.

Abramsky, S., & Brandenburger, A. (2011). The sheaf-theoretic structure of non-locality and contextuality. *New Journal of Physics*, 13, 113036. doi:10.1088/1367-2630/13/11/113036.

Townsend, R. G., Solomon, S. S., Chen, S. C., Pietersen, A. N. J., Martin, P. R., Solomon, S. G., & Gong, P. (2015). Emergence of complex wave patterns in primate cerebral cortex. *The Journal of Neuroscience*, 35(11), 4657-4662. doi:10.1523/JNEUROSCI.4509-14.2015.

Xu, Y., Long, X., Feng, J., & Gong, P. (2023). Interacting spiral wave patterns underlie complex brain dynamics and are related to cognitive processing. *Nature Human Behaviour*, 7(7), 1196-1215. doi:10.1038/s41562-023-01626-5.

Nicoletti, G., Di Terlizzi, I., & Busiello, D. M. (2026). Balancing information and dissipation with partially observed fluctuating signals. *Physical Review Letters*, 137(12), 127101. doi:10.1103/8gx2-rbns.

Kiyooka, D., Oomoto, I., Kitazono, J., Saito, Y., Kobayashi, M., Matsubara, C., Kobayashi, K., Murayama, M., & Oizumi, M. (2026). Single-cell resolution functional networks during unconsciousness are segregated into spatially intermixed modules. *Cell Reports*, 45(2), 116902. doi:10.1016/j.celrep.2025.116902.

Oomoto, I., Kiyooka, D., Oizumi, M., & Murayama, M. (2026). Multi-area single-cell calcium imaging dataset of the mouse cortex across wakefulness, sleep, and anesthesia. *Scientific Data*. doi:10.1038/s41597-026-08202-2.

Attwell, D., & Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism*, 21(10), 1133-1145. doi:10.1097/00004647-200110000-00001.

Harris, J. J., Jolivet, R., & Attwell, D. (2012). Synaptic energy use and supply. *Neuron*, 75(5), 762-777. doi:10.1016/j.neuron.2012.08.019.

Nicholls, D. G., & Ferguson, S. J. (2013). *Bioenergetics 4*. Academic Press.

Baer, S. M., Erneux, T., & Rinzel, J. (1989). The slow passage through a Hopf bifurcation: delay, memory effects, and resonance. *SIAM Journal on Applied Mathematics*, 49(1), 55–71. doi:10.1137/0149003.

Friedman, E. B., Sun, Y., Moore, J. T., Hung, H.-T., Meng, Q. C., Perera, P., Joiner, W. J., Thomas, S. A., Eckenhoff, R. G., Sehgal, A., & Kelz, M. B. (2010). A conserved behavioral state barrier impedes transitions between anesthetic-induced unconsciousness and wakefulness: evidence for neural inertia. *PLoS ONE*, 5(7), e11903. doi:10.1371/journal.pone.0011903.

Hudson, A. E., Calderon, D. P., Pfaff, D. W., & Proekt, A. (2014). Recovery of consciousness is mediated by a network of discrete metastable activity states. *Proceedings of the National Academy of Sciences*, 111(25), 9283–9288. doi:10.1073/pnas.1408296111.

Proekt, A., & Hudson, A. E. (2018). A stochastic basis for neural inertia in emergence from general anaesthesia. *British Journal of Anaesthesia*, 121(1), 86–94. doi:10.1016/j.bja.2018.02.035.

Bobadilla-Suarez, S. (2026). Unity requires physically enforced agreement. *PsyArXiv*. doi:10.31234/osf.io/dkqep_v1.