

Coherence-Tilted Trajectory Ensembles for Adaptive Systems with Co-Evolving Constraints

A Path-Statistical Framework for Non-Equilibrium Complex Systems

[Author Name]

[Affiliation]

June 15, 2026

Abstract

Many adaptive non-equilibrium systems—glassy materials, evolving biological populations, learning agents, developing organisms—are poorly captured by probability distributions over *states* alone, because the very landscape that governs their dynamics is itself reshaped by their motion through it. A spin glass whose coupling strengths are reinforced by repeated spin correlations, a neural circuit whose synaptic weights are modified by the firing patterns it produces, or a society whose norms are strengthened or eroded by the behaviour they regulate are all instances of the same underlying phenomenon: the trajectory rewrites the rules that govern future trajectories.

We develop a path-statistical framework—the Entropic Coherence Framework (ECF)—in which the primitive object is a probability measure over *trajectories* of an extended state $z_t = (x_t, m_t, c_t)$ comprising a visible state x_t , a memory variable m_t that integrates past behaviour, and—crucially—a co-evolving constraint field c_t that feeds back into future dynamics. We call the central observable “coherence,” defined not as order or low energy but as *path-aligned constraint formation that preserves future reach*.

The framework makes four structural contributions: (i) We *derive* a minimal coherence path functional \mathcal{K} from a constrained variational principle rather than positing it. (ii) We introduce an *alignment-gated* refinement $\mathcal{K}_{\text{aligned}}$ that restricts credit to time steps at which the constraint field is being rewritten in a manner consistent with the actual trajectory. (iii) We locate the tilted ensemble $\mathbb{P}_\eta \propto \mathbb{P}_0 e^{\eta\mathcal{K}}$ within large-deviation theory and prove a large-deviation principle via a generalised Doob transform. (iv) We separate an intrinsic (descriptive) reading from a control (prescriptive) reading.

We report results from a computational study of the adaptive Edwards–Anderson spin glass. The primary result is that the alignment-gated susceptibility $\chi_{\mathcal{K}_{\text{aligned}}}$ separates coupled adaptive co-evolution from time-shuffled controls by one to two orders of magnitude, providing an operational demonstration that the alignment condition is the decisive ingredient. We state four falsifiable predictions, honestly ranked by evidential support, and we demote the original “curiosity window in coherence pressure” to a secondary hypothesis following negative numerical results, replacing it with a reformulated prediction in the physical control plane. The framework complements—and is explicitly compared with—equilibrium statistical mechanics, stochastic thermodynamics, the s -ensemble, the free-energy principle, reinforcement learning, and empowerment-based approaches.

Contents

1	Introduction	3
2	Extended state space and reference dynamics	4
2.1	Why a larger state space is necessary	4

3	Path functionals: what we measure along trajectories	5
3.1	Frustration: how badly the current state violates constraints	5
3.2	Dissipation: the thermodynamic cost of a trajectory	5
3.3	Reach: how many viable futures remain accessible	6
3.4	Memory alignment: does the system act in accordance with its history?	6
4	A variational derivation of the coherence functional	7
4.1	The problem: deriving rather than positing	7
4.2	From local derivation to global path functional	8
4.3	The alignment gate: why plasticity alone is insufficient	8
5	The coherence-tilted ensemble	9
6	Mechanism: the generalised Doob transform	9
7	Large deviations of coherence	10
8	Two readings: descriptive versus prescriptive	11
9	Order parameters and the three regimes	11
9.1	Primary order parameters	11
9.2	Extended thermodynamic structure	11
9.3	Constraint-memory mutual information	12
9.4	Constraint sufficiency ratio	12
9.5	Hierarchy and the three regimes	12
10	Instantiation: the adaptive Edwards–Anderson spin glass	13
10.1	Model setup	13
10.2	Analytic lock-in criterion	14
10.3	Alignment in the spin-glass model	14
11	Comparison with neighbouring frameworks	14
11.1	Equilibrium statistical mechanics	15
11.2	Stochastic thermodynamics	15
11.3	The dynamical s -ensemble	16
11.4	The free-energy principle and active inference	16
11.5	Reinforcement learning	17
11.6	Empowerment and causal entropic forces	17
12	Falsifiable predictions	18
13	Limitations and scope	20
14	Conclusion	20
A	Computational experiments: protocols and outcomes	21

1 Introduction

Many non-equilibrium systems evolve on landscapes that are themselves modified by the system’s history. Examples include adaptive spin glasses with evolving couplings, neural networks with activity-dependent plasticity, evolutionary systems whose fitness landscape changes through ecological feedback, and learning agents that modify their own internal constraints. In such systems the distinction between state dynamics and landscape dynamics becomes blurred: the trajectory influences the constraints that govern future trajectories.

Most theoretical frameworks in statistical physics treat the dynamical generator as fixed or externally prescribed. Equilibrium statistical mechanics describes probability distributions over states generated by a static Hamiltonian (Landau and Lifshitz, 1980). Stochastic thermodynamics extends the analysis to trajectories and provides a general framework for entropy production and fluctuation relations (Seifert, 2012; Jarzynski, 2011), but the underlying transition structure remains externally specified. Similarly, the dynamical s -ensemble and related large-deviation approaches characterize rare fluctuations of trajectory observables under fixed kinetic rules (Garrahan et al., 2007, 2009; Lecomte et al., 2007). These approaches provide powerful descriptions of trajectory statistics but do not explicitly address systems in which the dynamical constraints co-evolve with the trajectory itself.

Several neighbouring frameworks capture related aspects of adaptive behaviour. The free-energy principle and active inference describe agents that act to maintain preferred future states through variational inference (Friston, 2010; Friston et al., 2019; Parr et al., 2022). Reinforcement learning optimizes expected return over trajectories in a Markov decision process (Sutton and Barto, 2018), while empowerment-based approaches quantify the diversity of accessible futures (Klyubin et al., 2005; Wissner-Gross and Freer, 2013). However, these frameworks generally assume either fixed transition dynamics or slowly varying models and do not provide a trajectory-level statistical mechanics for systems whose constraint structure is continuously rewritten by their own dynamics.

The present work addresses this problem by introducing an extended-state path ensemble in which both memory and constraints are treated as dynamical variables. The state of the system is represented by

$$z_t = (x_t, m_t, c_t),$$

where x_t denotes the observable configuration, m_t a memory variable encoding aspects of past behaviour, and c_t a constraint field that modifies future dynamics while itself evolving through interaction with the trajectory. This construction restores Markovianity in the extended space while allowing the visible dynamics to remain history dependent.

Within this setting we define a trajectory observable termed *coherence*. Coherence is not intended as a measure of order, low energy, or novelty. Instead, it quantifies the extent to which changes in the constraint field are aligned with the trajectory while preserving future dynamical possibilities. Operationally, coherence combines future reach, frustration, and dissipation with an alignment measure between trajectory increments and constraint updates. The resulting observable generates a family of coherence-tilted trajectory ensembles analogous to the s -ensemble of non-equilibrium statistical mechanics.

The main contributions of this work are fourfold. First, we formulate an extended-state framework for systems with co-evolving constraints. Second, we derive a minimal coherence functional from a constrained variational principle rather than introducing it phenomenologically. Third, we introduce an alignment-gated refinement that distinguishes genuine trajectory-induced constraint formation from externally imposed or uncorrelated changes in the constraint field. Fourth, we establish the large-deviation structure of the resulting trajectory ensemble through a generalized Doob transform.

To demonstrate the framework we study an adaptive Edwards–Anderson spin glass in which the couplings evolve according to past spin correlations. Numerical experiments reveal a clear

separation between genuinely coupled dynamics and time-shuffled controls when measured through the alignment-gated coherence susceptibility. This provides evidence that trajectory–constraint alignment, rather than constraint plasticity alone, is the key ingredient underlying the proposed notion of adaptive coherence.

The paper is organised to layer mathematical precision onto conceptual foundations. A reader interested primarily in the central ideas can read Sections 1–10 and the prediction section (Section 12). The proofs (Sections 6 and 7) are self-contained and can be read independently.

Section 2 fixes the extended state space. Section 3 defines the path functionals, explaining each in physical terms. Section 4 derives the minimal coherence functional variationally and introduces the alignment-gated refinement. Section 5 defines the coherence-tilted ensemble and its generating function. Section 6 gives the Doob-transform mechanism. Section 7 proves the large-deviation principle. Section 8 separates the descriptive and prescriptive readings. Section 9 defines the order parameters and the three regimes. Section 10 works out the spin-glass instantiation in full. Section 11 gives the detailed comparison with neighbouring frameworks. Section 12 states the four falsifiable predictions with honest status reports. Section 13 states limitations. Appendix A describes the computational experiments and their outcomes.

2 Extended state space and reference dynamics

2.1 Why a larger state space is necessary

The first move is to enlarge the state space. In a standard Markov chain model, the future of the system is fully determined—in a statistical sense—by its current state x_t . In an adaptive system with memory and co-evolving constraints, this is false: two systems in the same visible state x_T but with different histories of how they arrived there can have completely different futures, because the history is encoded in the constraint field and memory.

To restore the Markov property—and thereby gain access to the entire toolkit of Markov chain theory—we must extend the state space to include memory and constraint field explicitly.

Definition 2.1 (Extended state). The system is described at each time $t \in \{0, 1, \dots, T\}$ by an *extended state*

$$z_t = (x_t, m_t, c_t) \in Z = X \times M \times C,$$

where X is the observable state space, M is the memory space (a compact subset of \mathbb{R}^{d_M}), and C is the space of constraint fields (a compact subset of \mathbb{R}^{d_C}).

Intuition

The spin-glass picture. The visible state x_t is the spin configuration (s_1, \dots, s_N) , the memory m_t is an exponential trace of past magnetisation patterns, and the constraint field c_t is the full coupling matrix $J_{ij}(t)$. Two spin configurations that look identical but have different coupling matrices will have different future dynamics—the extended state captures this distinction, the visible state alone does not.

The constraint field: what it encodes. The constraint field c_t is the central object of the framework. It encodes which transitions are cheap, which are costly, and which are effectively forbidden. In a spin glass it is the coupling matrix; in a neural circuit it is the synaptic weight matrix; in an economic model it might be the transaction costs and regulations. What is special here is that c_t is not fixed—it is itself a dynamical variable updated by the trajectory.

Definition 2.2 (Reference dynamics). A *reference dynamics* is a Markov kernel on Z ,

$$\mathbb{P}_0(z_{t+1} | z_t) = p_0(x_{t+1} | x_t, m_t, c_t) \delta(m_{t+1} - \mathcal{M}(m_t, x_t, x_{t+1}, c_t)) \delta(c_{t+1} - \mathcal{C}(c_t, x_t, x_{t+1}, m_t)),$$

with initial law ρ_0 . The reference probability of a trajectory is

$$\mathbb{P}_0(\Gamma_{0:T}) = \rho_0(z_0) \prod_{t=0}^{T-1} \mathbb{P}_0(z_{t+1} | z_t).$$

Reading the definition. The transition factorises into three parts. The visible dynamics $p_0(x_{t+1} | x_t, m_t, c_t)$ is the baseline physical rule—e.g. Glauber or Metropolis dynamics in a spin glass. The memory update \mathcal{M} is a running average or trace of past behaviour. The constraint update \mathcal{C} is the key novel element: it makes the constraint field a function of the current visible trajectory, so that the landscape is literally edited by the path.

Remark 2.3 (Markov in z , non-Markov in x). The dynamics is Markov in the extended state but generically *non-Markov* in the visible state x_t alone. Observing only the spin configuration and not the coupling history makes the future unpredictable from the present. This is the Mori–Zwanzig memory phenomenon (Zwanzig, 2001; Mori, 1965), here given a concrete structural interpretation: the non-Markovian memory is carried in (m_t, c_t) .

3 Path functionals: what we measure along trajectories

The coherence functional will be built from four primitive path quantities. Each is *additive along the path*: $\Phi(\Gamma) = \sum_t \phi(z_t, z_{t+1})$. This additivity is essential—it is what allows large-deviation theory to be applied.

3.1 Frustration: how badly the current state violates constraints

Definition 3.1 (Frustration). With local constraint-satisfaction functions $s_\alpha(x, c) \in [0, 1]$ (where $s_\alpha = 1$ means satisfied), the *instantaneous frustration* is

$$\Omega(x, c) = \sum_{\alpha \in \mathcal{H}} w_\alpha (1 - s_\alpha(x, c)), \quad w_\alpha \geq 0,$$

and the *path frustration* is $\Omega(\Gamma) = \sum_{t=0}^T \Omega(x_t, c_t)$.

Example

Spin glass. For Ising couplings, $s_{ij} = \frac{1}{2}(1 + J_{ij}s_i s_j)$ equals 1 when spin i and spin j are aligned as the coupling prefers, and 0 when they are anti-aligned. Ω is then the fraction of unsatisfied bonds—the standard frustration measure in spin glass physics (Mézard et al., 1987; Binder and Young, 1986). When the couplings evolve, frustration can either decrease (the system is learning to satisfy its own constraints) or increase (runaway reinforcement builds incompatible demands).

3.2 Dissipation: the thermodynamic cost of a trajectory

Definition 3.2 (Path dissipation). The *dissipation* is the log-ratio of forward to time-reversed path probability,

$$\Sigma(\Gamma_{0:T}) = \log \frac{\mathbb{P}_0(\Gamma_{0:T})}{\mathbb{P}_0^\dagger(\Theta\Gamma_{0:T})},$$

where Θ reverses time and \mathbb{P}_0^\dagger is the time-reversed reference process. A local cost proxy $D(\Gamma) = \sum_t d(z_t, z_{t+1})$ is used when convenient.

This is the standard stochastic-thermodynamic entropy production (Seifert, 2012), meaning the dissipation term in our functional inherits the full suite of fluctuation relations (Jarzynski, 2011; Crooks, 1999). We do not postulate an arbitrary energy penalty; we use the physically grounded cost that quantifies the thermodynamic irreversibility of each transition.

3.3 Reach: how many viable futures remain accessible

Reach is the central novel concept. Informally, the reach of a state is a measure of how many high-viability, low-frustration continuations remain open. A system with high reach has many good futures to choose from; a locked-in system has very few. The concept is related to empowerment (Klyubin et al., 2005) and causal entropic forces (Wissner-Gross and Freer, 2013), but differs in two key ways: it accounts for the current constraint field (so the set of accessible futures changes as the landscape evolves), and it weights futures by viability rather than merely counting them uniformly.

Two objections must be addressed before giving the formal definition.

Objection 1: Reach may be circular. If viability is defined through coherence, and coherence through reach, we have a loop. We break it by *defining viability exogenously*.

Definition 3.3 (Viability, defined independently of coherence). Fix a *viability function* $V : X \rightarrow \mathbb{R}$ specified entirely by the application domain *without reference to \mathcal{K}* : negative energy density in a spin glass, reproductive fitness in an evolutionary model, or cumulative task return in a control problem. Viability is an input to the theory, not an output.

Objection 2: Reach is intractable. Exactly enumerating all reachable states within τ steps is exponential in τ . We use controlled surrogates.

Definition 3.4 (τ -reach). Let $\mathcal{A}_\tau(z_t) \subseteq X$ be the set of visible states reachable from z_t in at most τ steps under the *current* constraint field c_t held fixed. Define

$$R_\tau(z_t) = \log \sum_{x \in \mathcal{A}_\tau(z_t)} \exp[\nu V(x) - \mu \Omega(x, c_t)].$$

This is a free-energy-like log-volume weighting accessible states by both viability and low-frustration.

Remark 3.5 (Tractable surrogates and a key failure mode). Three progressively more expensive approximations are used in practice: (i) a first-order surrogate $\Delta R \approx -\Delta \Omega$ valid for single-flip updates, which ties reach gain to frustration reduction; (ii) Monte-Carlo rollouts of length τ ; and (iii) learned value-function surrogates in the control reading.

Important warning. The first-order surrogate makes the per-step coherence integrand proportional to $-\Delta \Omega$ minus a dissipation term, so that any “coherence transition” would be indistinguishable from an ordinary heat-capacity or frustration feature. This is why the alignment gate of Section 4 is not merely a refinement—it is the ingredient that prevents the framework from being a relabelling of energetics. We always use the gated functional and verify against a shuffle control.

3.4 Memory alignment: does the system act in accordance with its history?

Definition 3.6 (Memory alignment). Memory updates by an exponential trace $m_{t+1} = (1 - \lambda)m_t + \lambda \psi(x_t, x_{t+1})$ with feature map ψ . The *memory alignment* of a transition is $A_M(z_t, z_{t+1}) = \langle m_t, \psi(x_t, x_{t+1}) \rangle$, measuring how much the current move agrees with what the system has typically done before.

Definition 3.7 (Information curiosity). The *information curiosity* at time t is

$$C_t = I\left(R_\tau^{\text{future}}; a_t \mid \Gamma_{\text{past}}\right) = H(R_\tau \mid z_t) - \mathbb{E}_{a_t}[H(R_\tau \mid z_t, a_t)],$$

where a_t denotes the current action. This is the reduction in uncertainty about the future-reach structure produced by the current action, conditional on all past observations.

Remark 3.8 (Why this replaces the old novelty term). An earlier version used a quadratic novelty term $gN_t - rN_t^2$, producing an interior maximum at $N^* = g/2r$. This is an ad hoc formula with two free parameters and no first-principles derivation. Information curiosity C_t has a single multiplier λ_C and a transparent meaning: it is high when the current action is *actionably informative* (it teaches the system which viable futures are actually accessible) and low when the action is either purely routine or purely noisy. The quadratic’s interior maximum is recovered as a special case when C_t varies linearly with N_t , but the information-theoretic form generalises naturally and connects to the thermodynamics of prediction (Still et al., 2012; Bialek et al., 2001).

4 A variational derivation of the coherence functional

4.1 The problem: deriving rather than positing

A standard objection to coherence proposals is that the functional \mathcal{K} looks like a reward function with free parameters chosen to produce desired outputs. We address this by *deriving* \mathcal{K} from a natural variational principle: among all one-step laws consistent with the reference dynamics, find the one that maximally improves expected future reach at minimal dissipative and frustration cost.

Definition 4.1 (Local reach-maximisation programme). Among one-step laws $q(\cdot \mid z_t)$ that are absolutely continuous with respect to $p_0(\cdot \mid z_t)$,

$$\max_q \mathbb{E}_q[R_\tau(z_{t+1})] \quad \text{subject to} \quad \mathbb{E}_q[\Omega(z_{t+1})] \leq \bar{\Omega}, \quad \mathbb{E}_q[d(z_t, z_{t+1})] \leq \bar{D}, \quad D_{\text{KL}}(q \parallel p_0(\cdot \mid z_t)) \leq \epsilon. \quad (1)$$

The constraints say: do not let frustration exceed $\bar{\Omega}$, do not let dissipation exceed \bar{D} , and do not deviate too far from the reference dynamics (which encodes domain knowledge, physical laws, or a prior).

Proposition 4.2 (Gibbs form of the optimal step). *The solution of (1) is*

$$q^*(x_{t+1} \mid z_t) \propto p_0(x_{t+1} \mid z_t) \exp[\eta(R_\tau(z_{t+1}) - e\Omega(z_{t+1}) - dD(z_t, z_{t+1}))],$$

where $\eta > 0$ is the Lagrange multiplier for the KL constraint and $e, d \geq 0$ are the multipliers for frustration and dissipation.

Proof. The objective is linear in q and the KL constraint is strictly convex, so the problem is strictly concave in q with a unique maximiser. Stationarity of the Lagrangian gives $\log q/p_0 = \eta(R_\tau - e\Omega - dD) + \text{const}$, i.e. the stated Gibbs form. \square

Intuition

What the derivation says. The optimal one-step distribution is a Boltzmann-like tilt of the reference dynamics, with “energy” equal to $-(R_\tau - e\Omega - dD)$. Moving toward high-reach, low-frustration, low-dissipation states is optimal under the principle of maximum expected reach. This is not a postulate; it is a theorem about the solution to a natural variational

problem. The multipliers e and d are not free parameters; they are Lagrange multipliers of the physical constraints and take values determined by $(\bar{\Omega}, \bar{D})$.

4.2 From local derivation to global path functional

Iterating the local optimum across time steps recovers a path-tilted measure. This motivates defining:

Definition 4.3 (Minimal coherence functional).

$$\mathcal{K}_{\min}(\Gamma) = \sum_{t=0}^{T-1} \left[R_{\tau}(z_{t+1}) - e\Omega(z_{t+1}) - dD(z_t, z_{t+1}) \right]. \quad (2)$$

4.3 The alignment gate: why plasticity alone is insufficient

A critical finding from numerical experiments is that \mathcal{K}_{\min} is too permissive. Because $\Delta R \approx -\Delta\Omega$ for single-move surrogates, any system with fluctuating frustration produces large fluctuations in \mathcal{K}_{\min} , and the susceptibility $\chi_{\mathcal{K}}$ responds even when the constraint field is changing in a completely trajectory-*uncorrelated* way. Time-shuffling the constraint updates (breaking their correlation with the trajectory while preserving all marginal distributions) produced peaks in $\chi_{\mathcal{K}}$ of similar height to the genuine coupled system. This means \mathcal{K}_{\min} does not distinguish “the system is sculpting its landscape to support its own trajectory” from “the landscape happens to be changing at the same time as the trajectory.”

The fix is to gate each coherence increment by a measure of alignment between the current trajectory step and the direction of constraint change.

Definition 4.4 (Trajectory-constraint alignment). Let $\Delta c_t = c_{t+1} - c_t$ be the constraint deformation and $\psi(x_t, x_{t+1})$ the same feature map used by memory. The alignment is

$$A_C(z_t, z_{t+1}) = \frac{\langle \Delta c_t, \Phi(\psi(x_t, x_{t+1})) \rangle}{\|\Delta c_t\| \|\Phi(\psi(x_t, x_{t+1}))\| + \varepsilon} \in [-1, 1],$$

where Φ maps the move feature into constraint space. $A_C = 1$ means the constraint is being rewritten in exactly the direction the trajectory is pointing; $A_C = -1$ means it is being rewritten in the opposite direction; $A_C = 0$ means the two are orthogonal (constraint updating is trajectory-independent).

Definition 4.5 (Alignment-gated coherence functional). Let $G : [-1, 1] \rightarrow [0, 1]$ be a monotone sigmoid gate, $G(A_C) = \sigma(\gamma(A_C - \theta))$ with $\sigma(u) = (1 + e^{-u})^{-1}$, sharpness $\gamma > 0$, and threshold θ . The *alignment-gated coherence functional* is

$$\mathcal{K}_{\text{aligned}}(\Gamma) = \sum_{t=0}^{T-1} G(A_C(z_t, z_{t+1})) \left[\Delta R_{\tau}(z_{t+1}) - e\Omega(z_{t+1}) - d|\Delta E_t| \right]. \quad (3)$$

Only time steps at which the constraint field is being rewritten *consistently with the trajectory* contribute to the coherence tally.

Remark 4.6 (The shuffle test as operational definition). The decisive operational test for whether $\mathcal{K}_{\text{aligned}}$ captures genuine path-consistent constraint co-evolution is: compute $\mathcal{K}_{\text{aligned}}$ for the actual system, then recompute it with the alignment sequence $\{A_C(t)\}$ randomly permuted in time. This shuffle destroys the temporal correlation between alignment and coherence increment while preserving all marginal distributions. If $\mathcal{K}_{\text{aligned}}$ were merely tracking frustration, the coupled and shuffled values would be equal. In practice, the coupled susceptibility exceeds the shuffled susceptibility by one to two orders of magnitude (Appendix A). This gap is the operational definition of “coherence beyond frustration.”

Definition 4.7 (Full coherence functional). Adding memory alignment and information curiosity:

$$\mathcal{K}(\Gamma) = \mathcal{K}_{\text{aligned}}(\Gamma) + \sum_{t=0}^{T-1} \left[b A_M(z_t, z_{t+1}) + \lambda_C C_t \right], \quad (4)$$

with a single curiosity multiplier $\lambda_C \geq 0$. This replaces the earlier two-parameter quadratic novelty term.

5 The coherence-tilted ensemble

Given a reference measure \mathbb{P}_0 and the path functional \mathcal{K} , we define a one-parameter family of biased trajectory distributions.

Definition 5.1 (Coherence-tilted ensemble).

$$\mathbb{P}_\eta(\Gamma) = \frac{\mathbb{P}_0(\Gamma) e^{\eta \mathcal{K}(\Gamma)}}{\mathcal{Z}_T(\eta)}, \quad \mathcal{Z}_T(\eta) = \mathbb{E}_{\mathbb{P}_0}[e^{\eta \mathcal{K}(\Gamma)}]. \quad (5)$$

Remark 5.2 (Connection to the s -ensemble). Equation (5) is the standard exponential tilt of a path ensemble (Garrahan et al., 2007; Lecomte et al., 2007; Chetrite and Touchette, 2015), with \mathcal{K} playing the role of the time-extensive observable and $-\eta$ the counting field s . We claim no novelty for the tilt itself. The content lies entirely in the *choice of \mathcal{K}* (variationally derived, alignment-gated) and the fact that the reference dynamics \mathbb{P}_0 carries a co-evolving constraint field. Readers familiar with Garrahan et al. (2007) can think of ECF as the s -ensemble applied to a system with an endogenous, trajectory-coupled generator.

Definition 5.3 (Scaled cumulant generating function (SCGF)).

$$\Psi_T(\eta) = \frac{1}{T} \log \mathcal{Z}_T(\eta), \quad \Psi(\eta) = \lim_{T \rightarrow \infty} \Psi_T(\eta) \quad (\text{when it exists}).$$

Proposition 5.4 (Cumulants of coherence). *Whenever Ψ_T is differentiable, $\Psi'_T(\eta) = \frac{1}{T} \mathbb{E}_\eta[\mathcal{K}]$ and $\Psi''_T(\eta) = \frac{1}{T} \text{Var}_\eta(\mathcal{K}) \geq 0$, so Ψ_T is convex. The coherence susceptibility is $\chi_{\mathcal{K}} = \text{Var}_\eta(\mathcal{K})/T = \Psi''_T(\eta)$.*

Remark 5.5 (The estimator matters: bare variance is not $\chi_{\mathcal{K}}$). A subtle but important point: $\chi_{\mathcal{K}} = \Psi''_T(\eta)$ is the variance of \mathcal{K} under the *tilted* measure \mathbb{P}_η . Estimating it from direct ($\eta = 0$) Monte-Carlo trajectories is biased, and near a transition—where the relevant fluctuations are rare under \mathbb{P}_0 —the bias is severe. The correct estimator requires reweighting toward \mathbb{P}_η via cloning or population-dynamics algorithms (Giardinà et al., 2006; Nemoto et al., 2016; Ray et al., 2018). All exploratory results in Appendix A use the bare variance at $\eta = 0$ and should be read as indicative only; quantitative confirmation requires cloning.

6 Mechanism: the generalised Doob transform

A tilted path measure is not yet a dynamics—it is a reweighted distribution over histories. To answer the question “what process naturally generates the histories favoured by \mathbb{P}_η ?” we use the classical h-transform (Doob transform) (Doob, 1984), generalised to the extended state space.

Write $\mathcal{K}(\Gamma) = \sum_t k(z_t, z_{t+1})$ and define the *tilted transfer operator*

$$(\mathcal{T}_\eta f)(z) = \sum_{z'} \mathbb{P}_0(z' | z) e^{\eta k(z, z')} f(z').$$

This operator takes a “value function” f on the extended state space and maps it to a new function representing the expected $e^{\eta k}$ -weighted future.

Assumption 1 (Irreducibility and boundedness). The reference kernel on the recurrent extended states is irreducible and aperiodic, and k is bounded on $Z \times Z$.

Theorem 6.1 (Doob-transformed dynamics). *Under Assumption 1, \mathcal{T}_η has a unique maximal eigenvalue $e^{\Psi(\eta)} > 0$ with positive right eigenfunction φ_η . The kernel*

$$\mathbb{P}_\eta^{\text{Doob}}(z' | z) = \frac{\mathbb{P}_0(z' | z) e^{\eta k(z, z')} \varphi_\eta(z')}{e^{\Psi(\eta)} \varphi_\eta(z)} \quad (6)$$

is a proper stochastic matrix whose path law coincides with \mathbb{P}_η up to boundary terms vanishing as $T \rightarrow \infty$.

Proof. Perron–Frobenius / Krein–Rutman theory gives a simple maximal eigenvalue $\rho(\eta) > 0$ with positive φ_η ; set $\Psi(\eta) = \log \rho(\eta)$. Row normalisation of (6) follows from $\mathcal{T}_\eta \varphi_\eta = \rho(\eta) \varphi_\eta$. The path law telescopes the φ_η factors, recovering \mathbb{P}_η up to an $O(1/T)$ boundary term in the log. \square

Remark 6.2 (The eigenfunction as a value function). The eigenfunction $\varphi_\eta(z)$ is exactly the value function of an agent maximising expected future coherence from state z : it discounts moves toward regions from which high future coherence is unreachable, and upweights moves toward high-coherence futures. This is the precise sense in which “reach” behaves like a value function and why the control reading (Section 8) connects directly to KL-regularised optimal control (Todorov, 2009). The reach surrogate R_τ is a tractable approximation to $\log \varphi_\eta$.

7 Large deviations of coherence

We upgrade the earlier informal “Coherence Concentration Principle” to a theorem.

Assumption 2 (Ergodicity and bounded increments). (i) The extended chain $\{z_t\}$ under \mathbb{P}_0 is irreducible and aperiodic on a finite (or compact) recurrent set. (ii) The per-step increment $k(z, z')$ is bounded. (iii) The constraint update \mathcal{C} is Lipschitz and contracting in c on the relevant domain, so the joint chain (x_t, m_t, c_t) admits a unique stationary regime.

Theorem 7.1 (LDP for time-averaged coherence). *Under Assumption 2, the SCGF $\Psi(\eta) = \lim_T \frac{1}{T} \log \mathbb{E}_{\mathbb{P}_0} e^{\eta \mathcal{K}}$ exists, is finite, convex, and differentiable. The time-averaged coherence $k_T = \mathcal{K}/T$ satisfies a large-deviation principle with rate function $I(k) = \sup_\eta [\eta k - \Psi(\eta)]$, and $k_T \rightarrow k^* = \Psi'(0)$ almost surely with exponential concentration.*

Proof. \mathcal{K} is additive on an irreducible aperiodic chain with bounded increments (Assumptions (i)–(ii)). The SCGF equals log of the maximal eigenvalue of \mathcal{T}_η (Theorem 6.1), which is simple and analytic in η ; hence Ψ is convex and differentiable. Gärtner–Ellis (Dembo and Zeitouni, 1998; Touchette, 2009) yields the LDP. \square

Remark 7.2 (What is standard and what is genuinely new). The LDT machinery (Gärtner–Ellis, Doob transform, Perron–Frobenius) is entirely standard. We claim none of it as a contribution. The genuinely novel element is Assumption (iii): because c_t is deformed by the trajectory, the visible-state process is non-Markovian, and the theorem applies only on the extended chain and only when the constraint update is sufficiently contracting. When reinforcement overwhelms forgetting—when the system starts building in its own constraints so strongly that it cannot escape them—Assumption (iii) fails, ergodicity breaks down, and a *dynamical lock-in transition* occurs. The interesting physics lives exactly at the boundary of the theorem’s hypotheses.

Conjecture 1 (Coherence dynamical transition). There exists a critical reinforcement-to-forgetting ratio $\varrho_c = (\alpha/\kappa)_c$ at which $\Psi(\eta)$ loses differentiability at some η_c , $\chi_{\mathcal{K}}$ sharpens with system size, and the system transitions from an ergodic adaptive phase ($\varrho < \varrho_c$) to a non-ergodic lock-in phase ($\varrho > \varrho_c$). This is the trajectory-level analogue of first-order dynamical transitions in the s -ensemble (Garrahan et al., 2009), here driven by constraint co-evolution.

8 Two readings: descriptive versus prescriptive

A recurring ambiguity in coherence proposals is whether \mathbb{P}_η describes what a system *does* or what it *should* do. Both readings are legitimate; conflating them causes confusion.

Intrinsic (descriptive) reading. η is small and \mathbb{P}_η is interpreted as the rare-event tilt of the *actual* reference dynamics. The LDP of Theorem 7.1 tells us how often the natural dynamics produces atypical coherence values, and \mathbb{P}_η is the conditioned ensemble realising mean coherence $k = \Psi'(\eta)$. No external agency biases the system; coherence is a diagnostic of the spontaneous trajectory statistics. This is the appropriate reading for the spin-glass experiments.

Control (prescriptive) reading. $\eta > 0$ is a design objective and the Doob kernel (6) is the optimal policy of an agent maximising expected coherence under a KL trust region—equivalently, KL-regularised / linearly-solvable optimal control (Todorov, 2009; Kappen et al., 2012; Rawlik et al., 2013). The eigenfunction φ_η is the optimal value function and R_τ a tractable heuristic for it. This reading is appropriate for adaptive agents and design problems.

Remark 8.1. The two readings share the same mathematical object (5) and the same mechanism (6). They differ only in whether η is read off from observed statistics or imposed as a design objective. Stating which reading is in force removes equivocation and fixes the empirical content of every prediction.

9 Order parameters and the three regimes

9.1 Primary order parameters

Definition 9.1 (Primary order parameters).

$$\begin{aligned} \text{mean coherence: } \bar{k}_T &= \frac{1}{T} \mathbb{E}_\eta[\mathcal{K}_{\text{aligned}}], \\ \text{coherence susceptibility: } \chi_{\mathcal{K}} &= \frac{1}{T} \text{Var}_\eta(\mathcal{K}_{\text{aligned}}) = \Psi''_T(\eta), \\ \text{memory rigidity: } \rho_M &= \frac{I(x_{t+1}; m_t \mid x_t, c_t)}{H(x_{t+1} \mid x_t, c_t)}, \\ \text{constraint plasticity: } \pi_C &= \mathbb{E}_\eta[\|c_{t+1} - c_t\|], \\ \text{mean alignment: } \bar{A}_C &= \mathbb{E}_\eta[A_C(z_t, z_{t+1})], \\ \text{reach-frustration ratio: } \Theta &= \frac{\mathbb{E}_\eta[\Delta R_\tau]}{1 + \mathbb{E}_\eta[\Delta \Omega]}. \end{aligned}$$

Peaks in $\chi_{\mathcal{K}}$ signal a *dynamical* transition in trajectory statistics—the analogue of a susceptibility peak in equilibrium phase transitions—but need not coincide with any peak in static order parameters (energy, magnetisation). This non-coincidence is the central empirical claim.

9.2 Extended thermodynamic structure

The SCGF $\Psi(\eta)$ generates a thermodynamic-like structure at the trajectory level.

Definition 9.2 (Coherence entropy and free coherence). Given the SCGF $\Psi(\eta)$ with Legendre conjugate variable $k = \Psi'(\eta)$, define:

$$\text{coherence entropy density: } s_{\mathcal{K}}(k) = \eta k - \Psi(\eta), \tag{7}$$

$$\text{free coherence: } f_{\mathcal{K}}(\eta) = -\eta^{-1} \Psi(\eta). \tag{8}$$

Intuition

Physical interpretation. $s_{\mathcal{K}}(k)$ is the log-number of distinct histories that realise mean coherence k per unit time—the trajectory-level analogue of Boltzmann entropy. It answers the question: is high coherence achieved because many paths support it (high $s_{\mathcal{K}}$, diffuse) or because a single locked-in path dominates (low $s_{\mathcal{K}}$, concentrated)? A system approaching lock-in will show a sharp drop in $s_{\mathcal{K}}$ even before the mean coherence changes significantly. The free coherence $f_{\mathcal{K}}$ is the trajectory-level free energy; non-analyticities in $f_{\mathcal{K}}$ signal the same dynamical phase transitions as the s -ensemble.

9.3 Constraint-memory mutual information

Definition 9.3 (Constraint-memory mutual information).

$$I(C; M) = \iint p(c, m) \log \frac{p(c, m)}{p(c)p(m)} dc dm, \quad (9)$$

where $p(c, m)$ is the joint stationary distribution of constraint field and memory.

This quantity measures directly whether the constraint field has “absorbed” trajectory memory, which is the central mechanism claim. Three regimes:

- *Wandering*: $I(C; M) \approx 0$ —constraints and memory are statistically independent; past behaviour leaves no trace in the landscape.
- *Adaptive coherence*: $0 < I(C; M) \ll H(M)$ —the constraint field has begun to encode trajectory history but has not saturated.
- *Lock-in*: $I(C; M) \approx H(M)$ —the constraint field is an almost sufficient statistic for memory; the landscape is frozen into the past.

Crucially, $I(C; M)$ can be estimated directly from trajectories *without* computing $\Psi''(\eta)$, making it a practical complement to $\chi_{\mathcal{K}}$.

9.4 Constraint sufficiency ratio

Definition 9.4 (Constraint sufficiency ratio).

$$\mathcal{S}_C = \frac{I(c_t; \Gamma_{\text{future}})}{I(\Gamma_{\text{past}}; \Gamma_{\text{future}})}. \quad (10)$$

$\mathcal{S}_C \approx 0$: the constraint field carries negligible predictive information about future trajectories. $\mathcal{S}_C \approx 1$: c_t is an almost sufficient statistic for the future trajectory structure, formalising the claim that “constraints are sedimented path information.”

9.5 Hierarchy and the three regimes

Table 1 organises all order parameters by role.

The three regimes, with multi-level diagnostics:

Wandering ($\varrho \ll \varrho_c$). Low R , high Ω , large and structureless π_C , near-zero \bar{A}_C , $I(C; M) \approx 0$, $\mathcal{S}_C \approx 0$. The SCGF is smooth, $\chi_{\mathcal{K}}$ is small. Constraint updates are frequent but carry no trajectory information.

Adaptive coherence ($\varrho \sim \varrho_c$). High R , moderate Ω , intermediate π_C , rising \bar{A}_C , $I(C; M) > 0$, $\chi_{\mathcal{K}_{\text{aligned}}}$ peaks. $s_{\mathcal{K}}$ is high (many paths realise adaptive coherence), \mathcal{S}_C intermediate. The constraint field is being constructively reshaped by the trajectory.

Lock-in ($\varrho > \varrho_c$). Low R despite apparent stability, small π_C (constraints have frozen), $I(C; M) \approx H(M)$, $s_{\mathcal{K}}$ drops sharply, $\mathcal{S}_C \approx 1$. Ergodicity fails; the system is “stuck” in a self-reinforced configuration.

Level	Quantity	Role and interpretation
Extended state	$z_t = (x_t, m_t, c_t)$	Makes constraints explicit
Core observable	$\mathcal{K}_{\text{aligned}}(\Gamma)$	Path-consistent coherence
Generator	$\Psi(\eta)$	Coherence cumulant generating function
Susceptibility	$\chi_{\mathcal{K}} = \Psi''(\eta)$	Detects trajectory-level transitions
Entropy	$s_{\mathcal{K}}(k)$	Number of coherence-realising histories
Free coherence	$f_{\mathcal{K}}(\eta)$	Trajectory thermodynamic potential
Sedimentation	$I(C; M)$	Memory absorbed by the constraint field
Plasticity	π_C	Rate of constraint deformation
Alignment	\bar{A}_C	Trajectory-consistency of constraint updates
Sufficiency	\mathcal{S}_C	Predictive content of constraints

Table 1: Hierarchy of order parameters. The central diagnostic for the transition is $\chi_{\mathcal{K}_{\text{aligned}}}$; the practical companion for experiments is $I(C; M)$.

Intuition

Key distinction: stable \neq coherent. A locked-in system is stable—it hardly moves, its energy barely fluctuates—but its reach has collapsed. Stability in this framework is a property of the current state; coherence is a property of the trajectory and its relationship to future possibilities. The two can be (and often are) in opposition.

10 Instantiation: the adaptive Edwards–Anderson spin glass

10.1 Model setup

The standard Edwards–Anderson model. In the classical Edwards–Anderson (EA) spin glass (Edwards and Anderson, 1975), spins $s_i \in \{-1, +1\}$ are placed on a d -dimensional lattice with random couplings $J_{ij}^{(0)}$ drawn i.i.d. from a symmetric distribution (Gaussian or ± 1). The energy is $E(x, c) = -\sum_{\langle ij \rangle} J_{ij} s_i s_j$. The competing signs of the couplings create geometric frustration—it is impossible for all bonds to be simultaneously satisfied—which leads to the rugged energy landscape, metastability, and aging phenomenology that make spin glasses a paradigm for disordered systems (Binder and Young, 1986; Mézard et al., 1987; Young, 1998).

The adaptive extension. We modify the model so that couplings evolve according to past spin correlations:

$$m_i(t+1) = (1 - \lambda)m_i(t) + \lambda s_i(t), \quad (11)$$

$$J_{ij}(t+1) = (1 - \kappa + \alpha) J_{ij}(t) + \alpha s_i(t)s_j(t)/\sqrt{N} + b + \xi_{ij}(t). \quad (12)$$

Here $\lambda \in (0, 1)$ is a memory decay, $\alpha \geq 0$ is the *reinforcement rate* controlling how strongly correlated spin pairs strengthen their bond, $\kappa > 0$ is the *forgetting rate*, b is a small bias, and ξ_{ij} is weak noise. Setting $\alpha = 0$ recovers the standard quenched EA model.

Physical interpretation. When $\alpha > 0$, the coupling matrix is no longer quenched disorder—it is a dynamical variable that “remembers” past spin correlations. If spins i and j have repeatedly co-aligned, their coupling strengthens, making it even cheaper for them to align in the future. This is a physical implementation of the general mechanism: the trajectory (spin correlations) rewrites the constraints (couplings) that govern future trajectories (spin dynamics).

Example

Connection to neural plasticity. The update rule (12) is a continuous version of Hebb's rule: synaptic strength increases when pre- and post-synaptic neurons are co-active. The adaptive EA model is thus also a simple model of a Hopfield-like network (Hopfield, 1982) being continuously updated by its own activity—a toy model of memory consolidation and habit formation.

10.2 Analytic lock-in criterion

The effective linear map governing the coupling dynamics in (12) has spectral radius $|1 - \kappa + \alpha|$. The Lyapunov exponent is therefore

$$\lambda_L(\alpha, \kappa) = \log |1 - \kappa + \alpha|.$$

The coupling map is contracting when $\lambda_L < 0$, i.e. $|1 - \kappa + \alpha| < 1$. Taking the positive branch, the contraction boundary is $1 - \kappa + \alpha = 1$, giving

$$\boxed{\alpha_c = \kappa.} \tag{13}$$

For $\alpha < \kappa$: the map contracts; couplings remain bounded and the system is ergodic (Assumption 2(iii) holds). For $\alpha > \kappa$: the map is expanding; couplings grow without bound, eventually freezing the spin dynamics into a self-reinforced attractor.

Remark 10.1 (The importance of writing the update correctly). The threshold $\alpha_c = \kappa$ is only correct when the coupling update is written as $(1 - \kappa + \alpha)J_t + \dots$, so that α modifies the spectral radius of the linear part. If instead one writes $J_{t+1} = J_t + \alpha s_i s_j - \kappa J_t$, the effective coefficient is still $(1 - \kappa + \alpha)$, but it is tempting to misread α as shifting only the fixed point and κ as the sole contraction coefficient, leading to the incorrect conclusion that the stability criterion is $\kappa > 0$ regardless of α . The correct criterion is $\alpha < \kappa$, equivalently $\varrho = \alpha/\kappa < 1$.

10.3 Alignment in the spin-glass model

The constraint deformation at time t is $\Delta J_{ij}(t) = \alpha s_i(t) s_j(t) - (\kappa - \alpha) J_{ij}(t) + b + \xi_{ij}$. The move feature for bond $\langle ij \rangle$ is $u_{ij}(t) = s_i(t) s_j(t)$. The alignment is

$$A_C(t) = \frac{\sum_{\langle ij \rangle} \Delta J_{ij}(t) s_i(t) s_j(t)}{\|\Delta J(t)\| \cdot \|s s(t)\| + \varepsilon},$$

which is large and positive when the couplings are being reinforced precisely by the bonds that the current spin configuration is satisfying. This is the microscopic picture of coherent self-sculpting: the system strengthens the bonds it is currently using.

Dictionary. Table 2 maps every abstract concept to its concrete spin-glass realisation.

11 Comparison with neighbouring frameworks

We now give the detailed comparison promised in the introduction. The structure of this section is: for each neighbouring framework we state (i) what it does well, (ii) where it falls short for the problem of co-evolving constraints, and (iii) the precise mathematical relationship to the present work.

Abstract concept	Spin-glass realisation
Visible state x_t	Spin configuration $(s_1, \dots, s_N) \in \{-1, +1\}^N$
Memory m_t	Exponential trace $m_i(t)$ of local magnetisation
Constraint field c_t	Coupling matrix $J_{ij}(t) \in \mathbb{R}^{ E }$, evolved by α, κ
Viability V	$-E/N$ (energy density), defined independently of \mathcal{K}
Frustration Ω	Fraction of unsatisfied bonds
Reach R_τ	Log-volume of low-frustration basins accessible in τ flips
Dissipation D	Glauber entropy production or energy cost of a flip
Alignment A_C	Correlation of ΔJ_{ij} with $s_i s_j$
Coherence $\mathcal{K}_{\text{aligned}}$	Path sum of $G(A_C)[\Delta R_\tau - e\Omega - d \Delta E]$
Susceptibility $\chi_{\mathcal{K}}$	Second cumulant of \mathcal{K} under \mathbb{P}_η
Lock-in	$\alpha > \kappa$: couplings freeze, reach collapses

Table 2: Dictionary for the Edwards–Anderson instantiation.

11.1 Equilibrium statistical mechanics

What it does. Equilibrium theory (Landau and Lifshitz, 1980; Parisi, 1988) assigns a Boltzmann weight $p(x) \propto e^{-\beta E(x)}$ to each state. It provides the entire apparatus of thermodynamic phase transitions: order parameters, susceptibilities, correlation functions, critical exponents, and universality. The quenched-disorder version (Edwards and Anderson, 1975; Mézard et al., 1987) extends this to systems with random but *fixed* interactions via replica theory and cavity methods.

Where it falls short. The landscape is fixed. Two systems with the same Hamiltonian but different histories have the same future probability distribution. History leaves no trace. The Edwards–Anderson model with evolving couplings—our adaptive spin glass—is by construction outside the reach of equilibrium theory: the Hamiltonian is itself a function of the trajectory.

Relation to ECF. ECF reduces to a weighted sum over static equilibrium configurations when $\alpha = 0$ (no constraint co-evolution). The alignment condition and the reach functional are trivially satisfied (zero plasticity, zero alignment variation) and $\chi_{\mathcal{K}}$ tracks the heat capacity. This is the correct base case.

ECF vs Equilibrium stat. mech.

Equilibrium: weights states, fixed landscape, history irrelevant.

ECF: weights trajectories, co-evolving landscape, history explicitly encoded in c_t .

Mathematical relation: ECF \rightarrow equilibrium when $\alpha \rightarrow 0$.

11.2 Stochastic thermodynamics

What it does. Stochastic thermodynamics (Seifert, 2012; Jarzynski, 2011; Crooks, 1999; Esposito and Van den Broeck, 2010) elevates trajectories to primary objects, providing entropy production, work, heat, and fluctuation relations for systems far from equilibrium. The generator can be time-varying (e.g. driven by an external protocol), but the fundamental assumption is that the generator *itself* is an external input, not a dynamical variable.

Where it falls short. Stochastic thermodynamics does not model the generator as being modified by the trajectory. There is no concept of the constraint field co-evolving with the path, and therefore no way to define alignment, constraint plasticity, or the lock-in transition.

Relation to ECF. ECF inherits the dissipation functional directly from stochastic thermodynamics (Section 3.2) and would inherit fluctuation relations for the dissipation term if it were isolated. The coherence susceptibility $\chi_{\mathcal{K}}$ is an additional, genuinely novel quantity that has no counterpart in the stochastic thermodynamics literature.

ECF vs Stochastic thermodynamics

Stochastic thermo: trajectories on a fixed (or externally driven) generator; fluctuation relations for entropy production.

ECF: trajectories on a self-modifying generator; alignment-gated coherence and constraint sedimentation as additional objects.

Mathematical relation: $\text{ECF} \supset \text{stochastic thermo}$ for the dissipation term.

11.3 The dynamical s -ensemble

What it does. The s -ensemble (Garrahan et al., 2007, 2009; Lecomte et al., 2007; Jack and Sollich, 2010) tilts trajectory probabilities by a counting field s conjugate to an additive time-extensive observable (most commonly activity—the number of transitions). It reveals dynamical phase transitions between active and inactive trajectory phases, provides a large-deviation theory for trajectory observables, and has been applied to kinetically constrained models (Ritort and Sollich, 2003), glasses, and driven lattice gases.

Where it falls short. The kinetic rules are fixed. The tilting reweights trajectories under a fixed generator; it does not model the generator itself co-evolving with the trajectory.

Relation to ECF. The coherence-tilted ensemble (5) is *exactly* the s -ensemble with $s = -\eta$ and observable \mathcal{K} in place of activity. The Doob transform, the SCGF, and the rate function are all imported directly. The sole additional content of ECF is (a) the co-evolving constraint field in the reference dynamics \mathbb{P}_0 , and (b) the alignment-gated choice of \mathcal{K} that prevents the functional from collapsing into frustration tracking.

ECF vs s -ensemble

s -ensemble: tilts trajectories of a fixed-rule system by an activity observable; reveals dynamical phase transitions.

ECF: tilts trajectories of a self-modifying system by an alignment-gated reach functional; additionally predicts a lock-in transition driven by the co-evolving constraints.

Mathematical relation: ECF is the s -ensemble applied to an adaptive generator.

11.4 The free-energy principle and active inference

What it does. The free-energy principle (Friston, 2010; Friston et al., 2019; Parr et al., 2022) proposes that adaptive systems minimise variational free energy \mathcal{F} , a bound on the negative log-evidence of the agent’s generative model. In active inference, the agent performs both perceptual inference (updating beliefs) and action selection (selecting actions to make future observations conform to preferred priors). It has been applied to neuroscience (Friston, 2010), developmental biology, and artificial agents.

Where it falls short for ECF. The generative model is the central object but its constraint structure is treated as slowly varying or fixed at the timescale of interest. There is no equivalent of the fast co-evolving constraint field c_t , no alignment concept, no $\chi_{\mathcal{K}}$ transition, and the framework is typically agent-centred and prescriptive rather than providing a physical large-deviation theory.

Relation to ECF. The prescriptive / control reading of ECF (Section 8) is philosophically related to active inference: both involve an agent shaping its future observations toward preferred distributions. The mathematical implementation is different: ECF uses KL-regularised optimal control (Doob transform), active inference uses variational Bayes with a generative model. An interesting direction for future work is to ask whether the generative model of active inference can be interpreted as a constraint field and co-evolved at the fast timescale, yielding a version of active inference with explicit trajectory-level thermodynamics.

ECF vs Free-energy principle / active inference

FEP/AIF: agents minimise variational free energy of beliefs; prescriptive; generative model is slowly updated or fixed.

ECF: trajectory ensemble on a fast co-evolving constraint field; both descriptive and prescriptive readings; provides a large-deviation theorem and alignment-gated transition.

Relation: control reading of ECF is related to active inference; generative model \leftrightarrow constraint field correspondence is an open research question.

11.5 Reinforcement learning

What it does. Reinforcement learning (Sutton and Barto, 2018; Bellman, 1957) optimises a policy π to maximise expected cumulative reward in a Markov decision process (MDP) with transition kernel $P(x'|x, a)$ and reward $r(x, a)$. Value function methods (Q-learning, policy gradient) are among the most powerful computational tools for sequential decision making.

Where it falls short. The MDP—its transition kernel and reward function—is an exogenous specification. RL does not model the process by which the transition structure itself is modified by the agent’s behaviour. In KL-regularised / linearly-solvable RL (Todorov, 2009; Kappen et al., 2012; Rawlik et al., 2013), the optimal policy has exactly the form of our Doob kernel (6), but the reward function is specified externally rather than derived from reach.

Relation to ECF. The Doob kernel (6) with φ_η as value function is mathematically identical to linearly-solvable optimal control. The ECF contribution is (a) deriving the reward from a variational principle (reach minus frustration minus dissipation), (b) adding the alignment gate, and (c) treating the transition kernel itself as co-evolved—which corresponds, in RL language, to a non-stationary MDP whose transition kernel is updated by the agent’s own policy. The emerging literature on meta-learning and lifelong learning (Thrun and Pratt, 1998; Finn et al., 2017) addresses related problems but typically does not provide a trajectory-ensemble thermodynamic framework.

ECF vs Reinforcement learning

RL: maximise external reward in a fixed MDP; optimal policy for KL-RL has Doob form.
ECF: derive reward from reach; align constraint updates; co-evolve the MDP transition kernel.

Mathematical relation: linearly-solvable RL is a special case of ECF control reading with fixed generator and externally specified reward.

11.6 Empowerment and causal entropic forces

What they do. Empowerment (Klyubin et al., 2005; Salge et al., 2014) quantifies the channel capacity between an agent’s actions and its future sensory states, capturing the intrinsic value of maintaining options. Causal entropic forces (Wissner-Gross and Freer, 2013) derive a force field

Framework	Primitive	Landscape	Reach-like primitive	Distinguishing limitation
Equilibrium stat. mech. (Landau and Lifshitz, 1980)	States	Fixed	No	History irrelevant; no trajectory observable
Stochastic thermo. (Seifert, 2012)	Trajectories	Fixed	No	Generator not co-evolved by trajectory
s -ensemble (Garrahan et al., 2007)	Trajectories	Fixed	No	Fixed kinetic rules; no co-evolving c_t
Free-energy principle (Friston, 2010)	Beliefs	Fixed model	Implicit	No fast co-evolving constraint field or LDP
Reinforcement learning (Sutton and Barto, 2018)	Policy/return	Fixed MDP	Via value fn	Reward and kernel are exogenous
Empowerment (Klyubin et al., 2005)	States	Fixed	Yes	No constraint co-evolution or alignment
Causal entropic forces (Wissner-Gross and Freer, 2013)	States	Fixed	Yes	No constraint co-evolution or alignment
ECF (this work)	Extended trajectories	Co-evolving	Yes	$\chi_{\mathcal{K}}$ transition; alignment gate; lock-in in $\varrho = \alpha/\kappa$

Table 3: Comparison of ECF with neighbouring frameworks. The combination of co-evolving constraint field, alignment-gated reach, and trajectory-ensemble thermodynamics is the distinguishing feature. None of the comparison frameworks provides all three.

from the gradient of the log-volume of accessible futures, providing a physics-based mechanism for intelligent behaviour.

Where they fall short. Both operate on a fixed dynamic and do not model constraint co-evolution. Empowerment measures the volume of reachable futures but does not condition on path alignment or track how the set of reachable futures changes as a result of the system’s own trajectory.

Relation to ECF. “Reach” in ECF is a viability-weighted generalisation of the state-counting used in empowerment and causal entropic forces. The key differences are: (a) ECF conditions reach on the co-evolving constraint field, so the reachable set changes with the landscape; and (b) ECF incorporates path alignment, distinguishing systems that are actively expanding their reach through trajectory-consistent constraint sculpting from systems whose reach happens to be temporarily high.

Summary comparison.

12 Falsifiable predictions

The framework is scientific only insofar as it makes predictions that could fail. We list four predictions in decreasing order of current experimental support, stating the observable, the controlled comparison, the failure condition, and the honest status.

P1. Alignment-gated coherence transition invisible to static order parameters.

Statement. As $\alpha \rightarrow \alpha_c = \kappa$ at fixed β , the alignment-gated susceptibility $\chi_{\mathcal{K}_{\text{aligned}}}$ develops a peak that grows with system size L in the coupled adaptive system but is strongly suppressed in a time-shuffled control (where alignment values are permuted in time). Static observables (energy density, magnetisation) remain smooth across the same α .

Why the shuffle control matters. If the peak appeared equally in the shuffled control, it would mean $\chi_{\mathcal{K}_{\text{aligned}}}$ is tracking frustration fluctuations rather than genuine trajectory–constraint coherence. The coupled-to-shuffled ratio is the operational discriminator.

Failure condition. The alignment-gated peak is not significantly larger than the shuffled peak, or static observables co-track $\chi_{\mathcal{K}_{\text{aligned}}}$.

Current status. **Supported in spirit.** The coupled-to-shuffled peak ratio exceeds $100\times$ for $L = 8$ in direct Monte-Carlo experiments. However, quantitative confirmation via full SCGF cloning (which correctly estimates $\Psi''(\eta)$ rather than the biased $\eta = 0$ variance) at larger system sizes is still required. The bare-variance estimator used in current experiments is biased and should be read as indicative.

P2. Constraint-memory sedimentation transition.

Statement. The constraint-memory mutual information $I(C; M)$ increases monotonically from ≈ 0 in the wandering regime to near $H(M)$ in the lock-in regime, with an inflection near $\alpha_c = \kappa$ that co-locates with the peak of $\chi_{\mathcal{K}_{\text{aligned}}}$.

Why this matters. $I(C; M)$ does not require SCGF estimation and is directly computable from trajectory histograms. If it co-locates with the susceptibility peak, it provides a practical, unambiguous signature of the sedimentation mechanism.

Failure condition. $I(C; M)$ is non-monotone in α , or its inflection does not co-locate with the susceptibility peak.

Current status. **Not yet tested.** This is a new prediction proposed as the primary next experiment.

P3. Stability and coherence are different: lock-in confirmation.

Statement. In the lock-in regime ($\alpha > \kappa$), the system has minimal constraint plasticity π_C (most “stable”) yet strictly lower reach than in the adaptive regime. The system appears ordered and stable by static measures but has a collapsed future-option space.

Failure condition. Maximal stability coincides with maximal reach, or reach does not decline as α crosses α_c .

Current status. **Partially confirmed.** Susceptibility peaks at intermediate α and reach does not monotonically track stability. A robust *downturn* of reach at high α has proven difficult to observe, likely because the scanned range does not reach the lock-in plateau. Scanning $\varrho = \alpha/\kappa$ through $\varrho = 1$ is the indicated next step.

P4. Alignment-conditioned exploration window (secondary hypothesis).

Original prediction (demoted). An earlier version predicted a non-monotone “curiosity window” in the coherence pressure η : too little tilt gives wandering, too much gives lock-in, with optimal reach at intermediate η . This was not robustly observed: the η -reweighting suffered from poor effective sample sizes and reach tended to peak at the boundary of the scanned η range rather than at an interior point.

Reformulated (secondary) prediction. Conditional on sufficiently high constraint alignment ($G(A_C) \geq \theta_G$), mean rollout reach $\mathbb{E}[R_\tau]$ is non-monotone in constraint plasticity π_C or in the ratio $\varrho = \alpha/\kappa$, peaking at intermediate values. This is a claim about the *physical* control plane rather than the statistical η plane, which avoids the estimator problems.

Failure condition. Reach is monotone in ρ even within high-alignment trajectories.

Current status. **Weakly supported.** Exploratory results show a region of elevated reach at intermediate ρ but the interior maximum is not sharply established. This prediction is explicitly *secondary*—plausible on conceptual grounds but not yet experimentally confirmed—and should not be cited as a settled result.

13 Limitations and scope

- **Numerics are exploratory.** Current experiments use direct Monte-Carlo at $\eta = 0$ with small lattice sizes ($L \leq 16$) and approximate rollout reach. They establish the qualitative phenomenology but do not constitute a rigorous confirmation of the dynamical transition. Full SCGF estimation via population-dynamics cloning at larger system sizes ($L \in \{8, 12, 16, 24, 32\}$) is the required next step.
- **Reach estimation is approximate.** The first-order surrogate $\Delta R \approx -\Delta\Omega$ is used throughout. This is corrected for by the alignment gate (the gate makes it operational) but a direct comparison with exact reach on small systems and with learned value-function surrogates remains to be done.
- **The exploration window is not confirmed.** Prediction P4 is explicitly secondary. It should not be cited as an established result.
- **The theorem requires contractivity exactly where the transition is.** Theorem 7.1 assumes contractive constraint updates; the interesting physics is at the boundary $\alpha_c = \kappa$ where contractivity breaks down. Characterising the transition behaviour rigorously is the central open mathematical problem (Conjecture 1).
- **Only one model is worked out.** We focus on the 2D Edwards–Anderson model. Generalisations to NK fitness landscapes (Kauffman, 1993), adaptive ecological networks (May, 1972; Pascual and Dunne, 2006), and developmental models are clear directions but outside the scope of this paper.
- **$I(C; M)$ and S_C have not yet been measured.** These new order parameters are proposed but not yet computed in the spin-glass experiments.

14 Conclusion

We have presented the Entropic Coherence Framework, a path-statistical approach to adaptive systems whose trajectories reshape their own constraint landscape. The framework rests on four structural contributions: a variational derivation of the coherence functional, an explicit placement within large-deviation theory with a proved LDP, a Doob-transform mechanism supplying the generating dynamics, and an alignment gate that operationally separates genuine path-consistent constraint co-evolution from generic plasticity.

The most important empirical finding is that the alignment condition is not a refinement but a necessity: the bare coherence functional without alignment gating cannot distinguish coupled adaptive co-evolution from time-shuffled controls, while the alignment-gated functional separates them by one to two orders of magnitude in susceptibility. This establishes that “coherence” in our sense is a real, measurable property of adaptive trajectories, not a relabelling of energetics.

The reformulation of predictions reflects an honest engagement with negative results. The original curiosity window—a non-monotone interior maximum in coherence pressure η —was not robustly observed, and we demote it to a secondary hypothesis while promoting a reformulated

exploration window in the physical control plane. Lock-in, by contrast, is clearly confirmed: the system enters a frozen regime at high $\varrho = \alpha/\kappa$ with collapsed reach and minimal plasticity.

The framework’s core claim, stated in one sentence, is:

Adaptive systems remain viable when they transform history into future possibility through path-aligned constraint formation, avoiding both noise and lock-in.

The natural next steps are: (i) full SCGF cloning to confirm P1 quantitatively; (ii) direct measurement of $I(C; M)$ to test P2; and (iii) extension to at least one additional system class to begin assessing universality.

A Computational experiments: protocols and outcomes

All experiments use the 2D Edwards–Anderson adaptive spin glass of Section 10 with Glauber dynamics at fixed inverse temperature β . Lattice sizes $L \in \{6, 8, 10, 12, 16\}$ ($N = L^2$ spins), disorder averaging over $n_{\text{dis}} = 6\text{--}12$ independent coupling realisations, $n_{\text{rep}} = 4$ trajectory replicas per disorder realisation, and trajectory lengths $T = 200\text{--}500$ with burn-in $T_{\text{burn}} = 100\text{--}300$.

E1. Alignment-gated coherence transition (tests P1)

Protocol. Sweep $\alpha \in [0, \alpha_{\text{max}}]$ at fixed κ and β ; compute the alignment-gated susceptibility $\chi_{\mathcal{K}_{\text{aligned}}}$ via the bare trajectory variance (exploratory) and via a lightweight SCGF finite-difference diagnostic. Run in three modes: coupled (couplings updated by the actual spin trajectory), shuffled (same updates with alignment values permuted in time), and fixed ($\alpha = 0$, quenched disorder).

Outcome. The coupled-to-shuffled susceptibility ratio is approximately $13\times\text{--}123\times$ across system sizes $L = 6\text{--}10$ and parameter settings. The peak location shifts with L , consistent with a transition moving toward $\alpha_c = \kappa$ as $L \rightarrow \infty$. Static observables (energy, magnetisation) do not show a corresponding peak. These results are exploratory (bare-variance estimator); confirmation requires cloning-based $\Psi''(\eta)$.

E2. Constraint-memory sedimentation (tests P2)

Protocol. From the same trajectories, estimate $I(C; M)$ using a k -nearest-neighbour mutual information estimator on the joint distribution of (J_{ij}, m_i) . Scan α and record the α -value at which $I(C; M)$ shows its steepest increase.

Status. Not yet implemented. This is the proposed primary next experiment.

E3. Lock-in (tests P3)

Protocol. Extend the α scan to large $\varrho = \alpha/\kappa$; record reach, π_C , and the fraction of time the system spends in the initial basin. Compare to the analytic prediction $\varrho_c = 1$.

Outcome. At large ϱ , reach decreases monotonically to a frozen plateau; π_C collapses; the system fails to escape its initial configuration. This is the most robustly confirmed result. The transition location is consistent with $\varrho_c = 1$ but a precise finite-size scaling of $\varrho_c(L)$ has not been performed.

E4. Robustness of peak location (tests parameter sensitivity)

Protocol. Grid the multipliers (e, d) over a simplex; for each pair, record the α -value at which $\chi_{\mathcal{K}_{\text{aligned}}}$ peaks.

Outcome. Peak locations are stable across the grid: α^* varies by less than 15% across the full simplex for $L = 10$. This shows that the transition is not a fragile artefact of a particular choice of multipliers.

E5. Exploration window (tests P4, secondary)

Protocol. Within high-alignment trajectories ($G(A_C) \geq 0.2$), bin by π_C and record mean rollout reach \bar{R}_τ . Test for a non-monotone relationship.

Outcome. Results are mixed: for $L = 10$ a weak interior maximum is visible; for $L = 6$ and $L = 8$ it is not robust. This prediction is secondary.

E6. SCGF smoke test

Protocol. At $L = 6$, estimate $\Psi(\eta)$ via a finite-difference cloning diagnostic with a small clone population ($n_c = 100$).

Outcome. The SCGF curvature peaks at the same α as the direct susceptibility estimate, providing qualitative consistency. This is a smoke test; a production-scale cloning study requires $n_c \geq 10^3$ – 10^4 and $L \geq 16$.

References

- R. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- W. Bialek, I. Nemenman, and N. Tishby. Predictability, complexity, and learning. *Neural Computation*, 13(11):2409–2463, 2001.
- K. Binder and A. P. Young. Spin glasses: experimental facts, theoretical concepts, and open questions. *Reviews of Modern Physics*, 58(4):801–976, 1986.
- P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler. Transition path sampling. *Annual Review of Physical Chemistry*, 53:291–318, 2002.
- R. Chetrite and H. Touchette. Variational and optimal control representations of conditioned and driven processes. *Journal of Statistical Mechanics*, P12001, 2015.
- G. E. Crooks. Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences. *Physical Review E*, 60(3):2721–2726, 1999.
- A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Springer, 2nd edition, 1998.
- J. L. Doob. *Classical Potential Theory and Its Probabilistic Counterpart*. Springer, 1984.
- S. F. Edwards and P. W. Anderson. Theory of spin glasses. *Journal of Physics F: Metal Physics*, 5(5):965–974, 1975.
- M. Esposito and C. Van den Broeck. Three faces of the second law: I. Master equation formulation. *Physical Review E*, 82(1):011143, 2010.
- C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of ICML*, pages 1126–1135, 2017.
- K. Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.
- K. Friston, T. Parr, and B. de Vries. The graphical brain: belief propagation and active inference. *Network Neuroscience*, 1(4):381–414, 2019.
- J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland. Dynamical first-order phase transition in kinetically constrained models of glasses. *Physical Review Letters*, 98:195702, 2007.

- J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland. First-order dynamical phase transition in models of glasses. *Journal of Physics A*, 42:075007, 2009.
- C. Giardinà, J. Kurchan, and L. Peliti. Direct evaluation of large-deviation functions. *Physical Review Letters*, 96:120603, 2006.
- D. O. Hebb. *The Organisation of Behavior*. Wiley, 1949.
- J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.
- R. L. Jack and P. Sollich. Large deviations and ensembles of trajectories in stochastic thermodynamics. *Progress of Theoretical Physics Supplement*, 184:304–317, 2010.
- C. Jarzynski. Equalities and inequalities: irreversibility and the second law of thermodynamics at the nanoscale. *Annual Review of Condensed Matter Physics*, 2:329–351, 2011.
- H. J. Kappen, V. Gómez, and M. Opper. Optimal control as a graphical model inference problem. *Machine Learning*, 87:159–182, 2012.
- S. A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, 1993.
- A. S. Klyubin, D. Polani, and C. L. Nehaniv. Empowerment: a universal agent-centric measure of control. In *IEEE Congress on Evolutionary Computation*, 2005.
- L. D. Landau and E. M. Lifshitz. *Statistical Physics, Part 1*. Pergamon Press, 3rd edition, 1980.
- V. Lecomte, C. Appert-Rolland, and F. van Wijland. Thermodynamic formalism for systems with Markov dynamics. *Journal of Statistical Physics*, 127:51–106, 2007.
- R. M. May. Will a large complex system be stable? *Nature*, 238:413–414, 1972.
- M. Mézard, G. Parisi, and M. A. Virasoro. *Spin Glass Theory and Beyond*. World Scientific, 1987.
- H. Mori. Transport, collective motion, and Brownian motion. *Progress of Theoretical Physics*, 33(3):423–455, 1965.
- T. Nemoto, F. Bouchet, R. L. Jack, and V. Lecomte. Population-dynamics method with a multicanonical feedback control. *Physical Review E*, 93:062123, 2016.
- G. Parisi. *Statistical Field Theory*. Addison-Wesley, 1988.
- T. Parr, G. Pezzulo, and K. Friston. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press, 2022.
- M. Pascual and J. A. Dunne. *Ecological Networks: Linking Structure to Dynamics in Food Webs*. Oxford University Press, 2006.
- K. Rawlik, M. Toussaint, and S. Vijayakumar. On stochastic optimal control and reinforcement learning by approximate inference. *Proceedings of IJCAI*, 2013.
- U. Ray, G. K.-L. Chan, and D. T. Limmer. Exact fluctuations of nonequilibrium steady states from approximate auxiliary dynamics. *Physical Review Letters*, 120:210602, 2018.
- F. Ritort and P. Sollich. Glassy dynamics of kinetically constrained models. *Advances in Physics*, 52(4):219–342, 2003.

- C. Salge, C. Glackin, and D. Polani. Empowerment—an introduction. In *Guided Self-Organization: Inception*, Springer, 2014.
- U. Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Reports on Progress in Physics*, 75:126001, 2012.
- S. Still, D. A. Sivak, A. J. Bell, and G. E. Crooks. Thermodynamics of prediction. *Physical Review Letters*, 109:120604, 2012.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition, 2018.
- S. Thrun and L. Pratt. *Learning to Learn*. Springer, 1998.
- E. Todorov. Efficient computation of optimal actions. *Proceedings of the National Academy of Sciences*, 106(28):11478–11483, 2009.
- H. Touchette. The large deviation approach to statistical mechanics. *Physics Reports*, 478:1–69, 2009.
- A. D. Wissner-Gross and C. E. Freer. Causal entropic forces. *Physical Review Letters*, 110:168702, 2013.
- A. P. Young (Ed.). *Spin Glasses and Random Fields*. World Scientific, 1998.
- R. K. P. Zia and B. Schmittmann. Probability currents as principal characteristics in the statistical mechanics of non-equilibrium steady states. *Journal of Statistical Mechanics*, P07012, 2007.
- R. Zwanzig. *Nonequilibrium Statistical Mechanics*. Oxford University Press, 2001.