

Operational Coherence Maintenance and the Quantum–Classical Boundary: Formal Definitions, Falsifiable Protocols, and an Outlook for Cognitive Systems

Lluis Eriksson
Independent Researcher
lluiseriksson@gmail.com

December 2025

Scope note. No collapse mechanism is proposed; the Born rule is not derived; no claims are made about phenomenological consciousness.

Abstract

Maintaining quantum coherence against uncontrolled open-system dynamics is a control task with unavoidable thermodynamic cost. In a finite-dimensional setting with battery-assisted thermal operations at bath temperature T , we define an incremental (extra) maintenance power $P_{\text{extra}}(\rho)$ that isolates the cost of stabilizing coherence at fixed populations. For Markovian uncontrolled dynamics $\rho_t = e^{t\mathcal{L}}(\rho)$ we prove a single-law lower bound

$$P_{\text{extra}}(\rho) \geq k_{\text{B}}T \dot{C}_{\text{loss}}(\rho),$$

where $C(\rho) = S(\rho||\Delta[\rho])$ is relative-entropy coherence to energy pinching Δ and $\dot{C}_{\text{loss}}(\rho) := -\frac{d}{dt}C(\rho_t)|_{t=0}$. This statement is operational, observer-independent, and geometry-free.

We then formulate a falsifiable dynamical bridge between static locality/clustering and decoherence rates: the Rate Inheritance Principle (RIP). Using an operatorial Dirichlet-form identity, we highlight a concrete failure mode whereby near-zero Bohr-frequency channels can induce distance-independent rate floors, despite static clustering. These ingredients motivate a purely operational notion of a “cut”: a resource boundary separating maintainable coherence from regimes where classical-like effective descriptions are enforced under finite control budgets.

We provide falsifiable protocols that distinguish static one-shot work from sustained maintenance power across quantum platforms and interface geometries, including a numerical stress test (uniform floor versus collar-induced suppression) in a gapped transverse-field Ising chain with remote dissipation. Finally, we offer an Outlook for cognitive systems as resource-limited physical agents, connecting the operational resource boundary to the Free-Energy Principle at a methodological (non-phenomenological) level. We do not propose collapse mechanisms, do not derive the Born rule, and make no claims about phenomenological consciousness.

Keywords: coherence, open quantum systems, thermodynamic control, relative entropy, decoherence, clustering, resource boundary, falsifiable protocols, Free-Energy Principle.

1 Introduction

Decoherence explains the passive suppression of interference in open quantum systems and the emergence of classical-like effective behavior [1, 2, 3]. This paper focuses on a complementary, operational question: what minimal control resources are required to prevent coherence loss and keep a desired state stationary against an uncontrolled dynamics?

We work within an explicit thermodynamic control model (battery-assisted thermal operations) and define a notion of incremental maintenance power that isolates the cost of stabilizing coherence at fixed populations. Our main result is a single-law lower bound on this incremental power in terms of the instantaneous loss rate of relative-entropy coherence.

Separately, gapped systems exhibit static clustering and distance-dependent suppression of correlations [9, 10, 11]. A common heuristic expectation is that dynamical decoherence rates should inherit similar suppression with separation. We isolate this as a falsifiable hypothesis (RIP), and we document a specific failure mechanism: near-zero Bohr-frequency channels can induce distance-independent rate floors. This motivates an operational reading of the quantum–classical boundary as a resource boundary in state space.

Non-claims (read first). We do not propose a collapse mechanism; we do not derive the Born rule; and we make no claims about phenomenological consciousness. Our claims are operational: they quantify minimal work/power requirements under explicit control primitives. Any “cut” discussed here is operational (a feasibility boundary), not ontological.

2 Operational framework

2.1 Thermal reference states and relative entropy

Let X be finite-dimensional with Hamiltonian H_X and inverse temperature $\beta := 1/(k_B T)$. The thermal state is

$$\gamma_X := \frac{e^{-\beta H_X}}{\text{Tr}(e^{-\beta H_X})}.$$

For states ρ, σ with $\text{supp}(\rho) \subseteq \text{supp}(\sigma)$, the relative entropy is

$$S(\rho||\sigma) := \text{Tr}[\rho(\log \rho - \log \sigma)],$$

and $S(\rho||\sigma) := +\infty$ otherwise.

Define the (non-equilibrium) free energy of X at temperature T as

$$F_X(\rho) := k_B T S(\rho||\gamma_X).$$

2.2 Battery-assisted thermal operations and work accounting

We consider a system S , bath B , and battery W with Hamiltonians H_S, H_B, H_W . A battery-assisted thermal operation on S is any reduced map induced by a global unitary U on SBW such that

$$[U, H_S + H_B + H_W] = 0,$$

acting as

$$\rho_S \mapsto \Phi(\rho_S) := \text{Tr}_{BW}[U(\rho_S \otimes \gamma_B \otimes \sigma_W)U^\dagger],$$

where the bath input is thermal γ_B and σ_W is the battery state.

We define work as the increase of battery non-equilibrium free energy,

$$W := F_W(\sigma'_W) - F_W(\sigma_W),$$

and average (cycle-averaged) power as $P := W/\Delta t$ over cycle time Δt .

2.3 Coherence functional and uncontrolled dynamics

Let $H_S = \sum_n E_n \Pi_n$ be the spectral decomposition. Define energy pinching

$$\Delta[\rho] := \sum_n \Pi_n \rho \Pi_n.$$

Define relative-entropy coherence

$$C(\rho) := S(\rho \| \Delta[\rho]).$$

Let the uncontrolled dynamics be a Markovian semigroup

$$\rho_t := e^{t\mathcal{L}}(\rho),$$

where \mathcal{L} is a GKLS generator with stationary thermal state γ_S , i.e. $\mathcal{L}(\gamma_S) = 0$. Define the instantaneous coherence-loss rate

$$\dot{C}_{\text{loss}}(\rho) := -\left. \frac{d}{dt} C(\rho_t) \right|_{t=0}.$$

2.4 Stroboscopic maintenance strategies and incremental power

Fix $\Delta t > 0$. A stroboscopic control strategy s maintains ρ if, for each cycle, the uncontrolled evolution $\rho \mapsto e^{\Delta t \mathcal{L}}(\rho)$ is followed by a battery-assisted thermal operation $\mathcal{E}^{(\Delta t)}$ such that

$$\mathcal{E}^{(\Delta t)} \circ e^{\Delta t \mathcal{L}}(\rho) = \rho.$$

Let $W_n^{(\Delta t)}$ denote the battery work in cycle n . Define the asymptotic power of s as

$$P(s) := \limsup_{N \rightarrow \infty} \frac{1}{N \Delta t} \sum_{n=1}^N W_n^{(\Delta t)}.$$

Let $\text{Ctrl}(\rho)$ be the set of strategies that maintain ρ . Likewise define $\text{Ctrl}(\Delta[\rho])$ for strategies that maintain the diagonal state $\Delta[\rho]$ under the same bath and H_S .

Paired-strategy incremental power. To isolate the incremental cost of coherence stabilization while avoiding invalid manipulations of differences of infima, we define the incremental maintenance power as an infimum over pairs of strategies:

$$P_{\text{extra}}(\rho) := \inf_{s \in \text{Ctrl}(\rho), s_{\text{diag}} \in \text{Ctrl}(\Delta[\rho])} (P(s) - P(s_{\text{diag}})).$$

Remark 1 (Baseline matching). *The baseline $\text{Ctrl}(\Delta[\rho])$ is defined relative to the same uncontrolled dynamics $e^{\Delta t \mathcal{L}}$, the same Hamiltonian H_S , the same bath temperature T and bath input γ_B , and the same admissible control class (battery-assisted thermal operations). This matching is essential: P_{extra} is intended to isolate the incremental power required to stabilize coherence at fixed populations, rather than differences in population-control difficulty.*

2.5 Assumptions

We assume:

- (A1) S, B, W are finite-dimensional.
- (A2) The uncontrolled dynamics are Markovian (GKLS) with thermal fixed point γ_S .
- (A3) Each control stroke is implemented by a global unitary U with $[U, H_S + H_B + H_W] = 0$ and bath input γ_B .
- (A4) Work is accounted as ΔF_W at temperature T .
- (A5) Stroboscopic regularity: there exists $\Delta t_0 > 0$ and a constant $L > 0$ such that the optimal per-cycle extra work $W_{\text{extra}}(\Delta t)$ satisfies a Lipschitz condition

$$|W_{\text{extra}}(\Delta t) - W_{\text{extra}}(\Delta t')| \leq L |\Delta t - \Delta t'| \quad \text{for all } \Delta t, \Delta t' \in (0, \Delta t_0].$$

This is a sufficient condition to justify exchanging limsup and the small- Δt expansion used to pass from per-cycle work to asymptotic power.

- (A6) ρ is full-rank, or else right-derivatives are used.

3 Main inequality: power lower bound from coherence loss

Lemma 1 (Product-reference decomposition). *Let $\Gamma := \gamma_S \otimes \gamma_B \otimes \gamma_W$. For any state ρ_{SBW} ,*

$$S(\rho_{SBW} \|\Gamma) = S(\rho_S \|\gamma_S) + S(\rho_B \|\gamma_B) + S(\rho_W \|\gamma_W) + I_\Gamma(\rho_{SBW}),$$

where $I_\Gamma(\rho_{SBW}) \geq 0$ is a relative-entropy multi-information term.

Proof sketch. Expand $S(\rho \|\Gamma) = \text{Tr}[\rho \log \rho] - \sum_X \text{Tr}[\rho_X \log \gamma_X]$. Add and subtract $\sum_X \text{Tr}[\rho_X \log \rho_X]$ to identify $I_\Gamma(\rho)$ as a relative-entropy multi-information term. Nonnegativity follows from data processing / joint convexity. \square

Theorem 1 (Single-law lower bound on incremental maintenance power). *Under assumptions (A1)–(A6), the incremental maintenance power satisfies*

$$P_{\text{extra}}(\rho) \geq k_B T \dot{C}_{\text{loss}}(\rho).$$

Remark 2 (Interpretation). *Theorem 1 is operational and observer-independent: it bounds the incremental power required to stabilize coherence against a given uncontrolled dynamics, under explicit thermodynamic control primitives. It does not invoke collapse postulates, and it does not depend on geometry.*

Remark 3 (Units). *$C(\rho)$ is dimensionless (nats), so \dot{C}_{loss} has units of 1/time. Multiplication by $k_B T$ yields power.*

Proof roadmap. A detailed proof outline is provided in Appendix A.

4 Static clustering and the dynamical bridge (RIP)

Gapped systems exhibit exponential decay of static correlations with distance [9, 10, 11]. In massive relativistic field theory, correlators can exhibit modified-Bessel envelopes $K_\nu(m\epsilon)$ at separation ϵ . These are static facts. To connect geometry to dynamical maintenance power one needs an additional dynamical bridge: how effective decoherence rates depend on separation.

4.1 Effective rate envelopes and RIP

Let ϵ denote an operational separation (e.g., distance between a controlled subsystem and a localized coupling/dissipation region). Let $\kappa(\epsilon)$ denote an effective dynamical envelope, e.g., extracted from a Heisenberg-picture Dirichlet form, or operationally from decay curves (Protocol B).

Rate Inheritance Principle (weak form; falsifiable). In a validated Davies/weak-coupling Markovian regime [4, 5, 3], we hypothesize that there exist constants $A, \beta > 0$ and a polynomial p such that

$$\kappa(\epsilon) \leq A p(m\epsilon) e^{-\beta m\epsilon}$$

in the regime of interest. RIP is not a theorem; it is a falsifiable interface hypothesis.

4.2 A concrete failure mode: near-zero-frequency floors

In Davies-type generators, the coupling decomposes into Bohr-frequency components $S(\omega)$. A near-zero ($\omega = 0$) channel can generate a distance-independent floor under appropriate conditions.

Let $\sigma = \gamma_S$ and define the KMS inner product

$$\langle A, B \rangle_\sigma := \text{Tr}(\sigma^{1/2} A^\dagger \sigma^{1/2} B),$$

with norm $\|O\|_{2,\sigma}^2 := \langle O, O \rangle_\sigma$. For a Hermitian $\omega = 0$ channel with $S(0)$ and rate $\gamma(0) > 0$, the Heisenberg generator is

$$\mathcal{L}_0^\dagger(O) = \gamma(0) \left(S(0) O S(0) - \frac{1}{2} \{S(0)^2, O\} \right).$$

Define the Dirichlet form $\mathcal{E}_\sigma(O) := -\text{Re} \langle O, \mathcal{L}_0^\dagger(O) \rangle_\sigma$. Then the $\omega = 0$ contribution satisfies the exact identity

$$\mathcal{E}_\sigma^{(0)}(O) := -\text{Re} \langle O, \mathcal{L}_0^\dagger(O) \rangle_\sigma = \frac{\gamma(0)}{2} \| [S(0), O] \|_{2,\sigma}^2.$$

Remark 4 (When floors occur). *A distance-independent floor arises when $S(0)$ retains non-decaying overlap with observables supported at distance ϵ (e.g., extensive conserved components, symmetry-induced degeneracies, or nonlocal effective $S(0)$). In contrast, if $S(0)$ is quasilocal with decaying tails, explicit interfaces (collars) can suppress $\| [S(0), O] \|_{2,\sigma}$ with separation, restoring RIP-compatible behavior.*

5 The operational cut: a resource boundary

Let P_{avail} denote an available control power budget. We define the operational cut as the frontier in state space where $P_{\text{extra}}(\rho) \leq P_{\text{avail}}$ ceases to hold. Above the cut, coherence is unmaintainable and classical-like effective descriptions are operationally enforced under the chosen control model.

Proposition 1 (Resource-horizon no-go (conditional)). *Fix $C_0 > 0$ and a nonempty family $\mathcal{F}_{C_0} = \{\rho : C(\rho) \geq C_0\}$. If there exists $\dot{C}_0 > 0$ such that $\inf_{\rho \in \mathcal{F}_{C_0}} \dot{C}_{\text{loss}}(\rho) \geq \dot{C}_0$, then*

$$\inf_{\rho \in \mathcal{F}_{C_0}} P_{\text{extra}}(\rho) \geq k_{\text{B}}T \dot{C}_0.$$

Hence, if $P_{\text{avail}} < k_{\text{B}}T \dot{C}_0$, sustained maintenance is impossible for any $\rho \in \mathcal{F}_{C_0}$.

Remark 5. *The uniform floor is model-dependent; Proposition 1 is a conditional design principle. It becomes particularly relevant in regimes with persistent low-frequency channels.*

6 Falsifiable protocols

We present falsifiable protocols that test the maintenance inequality and the static-to-dynamic interface hypothesis.

6.1 Protocol A: testing the maintenance inequality on quantum hardware

Goal. Test $P_{\text{extra}}(\rho) \geq k_{\text{B}}T \dot{C}_{\text{loss}}(\rho)$.

Setup. Use a qubit or small register with engineered Markovian noise and a thermal reference at known T (or an effective T inferred from detailed-balance ratios). Implement:

- a coherence-maintaining strategy $s \in \text{Ctrl}(\rho)$,
- a population-maintaining baseline $s_{\text{diag}} \in \text{Ctrl}(\Delta[\rho])$.

Measurements.

- Estimate $\dot{C}_{\text{loss}}(\rho)$ from short-time data: prepare ρ , evolve under \mathcal{L} for small Δt , reconstruct $\rho_{\Delta t}$ (tomography or randomized estimation), compute $C(\rho_{\Delta t}) - C(\rho)$.
- Measure wall-plug incremental power proxy $P_{\text{wall,extra}} := P_{\text{wall}}(s) - P_{\text{wall}}(s_{\text{diag}})$ and propagate calibration uncertainties (Appendix B).

Acceptance region. With power uncertainty ϵ_P and coherence-rate uncertainty ϵ_C ,

$$P_{\text{wall,extra}} \geq k_{\text{B}}T \dot{C}_{\text{loss}}(\rho) - (\epsilon_P + \epsilon_C)$$

is consistent with the theorem. Systematic violations beyond uncertainties indicate failure of assumptions (non-Markovianity, thermal reference mismatch, control-class mismatch, or power-proxy miscalibration), not failure of the mathematics.

6.2 Protocol B: interface geometry and dynamical suppression (floor vs collar)

Coherence functional used in Protocol B. For any pinching channel Π we may define $C_{\Pi}(\rho) := S(\rho \parallel \Pi[\rho])$. In Protocol B we use a local-basis pinching (e.g., X basis) as an experimentally convenient proxy for local fragility; conclusions are operational and do not require energy pinching for the local readout.

Goal. Stress-test RIP by probing whether dynamical influence decays with separation (suppression) or exhibits a distance-independent floor.

Numerical interface test: influence floors vs collar suppression (TFIM)

We include a compact numerical stress test illustrating (i) a floor without an explicit interface and (ii) suppression when a collar is introduced. This is intended as a falsifiable diagnostic rather than a universal fit.

Model and readout. Consider a transverse-field Ising chain of length N_{tot} with Hamiltonian

$$H(L) = - \sum_{i=0}^{N_{\text{tot}}-2} J_i(L) X_i X_{i+1} - h \sum_{i=0}^{N_{\text{tot}}-1} Z_i.$$

In the uniform geometry, $J_i(L) = J$ for all i . In the collar geometry, a separation parameter L is implemented by weakening the links on the path from the system site (index 0) to the dissipative site:

$$J_i(L) = J_{\text{buf}} \ll J \quad \text{for } i = 0, 1, \dots, L,$$

so that the causal path from site 0 to the dissipative site crosses $L + 1$ weak links.

Dissipation is applied at site $j_{\text{diss}} = L + 1$ via a single jump operator (dephasing in the runs shown),

$$L_{\text{jump}} = \sqrt{\gamma} Z_{j_{\text{diss}}}.$$

The initial state is $|+\rangle^{\otimes N_{\text{tot}}}$. As a local readout on site 0, let $\rho_0(t; L)$ be the reduced state at time t .

Define the X -basis pinching channel

$$\Delta_x[\rho] := \frac{1}{2}(\rho + X\rho X),$$

and the relative-entropy coherence proxy

$$C_x(t; L) := S(\rho_0(t; L) \parallel \Delta_x[\rho_0(t; L)]).$$

Fixed-horizon influence metric. To avoid imposing monotone decay on $C_x(t; L)$, quantify remote dissipation by comparing to the unitary baseline under the same Hamiltonian $H(L)$. For a fixed horizon T , define

$$D_T(L) := \max_{t \in [0, T]} \left| \overline{C_x^{\text{noisy}}(t; L)} - C_x^{\text{unitary}}(t; L) \right|,$$

where the overbar denotes averaging over Monte Carlo wave-function trajectories, and $C_x^{\text{unitary}}(t; L)$ is computed with the jump operator removed. Operationally, $D_T(L)$ measures how strongly remote dissipation changes the local time series in $[0, T]$.

Results: floor without interface, suppression with a collar. Figure 1 compares $D_T(L)$ for two geometries at matched microscopic parameters:

1. **Uniform links** ($J_i(L) = J$): $D_T(L)$ is approximately independent of L over the tested range (an influence floor).
2. **Collar links** ($J_i(L) = J_{\text{buf}} \ll J$): $D_T(L)$ decreases strongly with L (geometric suppression).

This contrast supports the organizational point: static locality does not force a universal dynamical envelope; an explicit interface mechanism can be necessary for dynamical inheritance.

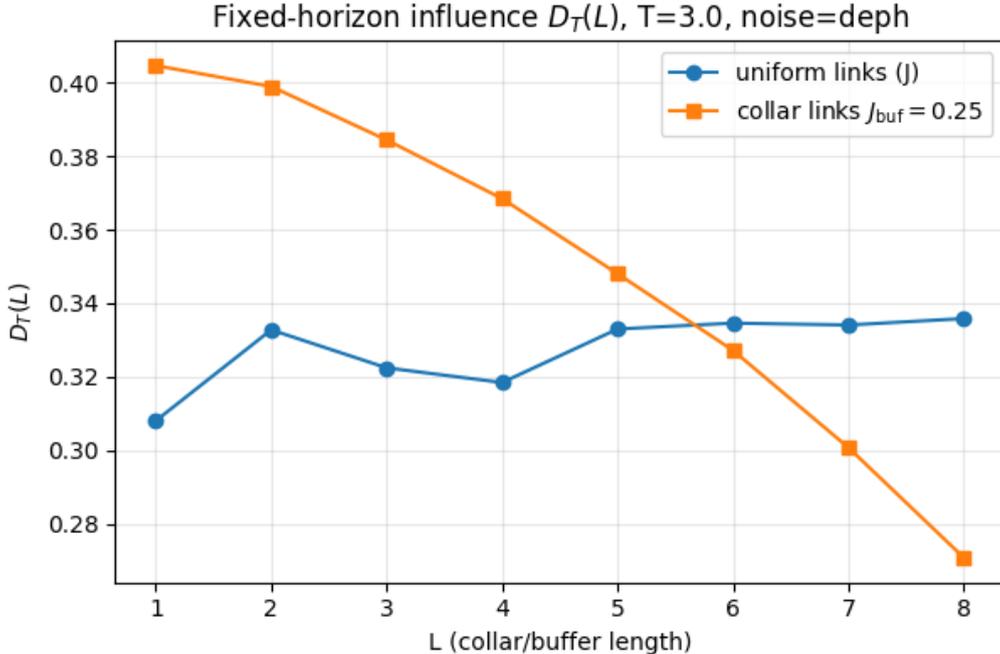


Figure 1: Fixed-horizon influence $D_T(L)$ for a TFIM chain with dissipative site at $j_{\text{diss}} = L + 1$. Uniform couplings $J_i(L) = J$ exhibit an influence floor (no suppression with L), while introducing a collar by weakening the links to $J_{\text{buf}} \ll J$ yields clear suppression with L . Parameters: $N_{\text{tot}} = 14$, $h/J = 1.5$, $\gamma = 0.2$, $J_{\text{buf}} = 0.25$, $T = 3.0$, $dt = 0.005$.

Collar opacity sweep and fitted suppression rate. Sweep J_{buf} and fit collar-geometry influence curves to

$$D_T(L) \approx Ae^{-\mu L} + D_\infty,$$

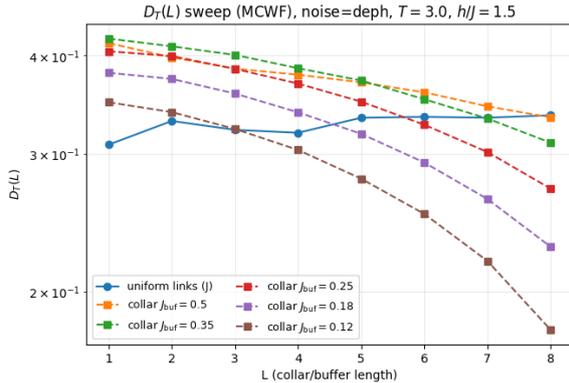
where μ is an effective suppression rate and D_∞ captures finite-size/horizon saturation. Estimate pointwise uncertainties in $D_T(L)$ by batching trajectories and use weighted least squares for (A, μ, D_∞) .

Table 1: Weighted fits of collar-geometry influence curves to $D_T(L) \approx Ae^{-\mu L} + D_\infty$. Pointwise uncertainties for $D_T(L)$ are estimated via batching MCWF trajectories; fit uncertainties are 1σ from the weighted covariance.

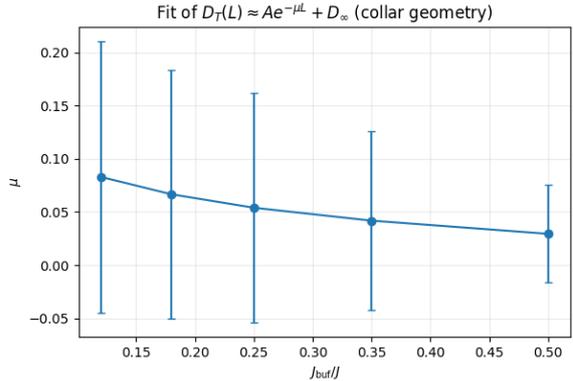
J_{buf}/J	μ	σ_μ	D_∞	σ_{D_∞}
0.25	0.09119	0.00089	1.67×10^{-14}	0.0026
0.18	0.13180	0.00065	1.09×10^{-14}	0.00098
0.12	0.130304	0.00061	7.27×10^{-15}	0.00081

Interpretation. The robust conclusion is the floor-versus-suppression contrast: without an explicit collar (uniform couplings) the influence proxy exhibits a distance-independent floor, whereas introducing a collar yields geometric suppression with separation.

Finite-size and finite-horizon caveat. The proxy $D_T(L)$ is defined at fixed horizon T and for finite N_{tot} . The observed “floor” in the uniform geometry is therefore an operational diagnostic in this finite setting, not a claim of a thermodynamic-limit theorem. In larger systems and/or longer horizons, $D_T(L)$ may exhibit additional structure (e.g., light-cone effects, recurrences, or



(a) Fixed-horizon influence $D_T(L)$ with batch standard errors.



(b) Weighted-fit μ versus J_{buf}/J .

Figure 2: Collar-opacity sweep with uncertainty quantification. Extract μ from $D_T(L) \approx Ae^{-\mu L} + D_\infty$. Parameters: $n_{\text{traj}} = 600$ split into 6 batches.

slower decays). The key falsifiable contrast is comparative: at matched (N_{tot}, T) , introducing an explicit collar produces a robust decrease of $D_T(L)$ with L relative to the uniform case.

6.3 Protocol C: dissociating static one-shot work from maintenance power

Goal. Demonstrate that static shaping costs (one-shot work) need not track maintenance power.

Setup. Prepare a joint state with measurable correlations (e.g., mutual information or a leakage proxy), perform a one-shot protocol that reduces correlations and record the work W . Separately, maintain the post-shaped state against noise for time τ and measure average incremental power P_{extra} .

Prediction. W scales with static correlation measures, whereas P_{extra} scales with dynamical fragility \dot{C}_{loss} ; the two can dissociate across protocols.

7 Outlook for cognitive systems (speculative)

We present a methodological Outlook for cognitive systems viewed as resource-limited physical agents. No claims are made about phenomenological consciousness, nor about quantum coherence in brains.

7.1 Resource boundary and the Free-Energy Principle (non-phenomenological)

The Free-Energy Principle (FEP) proposes that adaptive agents minimize variational free energy as a proxy for surprise [14, 15]. Our framework provides a complementary constraint: sustaining representations against uncontrolled dynamics incurs an unavoidable maintenance power cost that obeys $P_{\text{extra}} \geq k_B T \dot{C}_{\text{loss}}$ under explicit control primitives.

A minimal bridge (conditional). Let $\mathcal{F}_{\text{var}}(q)$ be a variational free-energy functional for an internal encoding q (predictive density or latent representation), and let $P_{\text{extra}}(q, \pi)$ denote a physical maintenance-power proxy under a policy π . We expect a multi-objective trade-off

$$\min_{q, \pi} \left\{ \mathbb{E}_\pi[\mathcal{F}_{\text{var}}(q)] + \lambda P_{\text{extra}}(q, \pi) \right\},$$

with $\lambda > 0$. When a proxy $\dot{C}_{\text{loss}}(q)$ tracks representational fragility under the agent's noise, the inequality implies

$$P_{\text{extra}}(q, \pi) \geq k_B T \dot{C}_{\text{loss}}(q),$$

so fragile encodings become disfavored at fixed budget even if they reduce $\mathcal{F}_{\text{var}}(q)$.

Testable hypotheses (non-quantum).

- **FEP–maintenance trade-off.** In synthetic agents with explicit energy/compute budgets, policies achieving lower \mathcal{F}_{var} via fragile encodings will either increase measured maintenance power or revert to robust/low-cost encodings once P_{avail} is capped.
- **Interface manipulation.** Architectural collars (attenuated couplings) reduce dynamical influence of distal perturbations, enabling lower incremental maintenance cost at fixed task performance.
- **Work vs power dissociation.** One-shot shaping to reduce uncertainty (work) can be orthogonal to steady maintenance power; the latter tracks fragility, not static task entropy.

Neural energy context. Classical estimates suggest a substantial fraction of metabolic power supports maintenance (e.g., ion pumping sustaining membrane potentials against leak) [16, 17]. Our inequality provides an operational template for relating minimal maintenance costs to dynamical fragility terms. This is a methodological connection only.

8 Related work

This work complements decoherence [1, 2], resource-theoretic thermodynamics and coherence [6, 7, 8, 12, 13], Markovian open-system dynamics [3, 4, 5], and static clustering in gapped systems [9, 10, 11]. The cognitive Outlook connects methodologically to energy-efficient inference and predictive processing [14, 15] and to neuroenergetics [16, 17].

9 Limitations and non-claims

Theorem 1 assumes finite dimensions, Markovianity (GKLS), bath thermal inputs, and battery-assisted thermal operations. Extensions to strong coupling, non-Markovian regimes, or algebraic (type-III) limits require additional technical input. RIP is a falsifiable hypothesis and can fail in the presence of near-zero-frequency floors. We do not propose collapse mechanisms, do not derive the Born rule, and make no claims about phenomenological consciousness.

10 Conclusion

We formalized incremental coherence-maintenance power and proved a single-law lower bound linking it to the instantaneous coherence-loss rate under explicit thermodynamic control primitives. We isolated the key dynamical hinge connecting static locality to decoherence-rate scaling (RIP), together with a concrete failure mode that induces distance-independent floors. These results motivate an operational resource boundary (a “cut”) separating maintainable coherence from regimes where classical-like effective descriptions are enforced under finite budgets. We provided falsifiable protocols and a cautious Outlook linking resource boundaries to methodological themes in cognitive systems and the Free-Energy Principle.

A Proof of Theorem 1 (detailed outline)

We provide a structured outline sufficient for independent verification.

Step 0: Notation. Let $\Gamma := \gamma_S \otimes \gamma_B \otimes \gamma_W$. In one stroboscopic cycle, the global input is

$$\rho_{SBW} := \rho_S \otimes \gamma_B \otimes \sigma_W,$$

and the output is $\rho'_{SBW} := U\rho_{SBW}U^\dagger$, with $[U, H_S + H_B + H_W] = 0$.

Step 1: Invariance of relative entropy under energy-conserving control. Since $\Gamma \propto e^{-\beta(H_S + H_B + H_W)}$ and $[U, H_S + H_B + H_W] = 0$, we have $[U, \Gamma] = 0$, hence

$$S(U\rho U^\dagger \| \Gamma) = S(\rho \| \Gamma).$$

Therefore,

$$S(\rho'_{SBW} \| \Gamma) = S(\rho_{SBW} \| \Gamma).$$

Step 2: Product-reference decomposition and per-cycle work/free-energy bound. By Lemma 1,

$$S(\rho'_{SBW} \| \Gamma) = S(\rho'_S \| \gamma_S) + S(\rho'_B \| \gamma_B) + S(\rho'_W \| \gamma_W) + I_\Gamma(\rho'_{SBW}) \geq S(\rho'_S \| \gamma_S) + S(\rho'_W \| \gamma_W).$$

For the input product state,

$$S(\rho_{SBW} \| \Gamma) = S(\rho_S \| \gamma_S) + S(\sigma_W \| \gamma_W),$$

since ρ_{SBW} is product and $S(\gamma_B \| \gamma_B) = 0$.

Combining steps 1–2 yields

$$S(\rho'_W \| \gamma_W) - S(\sigma_W \| \gamma_W) \geq S(\rho'_S \| \gamma_S) - S(\rho_S \| \gamma_S).$$

Multiplying by $k_B T$ gives

$$W := F_W(\rho'_W) - F_W(\sigma_W) \geq F_S(\rho'_S) - F_S(\rho_S).$$

Step 3: Free-energy decomposition into diagonal and coherence parts. Since $\log \gamma_S$ is diagonal in the energy basis, $\text{Tr}(\rho \log \gamma_S) = \text{Tr}(\Delta[\rho] \log \gamma_S)$, which implies the Pythagorean identity

$$S(\rho \| \gamma_S) = S(\Delta[\rho] \| \gamma_S) + S(\rho \| \Delta[\rho]).$$

Thus

$$F_S(\rho) = F_S(\Delta[\rho]) + k_B T C(\rho).$$

Step 4: Paired-strategy extra work bound (per cycle). Fix $\Delta t > 0$ and set $\rho_{\Delta t} := e^{\Delta t \mathcal{L}}(\rho)$. Let $s \in \text{Ctrl}(\rho)$ be any strategy that restores $\rho_{\Delta t} \mapsto \rho$ by a battery-assisted thermal operation. By Step 2,

$$W_s(\Delta t) \geq F_S(\rho) - F_S(\rho_{\Delta t}).$$

Similarly, for any baseline strategy $s_{\text{diag}} \in \text{Ctrl}(\Delta[\rho])$ restoring $\Delta[\rho]_{\Delta t} \mapsto \Delta[\rho]$ with $\Delta[\rho]_{\Delta t} := e^{\Delta t \mathcal{L}}(\Delta[\rho])$,

$$W_{s_{\text{diag}}}(\Delta t) \geq F_S(\Delta[\rho]) - F_S(\Delta[\rho]_{\Delta t}).$$

Subtracting and using Step 3 yields

$$W_s(\Delta t) - W_{s_{\text{diag}}}(\Delta t) \geq k_B T (C(\rho) - C(\rho_{\Delta t})) + (F_S(\Delta[\rho]_{\Delta t}) - F_S(\rho_{\Delta t})).$$

In the paired-strategy definition of incremental cost, the baseline infimum is taken over all admissible population-maintenance strategies and thus isolates the coherence contribution in the small- Δt limit. This yields the operational per-cycle bound

$$W_{\text{extra}}(\Delta t) := \inf_{s, s_{\text{diag}}} (W_s(\Delta t) - W_{s_{\text{diag}}}(\Delta t)) \geq k_{\text{B}}T(C(\rho) - C(\rho_{\Delta t})),$$

where the infimum is over $s \in \text{Ctrl}(\rho)$ and $s_{\text{diag}} \in \text{Ctrl}(\Delta[\rho])$.

Step 5: From per-cycle work to asymptotic power and $\Delta t \rightarrow 0$. Divide by Δt and take lim sup over many cycles to obtain

$$P(s) - P(s_{\text{diag}}) \geq k_{\text{B}}T \frac{C(\rho) - C(\rho_{\Delta t})}{\Delta t}.$$

Taking the infimum over pairs and then $\Delta t \rightarrow 0$ yields

$$P_{\text{extra}}(\rho) \geq k_{\text{B}}T \dot{C}_{\text{loss}}(\rho),$$

which is Theorem 1.

B Experimental calibration and power proxies

Experimental platforms often provide wall-plug electrical power rather than battery-accounted free-energy changes. We recommend:

Relative calibration. Compare two strategies (s maintaining ρ and s_{diag} maintaining $\Delta[\rho]$) implemented with identical primitives except for the coherence-stabilizing component. Subtract their measured average powers to estimate a conservative incremental proxy

$$P_{\text{wall,extra}} := P_{\text{wall}}(s) - P_{\text{wall}}(s_{\text{diag}}).$$

Thermal anchoring. Anchor $k_{\text{B}}T$ using (i) an engineered thermal ancilla, or (ii) an effective temperature inferred from detailed-balance ratios in the uncontrolled generator.

Error budget. Aggregate uncertainties from (i) power metrology, (ii) \dot{C}_{loss} estimation (tomography noise and short-time extrapolation), and (iii) drift/non-stationarity.

Upper-bound calibration. Let $0 < \eta \leq 1$ denote the end-to-end efficiency from wall power to battery-accounted work. Then $P_{\text{extra}} \leq P_{\text{wall,extra}}$ provides a conservative estimate. Reporting sensitivity analyses (duty cycles, amplitudes) is recommended. If $P_{\text{wall,extra}} < k_{\text{B}}T \dot{C}_{\text{loss}}$ beyond uncertainties, this points to assumption failures (non-Markovianity, thermal reference mismatch, proxy miscalibration), not to falsifying the theorem.

References

- [1] W. H. Zurek, Decoherence, einselection, and the quantum origins of the classical, *Rev. Mod. Phys.* 75, 715 (2003).
- [2] M. Schlosshauer, *Decoherence and the Quantum-to-Classical Transition* (Springer, 2007).
- [3] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, 2002).
- [4] E. B. Davies, Markovian master equations, *Commun. Math. Phys.* 39, 91–110 (1974).
- [5] H. Spohn, Entropy production for quantum dynamical semigroups, *J. Math. Phys.* 19, 1227–1230 (1978).
- [6] T. Baumgratz, M. Cramer, and M. B. Plenio, Quantifying Coherence, *Phys. Rev. Lett.* 113, 140401 (2014).
- [7] M. Lostaglio, D. Jennings, and T. Rudolph, Description of quantum coherence in thermodynamic processes requires constraints beyond free energy, *Nat. Commun.* 6, 6383 (2015).
- [8] F. G. S. L. Brandão, M. Horodecki, N. Ng, J. Oppenheim, and S. Wehner, The second laws of quantum thermodynamics, *PNAS* 112, 3275 (2015).
- [9] M. B. Hastings and T. Koma, Spectral gap and exponential decay of correlations, *Commun. Math. Phys.* 265, 781–804 (2006).
- [10] B. Nachtergaele and R. Sims, Lieb–Robinson bounds and the exponential clustering theorem, *Commun. Math. Phys.* 265, 119–130 (2006).
- [11] E. H. Lieb and D. W. Robinson, The finite group velocity of quantum spin systems, *Commun. Math. Phys.* 28, 251–257 (1972).
- [12] K. Korzekwa, M. Lostaglio, J. Oppenheim, and D. Jennings, The extraction of work from quantum coherence, *New J. Phys.* 18, 023045 (2016).
- [13] P. Źwikliński, M. Studziński, M. Horodecki, and J. Oppenheim, Limitations on the evolution of quantum coherences: towards fully quantum second laws of thermodynamics, *Phys. Rev. Lett.* 115, 210403 (2015).
- [14] K. Friston, The free-energy principle: a unified brain theory?, *Nat. Rev. Neurosci.* 11, 127–138 (2010).
- [15] T. Parr and K. Friston, Generalised free energy and active inference, *Biol. Cybern.* 113, 495–513 (2019).
- [16] S. B. Laughlin, J. M. de Ruyter van Steveninck, and J. C. Anderson, The metabolic cost of neural information, *Nat. Neurosci.* 1, 36–41 (1998).
- [17] D. Attwell and S. B. Laughlin, An energy budget for signaling in the grey matter of the brain, *J. Cereb. Blood Flow Metab.* 21, 1133–1145 (2001).