

The Conditional Maintenance Work Theorem: Operational Power Lower Bounds from Energy Pinching and a Split-Inclusion Blueprint for Type III AQFT

Lluis Eriksson
Independent Researcher
lluiseriksson@gmail.com

July 4, 2026 (v2; v1: December 2025)

Abstract

We derive operational lower bounds on the minimal thermodynamic power required to maintain quantum coherence against uncontrolled open-system dynamics. Work is defined as the *consumption* of non-equilibrium free energy stored in an explicit battery at bath temperature T , namely $W := F_W(\sigma_W) - F_W(\sigma'_W)$ with $F_W(\sigma) = k_B T S(\sigma \| \gamma_W)$, and allowed controls are battery-assisted thermal operations implemented by energy-conserving global unitaries. Coherence is quantified as a relative entropy to a conditional expectation, $C^E(\rho) := S(\rho \| E[\rho])$. Our main proved results are: (i) an *unconditional* maintenance bound $P_{\min}(\rho) \geq k_B T \sigma(\rho)$, where $\sigma(\rho) \geq 0$ is the entropy production rate of the target under the uncontrolled semigroup (Spohn), which splits via the pinching Pythagorean identity as $\sigma = \dot{F}_{\text{diag}}^\Delta + \dot{C}_{\text{loss}}^\Delta$; (ii) the coherence-power bound $P_{\min} \geq k_B T \dot{C}_{\text{loss}}^\Delta(\rho)$ under a diagonal-contraction hypothesis satisfied by Davies generators; and (iii) extra-power bounds for coherence stabilization relative to *thermodynamically efficient* population-maintenance baselines — version 1’s “assumption-free” paired-infimum formulation is shown to be vacuous (the unlinked-pairs infimum is $-\infty$), and the corrected statement is genuinely assumption-free precisely when the pinched target is stationary (e.g., pure dephasing), where the baseline is free. We further provide a conditional extension to general γ_S -preserving conditional expectations and a Type III split blueprint, with external geometric inputs encoded as explicit interface assumptions. All finite-dimensional claims are verified by an exact numerical suite distributed with the paper (Appendix C). These results provide an operational resource criterion for quantum-to-classical behavior, not a collapse theory.

Version 2. Corrects a sign error in the work definition and in the final step of v1’s Appendix A, replaces v1’s vacuous extra-power definition (Remark 4.13), and strengthens the main bound to an unconditional entropy-production form. See Appendix D.

Contents

1	Introduction	2
1.1	Problem	2
1.2	Contribution and architecture	2
1.3	Editorial status and dependency map	3
2	Operational framework: system–bath–battery	3
3	Conditional expectations and coherence	4
3.1	Energy pinching	4

4	Single-law maintenance bounds	4
4.1	Uncontrolled dynamics, maintenance, and power	4
4.2	Unconditional entropy-production bound and its pinching split	5
4.3	Extra power: efficient baselines	6
5	Qubit specialization	7
6	Conditional extension to general conditional expectations	7
7	Type III split blueprint (conditional)	7
8	Geometric interface and scaling laws	8
8.1	One-shot work vs maintenance power	8
9	Interpretation and scope: an operational “cut” (not a collapse theory)	9
9.1	On a proposed “master inequality”	9
10	Protocols and reproducibility	9
10.1	Protocol sketch: incremental coherence stabilization	9
11	Conclusion	9
A	Work bound proof via product-reference decomposition	10
B	Exact qubit formula for $\dot{C}_{\text{loss}}^{\Delta}$ under pure dephasing	10
C	Numerical verification	11
D	Changes relative to version 1	11

1 Introduction

1.1 Problem

Maintaining coherence is a dynamical control task: an uncontrolled environment drives a system toward a thermal fixed point, while a controller expends work to hold a target state stationary. This paper asks for operational lower bounds on the minimal power required to maintain a target state under explicit thermodynamic control primitives.

1.2 Contribution and architecture

We use a deliberately modular structure:

- (C1) **From static work to dynamical power.** Static free-energy/work inequalities are translated into maintenance-power lower bounds via stroboscopic control cycles.
- (C2) **Entropy production as the unconditional floor.** The minimal maintenance power is bounded below by $k_{\text{B}}T$ times the entropy production rate of the target, with no structural assumption on the generator beyond a thermal fixed point (Theorem 4.8). The pinching Pythagorean identity splits this floor into a population term and a coherence term (Proposition 4.9).
- (C3) **Paired-strategy rigor.** The incremental cost of coherence stabilization is formulated against *thermodynamically efficient* baselines. We show that the naive infimum over

unlinked strategy pairs is $-\infty$ (Remark 4.13), identify the exact assumption (baseline achievability) under which the difference-of-infima bound holds, and prove that no assumption is needed when the pinched target is stationary.

- (C4) **Proved/conditional/interface separation.** Finite-dimensional results are proved; general- E and Type III statements are conditional; geometric inputs appear only as explicit interface assumptions.

1.3 Editorial status and dependency map

This manuscript is self-contained at the level of proved results.

- (S1) *Closed results proved here:* Proposition 2.8, Proposition 3.5, Theorem 4.8, Proposition 4.9, Corollary 4.12, Theorem 4.16, Proposition 4.20, and Corollary 5.2, with full appendices.
- (S2) *Conditional extensions:* Assumption 6.1 and Theorem 6.3; the Type III blueprint in Section 7; the achievability-based Corollary 4.18.
- (S3) *External geometric inputs:* encoded by Assumptions 8.1 and 8.3.

2 Operational framework: system–bath–battery

Remark 2.1 (Entropy convention). We use natural logarithms. All entropies are measured in nats.

Definition 2.2 (Thermal state). Let X be finite-dimensional with Hamiltonian H_X and inverse temperature $\beta := 1/(k_B T)$. Define $\gamma_X := e^{-\beta H_X} / \text{Tr}(e^{-\beta H_X})$.

Definition 2.3 (Relative entropy). For states ρ, σ , define $S(\rho||\sigma) := \text{Tr}[\rho(\log \rho - \log \sigma)]$ when $\text{supp}(\rho) \subseteq \text{supp}(\sigma)$ and $S(\rho||\sigma) = +\infty$ otherwise.

Definition 2.4 (Non-equilibrium free energy). For any finite-dimensional system X at temperature T , define $F_X(\sigma) := k_B T S(\sigma||\gamma_X)$.

Definition 2.5 (Battery-assisted thermal operation). 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

$$\rho_S \mapsto \text{Tr}_{BW} \left(U(\rho_S \otimes \gamma_B \otimes \sigma_W) U^\dagger \right). \quad (1)$$

Definition 2.6 (Work and power). If the battery changes from σ_W to σ'_W , define the *work consumed*

$$W := F_W(\sigma_W) - F_W(\sigma'_W), \quad (2)$$

i.e., the decrease of battery free energy, and average power over time t as $P := W/t$.

Remark 2.7 (Sign correction relative to v1). Version 1 defined $W := F_W(\sigma'_W) - F_W(\sigma_W)$ (battery *increase*) and asserted $W \geq \Delta F_S$. With that sign the inequality is false — the correct consequence of energy conservation and relative-entropy invariance is that the battery must be *drained* by at least the system’s free-energy gain; the final line of v1’s Appendix A reversed an inequality. With Definition 2.6 the statement and proof below are correct, and have been verified numerically over random energy-conserving unitaries (Appendix C).

Proposition 2.8 (Work pays for system free energy). *Let a battery-assisted thermal operation map $\rho_S \mapsto \rho'_S$ while the battery changes $\sigma_W \mapsto \sigma'_W$. Then*

$$W \geq F_S(\rho'_S) - F_S(\rho_S). \quad (3)$$

Proof. See Appendix A. □

3 Conditional expectations and coherence

Definition 3.1 (Finite-dimensional conditional expectation). Let $\mathcal{A} \subseteq B(\mathcal{H}_S)$ be a unital $*$ -subalgebra. A map $E : B(\mathcal{H}_S) \rightarrow \mathcal{A}$ is a conditional expectation if it is completely positive, trace-preserving, unital, idempotent, and satisfies the bimodule property $E(AXB) = A E(X) B$ for all $A, B \in \mathcal{A}$, $X \in B(\mathcal{H}_S)$.

Definition 3.2 (γ_S -preserving conditional expectation). A conditional expectation E is γ_S -preserving if $E^\dagger(\gamma_S) = \gamma_S$, equivalently $\text{Tr}(\gamma_S E(X)) = \text{Tr}(\gamma_S X)$ for all X .

Definition 3.3 (Relative-entropy coherence to E). Define $C^E(\rho) := S(\rho \| E[\rho])$.

Assumption 3.4 (Faithfulness). In Sections 4–5, the target state ρ is assumed faithful (full rank), so that $C^E(\rho) < \infty$.

3.1 Energy pinching

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

$$\Delta[\rho] := \sum_n \Pi_n \rho \Pi_n. \quad (4)$$

Proposition 3.5 (Pythagorean identity for pinching). Let $\gamma_S \propto e^{-\beta H_S}$. Then

$$S(\rho \| \gamma_S) = S(\Delta[\rho] \| \gamma_S) + S(\rho \| \Delta[\rho]), \quad \text{equivalently} \quad F_S(\rho) = F_S(\Delta[\rho]) + k_B T C^\Delta(\rho). \quad (5)$$

Proof. Because $\log \gamma_S = -\beta H_S - \log Z$ is diagonal in the $\{\Pi_n\}$ decomposition, $\text{Tr}(\rho \log \gamma_S) = \text{Tr}(\Delta[\rho] \log \gamma_S)$, hence $S(\rho \| \gamma_S) - S(\Delta[\rho] \| \gamma_S) = \text{Tr}(\rho \log \rho) - \text{Tr}(\Delta[\rho] \log \Delta[\rho]) = S(\rho \| \Delta[\rho])$. $\square$

4 Single-law maintenance bounds

4.1 Uncontrolled dynamics, maintenance, and power

Assumption 4.1 (Markovian setting). The uncontrolled evolution is a quantum dynamical semigroup $\rho_t = e^{t\mathcal{L}}(\rho)$ with a stationary thermal state γ_S , i.e. $\mathcal{L}(\gamma_S) = 0$.

Definition 4.2 (Coherence-loss and entropy-production rates). Let $\rho_t = e^{t\mathcal{L}}(\rho)$. Define

$$\dot{C}_{\text{loss}}^E(\rho) := -\frac{d}{dt} C^E(\rho_t) \Big|_{t=0}, \quad \sigma(\rho) := -\frac{d}{dt} S(\rho_t \| \gamma_S) \Big|_{t=0}, \quad \dot{F}_{\text{diag}}^\Delta(\rho) := -\frac{d}{dt} S(\Delta[\rho_t] \| \gamma_S) \Big|_{t=0}. \quad (6)$$

By Spohn's theorem [2], $t \mapsto S(\rho_t \| \gamma_S)$ is nonincreasing, so $\sigma(\rho) \geq 0$: it is the entropy production rate of the target under the uncontrolled dynamics.

Definition 4.3 (Maintenance task; stroboscopic strategies). A $(\delta t$ -stroboscopic) control strategy alternates free evolution for time δt with a battery-assisted thermal operation (Definition 2.5); it *maintains* ρ if the controlled reduced state returns to ρ at the end of every cycle. $\text{Ctrl}(\rho)$ denotes the set of stroboscopic maintaining strategies (all $\delta t > 0$).

Remark 4.4. Continuous-control protocols can be treated as limits of stroboscopic ones; we fix the stroboscopic model to keep the accounting elementary. This restriction was implicit in v1 (its Lemma 4.6 covers only stroboscopic cycles) and is now explicit.

Definition 4.5 (Asymptotic power of a strategy). For a strategy s with battery state $\sigma_W^s(t)$, define

$$P(s) := \limsup_{t \rightarrow \infty} \frac{1}{t} (F_W(\sigma_W^s(0)) - F_W(\sigma_W^s(t))), \quad (7)$$

the asymptotic rate of battery free-energy consumption (cf. Remark 2.7).

Definition 4.6 (Minimal maintenance power). $P_{\min}(\rho; \mathcal{L}, T) := \inf_{s \in \text{Ctrl}(\rho)} P(s)$.

Lemma 4.7 (Infinitesimal to asymptotic). *Fix $\delta t > 0$ and consider δt -stroboscopic maintenance. If each cycle consumes at least $w(\delta t)$, then any such strategy satisfies $\limsup_{t \rightarrow \infty} W(t)/t \geq w(\delta t)/\delta t$. If $w(\delta t) = a \delta t + o(\delta t)$ as $\delta t \rightarrow 0$, then $P_{\min} \geq a$.*

Proof. After N cycles, $t = N\delta t$ and $W(t) \geq Nw(\delta t)$, so $W(t)/t \geq w(\delta t)/\delta t$. Take $\limsup$ in t ; then optimize the cycle time. $\square$

4.2 Unconditional entropy-production bound and its pinching split

Theorem 4.8 (Entropy-production maintenance bound (proved, unconditional)). *Assume Assumptions 3.4 and 4.1. Then*

$$P_{\min}(\rho; \mathcal{L}, T) \geq k_B T \sigma(\rho) \geq 0. \quad (8)$$

Proof. Fix a cycle time $\delta t > 0$ and write $\rho_{\delta t} = e^{\delta t \mathcal{L}}(\rho)$. One maintenance cycle must implement $\rho_{\delta t} \mapsto \rho$ by battery-assisted thermal operations. By Proposition 2.8,

$$w(\delta t) \geq F_S(\rho) - F_S(\rho_{\delta t}) = k_B T [S(\rho \| \gamma_S) - S(\rho_{\delta t} \| \gamma_S)]. \quad (9)$$

Divide by δt , let $\delta t \rightarrow 0$, and apply Lemma 4.7. $\square$

Proposition 4.9 (Rate Pythagoras). *Under Assumptions 3.4 and 4.1,*

$$\sigma(\rho) = \dot{F}_{\text{diag}}^{\Delta}(\rho) + \dot{C}_{\text{loss}}^{\Delta}(\rho). \quad (10)$$

Proof. Differentiate the Pythagorean identity of Proposition 3.5 along ρ_t at $t = 0$ (all three terms are differentiable for faithful ρ under a semigroup). $\square$

Assumption 4.10 (Diagonal contraction). $t \mapsto S(\Delta[\rho_t] \| \gamma_S)$ is nonincreasing; equivalently $\dot{F}_{\text{diag}}^{\Delta}(e^{t\mathcal{L}}\rho) \geq 0$ for all $t \geq 0$.

Proposition 4.11 (Davies generators satisfy diagonal contraction). *If $\mathcal{L}$ is a Davies generator obtained by weak coupling and secular approximation in the energy basis [1], then populations evolve as a classical detailed-balance Markov semigroup with stationary distribution given by the diagonal of γ_S , decoupled from coherences. Consequently Assumption 4.10 holds, and moreover $\Delta \circ e^{t\mathcal{L}} = e^{t\mathcal{L}} \circ \Delta$ (population autonomy).*

Proof sketch. Under the Davies form the diagonal algebra is invariant and decouples from coherences; the induced classical semigroup satisfies detailed balance, and the classical H-theorem gives monotonicity of $S(\Delta[\rho_t] \| \gamma_S)$ [1, 2]. $\square$

Corollary 4.12 (Total-power coherence bound (v1 Theorem 4.9)). *Assume Assumptions 3.4, 4.1, and 4.10. Then*

$$P_{\min}(\rho; \mathcal{L}, T) \geq k_B T \dot{C}_{\text{loss}}^{\Delta}(\rho). \quad (11)$$

Proof. By Theorem 4.8 and Proposition 4.9, $P_{\min} \geq k_B T (\dot{F}_{\text{diag}}^{\Delta} + \dot{C}_{\text{loss}}^{\Delta})$, and $\dot{F}_{\text{diag}}^{\Delta} \geq 0$ by Assumption 4.10. $\square$

4.3 Extra power: efficient baselines

Version 1 defined the extra power of coherence as $\inf_{s \in \text{Ctrl}(\rho), s_{\text{diag}} \in \text{Ctrl}(\Delta[\rho])} [P(s) - P(s_{\text{diag}})]$ and claimed an assumption-free lower bound. That definition is vacuous:

Remark 4.13 (The unlinked-pairs infimum is $-\infty$). Fix any $s \in \text{Ctrl}(\rho)$. A population-maintenance strategy may append *battery-burning* cycles: a unitary acting only on WB that transfers battery free energy into the bath is energy-conserving, acts trivially on S , and hence preserves maintenance of $\Delta[\rho]$, while making $P(s_{\text{diag}})$ arbitrarily large (assuming, as is generic, that the bath–battery pair contains resonant sectors permitting free-energy dumping at any rate; otherwise append an auxiliary work sink with dense spectrum to the bath). Therefore $\inf_{(s, s_{\text{diag}})} [P(s) - P(s_{\text{diag}})] = -\infty$, and v1’s Theorem 4.12 cannot hold as stated: its proof subtracts *lower* bounds for both strategies, which is invalid for the subtrahend. The corrected statements below compare against baselines that are required to be thermodynamically efficient.

Assumption 4.14 (Population autonomy). $\Delta \circ e^{t\mathcal{L}} = e^{t\mathcal{L}} \circ \Delta$ for all $t \geq 0$ (automatic for Davies generators, Proposition 4.11). Then the diagonal task’s entropy production is $\sigma(\Delta[\rho]) = \dot{F}_{\text{diag}}^{\Delta}(\rho)$.

Definition 4.15 (ε -efficient baseline). For $\varepsilon \geq 0$, a strategy $s_{\text{diag}} \in \text{Ctrl}(\Delta[\rho])$ is ε -efficient if $P(s_{\text{diag}}) \leq k_{\text{B}}T \sigma(\Delta[\rho]) + \varepsilon$, i.e., its power exceeds the thermodynamic floor of its own task (Theorem 4.8) by at most ε .

Theorem 4.16 (Extra power against efficient baselines (proved)). *Assume Assumptions 3.4, 4.1, and 4.14. Then for every $s \in \text{Ctrl}(\rho)$ and every ε -efficient $s_{\text{diag}} \in \text{Ctrl}(\Delta[\rho])$,*

$$P(s) - P(s_{\text{diag}}) \geq k_{\text{B}}T \dot{C}_{\text{loss}}^{\Delta}(\rho) - \varepsilon. \quad (12)$$

Proof. By Theorem 4.8 and Proposition 4.9, $P(s) \geq k_{\text{B}}T [\dot{F}_{\text{diag}}^{\Delta}(\rho) + \dot{C}_{\text{loss}}^{\Delta}(\rho)]$. By Definition 4.15 and Assumption 4.14, $P(s_{\text{diag}}) \leq k_{\text{B}}T \sigma(\Delta[\rho]) + \varepsilon = k_{\text{B}}T \dot{F}_{\text{diag}}^{\Delta}(\rho) + \varepsilon$. Subtract. $\square$

Definition 4.17 (Extra power). $P_{\text{extra}}(\rho; \mathcal{L}, T) := P_{\min}(\rho; \mathcal{L}, T) - P_{\min}(\Delta[\rho]; \mathcal{L}, T)$, assuming both control sets are non-empty and the second infimum is finite.

Corollary 4.18 (Difference of infima (conditional on achievability)). *Assume additionally that the diagonal task is Landauer-achievable: ε -efficient baselines exist for every $\varepsilon > 0$, equivalently $P_{\min}(\Delta[\rho]) = k_{\text{B}}T \sigma(\Delta[\rho])$. Then*

$$P_{\text{extra}}(\rho; \mathcal{L}, T) \geq k_{\text{B}}T \dot{C}_{\text{loss}}^{\Delta}(\rho). \quad (13)$$

Proof. $P_{\min}(\rho) \geq k_{\text{B}}T [\dot{F}_{\text{diag}}^{\Delta} + \dot{C}_{\text{loss}}^{\Delta}]$ and $P_{\min}(\Delta[\rho]) = k_{\text{B}}T \dot{F}_{\text{diag}}^{\Delta}$ by achievability and Assumption 4.14. $\square$

Remark 4.19. Achievability of the classical (population) Landauer floor is a statement about optimal classical stochastic control of a detailed-balance Markov chain in the slow-driving limit; it is plausible but not proved here, which is why Corollary 4.18 is listed as conditional. The experimentally meaningful statement is Theorem 4.16: the measured incremental power of a coherence stabilizer over a *well-implemented* population stabilizer is at least $k_{\text{B}}T \dot{C}_{\text{loss}}^{\Delta}$.

Proposition 4.20 (Free baseline: stationary pinched target (proved, assumption-free)). *Assume Assumptions 3.4 and 4.1, and suppose both $\mathcal{L}(\Delta[\rho]) = 0$ (the pinched target is stationary) and $\Delta[\mathcal{L}(\rho)] = 0$ (the populations of the target are stationary along the flow); both hold simultaneously for pure dephasing and, more generally, in population-autonomous settings. Then the do-nothing strategy is a 0-efficient baseline, $P_{\min}(\Delta[\rho]) = 0$, $\dot{F}_{\text{diag}}^{\Delta}(\rho) = 0$, and*

$$P_{\text{extra}}(\rho; \mathcal{L}, T) = P_{\min}(\rho; \mathcal{L}, T) \geq k_{\text{B}}T \sigma(\rho) = k_{\text{B}}T \dot{C}_{\text{loss}}^{\Delta}(\rho). \quad (14)$$

Proof. If $\mathcal{L}(\Delta[\rho]) = 0$, the strategy that never touches the battery maintains $\Delta[\rho]$ with $P = 0$; by Theorem 4.8 applied to the diagonal task, $P_{\min}(\Delta[\rho]) \geq k_B T \sigma(\Delta[\rho]) = 0$, so $P_{\min}(\Delta[\rho]) = 0$. The second hypothesis gives $\frac{d}{dt}\Delta[\rho_t]|_{t=0} = \Delta[\mathcal{L}(\rho)] = 0$, hence $\dot{F}_{\text{diag}}^\Delta(\rho) = 0$ and $\sigma = \dot{C}_{\text{loss}}^\Delta$ by Proposition 4.9. $\square$

5 Qubit specialization

Assumption 5.1 (Qubit pure dephasing). Let S be a qubit in the energy basis such that under uncontrolled dynamics $\frac{d}{dt}\rho_{01}(t) = -\Gamma_\phi \rho_{01}(t)$ and populations are constant.

Corollary 5.2 (Qubit law (exact and leading order; assumption-free)). *Under Assumptions 3.4, 4.1, and 5.1, with $\rho = \begin{pmatrix} p & c \\ \bar{c} & 1-p \end{pmatrix}$, $z := 2p - 1$, $r := \sqrt{z^2 + 4|c|^2}$:*

$$P_{\text{extra}}(\rho) = P_{\min}(\rho) \geq k_B T \dot{C}_{\text{loss}}^\Delta(\rho) = k_B T \frac{2\Gamma_\phi |c|^2}{r} \log \frac{1+r}{1-r} = 2k_B T \Gamma_\phi C^\Delta(\rho) (1 + O(|c|^2)). \quad (15)$$

No diagonal-contraction or achievability assumption is needed: this is the free-baseline case of Proposition 4.20. The exact formula for $\dot{C}_{\text{loss}}^\Delta$ is derived in Appendix B and verified numerically to six digits (Appendix C).

6 Conditional extension to general conditional expectations

Assumption 6.1 (Pythagorean identity for general E). Let $E : B(\mathcal{H}_S) \rightarrow \mathcal{A}$ be a γ_S -preserving conditional expectation. Assume the Pythagorean identity holds for all faithful ρ :

$$S(\rho||\gamma_S) = S(E[\rho]||\gamma_S) + S(\rho||E[\rho]). \quad (16)$$

Remark 6.2. Assumption 6.1 is automatic for pinching $E = \Delta$ and is related to sufficiency and Petz recovery structure; see [3, 4].

Theorem 6.3 (Single-law bound (conditional general- E)). *Assume Assumptions 3.4, 4.1, and 6.1, and the E -analogue of Assumption 4.10 (contraction of $t \mapsto S(E[\rho_t]||\gamma_S)$). Then*

$$P_{\min}(\rho; \mathcal{L}, T) \geq k_B T \dot{C}_{\text{loss}}^E(\rho). \quad (17)$$

Unconditionally (without the contraction hypothesis), $P_{\min} \geq k_B T \sigma(\rho)$ with the E -split $\sigma = \dot{F}^E + \dot{C}_{\text{loss}}^E$.

Proof. Repeat the proofs of Theorem 4.8, Proposition 4.9, and Corollary 4.12, replacing Δ by E and using Assumption 6.1. $\square$

7 Type III split blueprint (conditional)

Remark 7.1 (Blueprint status). This section is a blueprint for a full Type III treatment: no new Type III technical proofs (Haagerup L^p , domain control, energy constraints) are claimed here.

Assumption 7.2 (Split inclusion and vacuum-preserving expectation). Let $\mathcal{A}(\mathcal{O}_1) \subset \mathcal{A}(\mathcal{O}_2) =: \mathcal{M}$ be a split inclusion with intermediate Type I factor $\mathcal{N}$ and vacuum state ω_0 . Assume $\mathcal{N}$ is $\sigma_t^{\omega_0}$ -invariant, hence there exists a unique ω_0 -preserving conditional expectation $E_{\omega_0} : \mathcal{M} \rightarrow \mathcal{N}$.

Definition 7.3 (Split-relative entropy functional (Type III, conditional)). Define $C^{\text{split}}(\omega) := S(\omega || \omega \circ E_{\omega_0})$, where S is Araki relative entropy.

Remark 7.4 (Relation to the companion paper). The geometric inputs of Section 8 and the existence statement in Assumption 7.2 are developed in the companion preprint [11]. Note that v2 of [11] corrects its own Proposition 2.14 (the Gaussian formula there describes a conditional *reattachment* reconstruction, not the Petz map); nothing in the present paper depends on that formula — we use only (i) the existence of the ω_0 -preserving conditional expectation on the canonical split factor (unaffected; it follows from modular invariance of the nuclearity-map construction) and (ii) the Bessel-type collar suppression of vacuum correlations (unaffected).

8 Geometric interface and scaling laws

8.1 One-shot work vs maintenance power

We distinguish *one-shot shaping* (reduce a leakage proxy at fixed collar width ϵ , yielding a bound on W) from *sustained maintenance* (keep a target stationary in time, yielding a bound on P).

Assumption 8.1 (Geometric leakage amplitude interface). In a bosonic Gaussian split regime with collar width ϵ and mass gap $m > 0$, assume there exists a bounded leakage proxy $\text{Leak}(\epsilon)$ satisfying the *two-sided* interface bound

$$a(m\epsilon)^\alpha K_\nu(m\epsilon) \leq \text{Leak}(\epsilon) \leq A(m\epsilon)^\alpha K_\nu(m\epsilon) \quad (18)$$

for constants $A \geq a > 0$, $\alpha \geq 0$, with K_ν the modified Bessel function, and weak-correlation quadraticity of mutual information, $I(S : \text{Env}) \asymp \text{Leak}(\epsilon)^2$. (The upper bound alone — all that v1 assumed — yields upper, not lower, scalings; the lower bound encodes sharpness of the geometric suppression and is what the “in particular” scalings of Corollaries 8.2 and 8.4 consume.)

Corollary 8.2 (One-shot geometric work scaling). *Under Assumption 8.1, any thermal-operation protocol at temperature T that reduces the leakage proxy from its baseline value $\text{Leak}(\epsilon)$ to a target level δ satisfies a Landauer-type lower bound of the form*

$$W \geq k_B T c_I [\text{Leak}(\epsilon)^2 - \delta^2]_+, \quad (19)$$

for a state-family constant $c_I > 0$; in particular $W \gtrsim k_B T K_\nu(m\epsilon)^2$ up to polynomial prefactors in $m\epsilon$. (In v1 the constant here was denoted κ , clashing with the relaxation rate of Assumption 8.3; renamed.)

Assumption 8.3 (Bessel-suppressed relaxation rate, two-sided). Assume the uncontrolled relaxation rate satisfies

$$b(m\epsilon)^\beta K_\nu(m\epsilon) \leq \kappa(\epsilon) \leq B(m\epsilon)^\beta K_\nu(m\epsilon) \quad (20)$$

for constants $B \geq b > 0$, $\beta \geq 0$ within the regime of interest.

Corollary 8.4 (Conditional maintenance-power scaling). *Assume the uncontrolled evolution yields $C^\Delta(\rho_t) = C^\Delta(\rho) e^{-\kappa(\epsilon)t}$ for t near 0, population autonomy in the regime of interest, and Assumption 8.3. Then, against ϵ -efficient baselines (Theorem 4.16) or in the free-baseline regime (Proposition 4.20),*

$$P_{\text{extra}}(\epsilon) \geq k_B T \kappa(\epsilon) C^\Delta(\rho) \quad \text{and in particular} \quad P_{\text{extra}}(\epsilon) \gtrsim k_B T K_\nu(m\epsilon) \quad (21)$$

up to polynomial prefactors in $m\epsilon$.

9 Interpretation and scope: an operational “cut” (not a collapse theory)

Remark 9.1 (What this paper does and does not claim). This work does not propose a collapse mechanism, does not derive the Born rule, and does not claim to solve the measurement problem in the strong foundational sense. The results are operational: they quantify minimal work/power requirements for maintaining a target state against uncontrolled dynamics under explicit thermodynamic control restrictions.

Remark 9.2 (Quantum-to-classical transition as resource infeasibility). Theorem 4.8, Corollary 4.12, and Theorem 4.16 imply that sustaining coherence against uncontrolled noise requires a minimal power budget proportional to $k_B T$ times an intrinsic dissipation rate. When this required power exceeds available control resources, coherence cannot be maintained in steady state. This is an operational statement about control feasibility, not an ontological account of definite outcomes. The resource-boundary reading of the Heisenberg cut built on these bounds — including an exact rate-inheritance law for gapped buffers, $\kappa(\epsilon) \propto e^{-2q(\omega_b)\epsilon}$, verified against exact Liouvillian spectra — is developed in the companion paper [12].

9.1 On a proposed “master inequality”

Remark 9.3 (Work vs power, amplitude vs rate). This paper does not prove a universal inequality of the literal form $I \lesssim E e^{-2m\epsilon}$. What it supports, under explicit assumptions, is two distinct statements: one-shot work scaling $W \gtrsim k_B T K_\nu(m\epsilon)^2 \sim k_B T e^{-2m\epsilon}$ (Corollary 8.2), and maintenance-power scaling $P_{\text{extra}} \gtrsim k_B T K_\nu(m\epsilon) \sim k_B T e^{-m\epsilon}$ (Corollary 8.4), each with its own hypotheses.

10 Protocols and reproducibility

Remark 10.1. This paper provides theoretical bounds and reproducibility templates; it does not claim new direct measurements of battery-accounted work in hardware.

10.1 Protocol sketch: incremental coherence stabilization

- (P1) Stabilize populations with a *well-implemented* baseline: an ε -efficient strategy in $\text{Ctrl}(\Delta[\rho])$ (Definition 4.15); calibrate its power against the classical floor $k_B T \sigma(\Delta[\rho])$. In dephasing-dominated platforms the baseline is free (Proposition 4.20).
- (P2) Turn on phase/coherence stabilization to implement a strategy in $\text{Ctrl}(\rho)$.
- (P3) Measure the incremental controller power $P(s) - P(s_{\text{diag}})$ and compare against $k_B T \dot{C}_{\text{loss}}^\Delta(\rho) - \varepsilon$ (Theorem 4.16).

11 Conclusion

We derived single-law operational lower bounds on minimal coherence-maintenance power under explicit thermodynamic control constraints. The unconditional floor is $k_B T$ times the entropy production rate of the target, which the pinching Pythagorean identity splits into population and coherence contributions; the coherence term alone bounds the total power under diagonal contraction (automatic for Davies generators), and bounds the *incremental* power against thermodynamically efficient population-maintenance baselines — with no assumptions at all when the pinched target is stationary, as in pure dephasing. Conditional extensions to

general expectations and Type III split configurations are provided as blueprints, and geometric inputs are incorporated only via explicit interface assumptions separating one-shot work from maintenance power. All finite-dimensional claims are verified by an exact numerical suite distributed with the paper.

A Work bound proof via product-reference decomposition

Definition A.1 (Mutual information and multi-information). For a state ρ_{XY} define $I_\rho(X : Y) := S(\rho_{XY} \| \rho_X \otimes \rho_Y) \geq 0$; for ρ_{XYZ} define $I_\rho(X : Y : Z) := S(\rho_{XYZ} \| \rho_X \otimes \rho_Y \otimes \rho_Z) \geq 0$.

Lemma A.2 (Relative entropy decomposition for product references). *Let σ_X, σ_Y be states and ρ_{XY} any joint state. Then*

$$S(\rho_{XY} \| \sigma_X \otimes \sigma_Y) = S(\rho_X \| \sigma_X) + S(\rho_Y \| \sigma_Y) + I_\rho(X : Y), \quad (22)$$

and analogously for three parties.

Proof. Expand definitions and use $\text{Tr}(\rho_{XY} \log \sigma_X) = \text{Tr}(\rho_X \log \sigma_X)$, etc., then add and subtract the local entropy terms to identify the mutual information. $\square$

Proof of Proposition 2.8. Let the initial global state be $\rho_{SBW} := \rho_S \otimes \gamma_B \otimes \sigma_W$ and the final state $\rho'_{SBW} := U \rho_{SBW} U^\dagger$, with $[U, H_S + H_B + H_W] = 0$. Let $\gamma_{SBW} := \gamma_S \otimes \gamma_B \otimes \gamma_W$. Invariance of relative entropy under unitaries and $U \gamma_{SBW} U^\dagger = \gamma_{SBW}$ give

$$S(\rho'_{SBW} \| \gamma_{SBW}) = S(\rho_{SBW} \| \gamma_{SBW}). \quad (23)$$

By Lemma A.2 applied to the final state,

$$S(\rho'_{SBW} \| \gamma_{SBW}) \geq S(\rho'_S \| \gamma_S) + S(\rho'_W \| \gamma_W), \quad (24)$$

discarding $S(\rho'_B \| \gamma_B) \geq 0$ and $I_{\rho'} \geq 0$; and applied to the initial product state,

$$S(\rho_{SBW} \| \gamma_{SBW}) = S(\rho_S \| \gamma_S) + S(\sigma_W \| \gamma_W), \quad (25)$$

since $S(\gamma_B \| \gamma_B) = 0$ and the initial multi-information vanishes. Combining,

$$S(\rho'_S \| \gamma_S) + S(\rho'_W \| \gamma_W) \leq S(\rho_S \| \gamma_S) + S(\sigma_W \| \gamma_W), \quad (26)$$

hence

$$S(\sigma_W \| \gamma_W) - S(\rho'_W \| \gamma_W) \geq S(\rho'_S \| \gamma_S) - S(\rho_S \| \gamma_S). \quad (27)$$

Multiplying by $k_B T$ and using Definition 2.6 yields $W = F_W(\sigma_W) - F_W(\rho'_W) \geq F_S(\rho'_S) - F_S(\rho_S)$. (v1 stated the last display with both signs flipped on the left; see Remark 2.7.) $\square$

B Exact qubit formula for $\dot{C}_{\text{loss}}^\Delta$ under pure dephasing

Let

$$\rho = \begin{pmatrix} p & c \\ \bar{c} & 1-p \end{pmatrix}, \quad \Delta[\rho] = \begin{pmatrix} p & 0 \\ 0 & 1-p \end{pmatrix}. \quad (28)$$

Let $z := 2p - 1$ and $r := \sqrt{z^2 + 4|c|^2}$. Under pure dephasing with rate Γ_ϕ , p is constant and $|c(t)| = |c|e^{-\Gamma_\phi t}$. Then

$$\dot{C}_{\text{loss}}^\Delta(\rho) = \frac{2\Gamma_\phi |c|^2}{r} \log \frac{1+r}{1-r}. \quad (29)$$

For $|c| \ll 1$ at fixed p , $\dot{C}_{\text{loss}}^\Delta(\rho) = 2\Gamma_\phi C^\Delta(\rho) + O(|c|^4)$.

Derivation. $C^\Delta(\rho) = H(p) - H(\frac{1+r}{2})$ with H the binary entropy in nats; differentiating through $r(t) = \sqrt{z^2 + 4|c|^2}e^{-2\Gamma_\phi t}$ at $t = 0$ gives the display. Verified numerically to six digits (Appendix C).

C Numerical verification

An exact finite-dimensional suite (`verification/verify_0061.py` in the companion repository) checks:

1. **Proposition 2.8 and the sign correction.** Over 200 random energy-conserving unitaries on a qubit–qutrit–ququart SBW system (block-Haar within total-energy eigenspaces, random faithful ρ_S and σ_W): $\min(W - \Delta F_S) = +0.48 > 0$ with the v2 sign (consumption), while with v1’s sign the difference reaches -1.12 , exhibiting the error of Remark 2.7.
2. **Appendix B formula.** Central-difference derivative of $C^\Delta(\rho_t)$ matches the closed formula to six digits at $(p, c) = (0.3, 0.2), (0.5, 0.1), (0.62, 0.31)$, and the leading-order form $2\Gamma_\phi C^\Delta$ within the expected $O(|c|^4)$ deviation.
3. **Rate Pythagoras, Spohn, Davies.** For a random qutrit Davies-type generator with detailed-balance jump rates and energy-basis dephasing ($\|\mathcal{L}(\gamma_S)\| \sim 10^{-16}$): $\sigma = \dot{F}_{\text{diag}}^\Delta + \dot{C}_{\text{loss}}^\Delta$ holds to 10^{-13} ; $\sigma \geq 0$; $\dot{F}_{\text{diag}}^\Delta \geq 0$; and $t \mapsto S(\Delta[\rho_t]||\gamma_S)$ is monotone along the flow (Assumption 4.10).
4. **Vacuity of the v1 pair-infimum** (Remark 4.13): battery-burning baselines make $P(s_{\text{diag}})$ arbitrary, so the unlinked-pairs infimum is $-\infty$.

D Changes relative to version 1

1. **Work sign corrected.** v1’s Definition 2.6 defined W as the battery free-energy *increase* and its Appendix A reversed the final inequality to conclude $W \geq \Delta F_S$; with that sign the claim is false (numerically: $W - \Delta F_S$ reaches -1.12 on random energy-conserving unitaries). v2 defines work as battery free-energy *consumption* (Definition 2.6, Remark 2.7) and proves Proposition 2.8 with the correct sign; all downstream power definitions are consumption rates.
2. **v1 Definition 4.11 was vacuous.** The infimum over unlinked strategy pairs equals $-\infty$ (battery-burning baselines; Remark 4.13), so v1’s “assumption-free” Theorem 4.12 could not hold as stated; its proof subtracted lower bounds in the invalid direction. v2 replaces it with: (a) Theorem 4.16 against ε -efficient baselines (proved, under population autonomy); (b) Corollary 4.18 for the difference of infima under an explicit Landauer-achievability assumption; and (c) Proposition 4.20, which is genuinely assumption-free when the pinched target is stationary (pure dephasing), the case relevant to Corollary 5.2.
3. **Unconditional strengthening.** New Theorem 4.8: $P_{\min} \geq k_B T \sigma(\rho)$ with σ the entropy production rate (Spohn [2]), plus the rate-Pythagoras split (Proposition 4.9); v1’s Theorem 4.9 becomes Corollary 4.12.
4. **Stroboscopic model made explicit** (Definition 4.3); v1’s Lemma 4.6 implicitly assumed it.
5. **Notation fix** in Section 8: the constant of Corollary 8.2 is now c_I (v1 used κ , clashing with the relaxation rate $\kappa(\epsilon)$ of Assumption 8.3).
6. **Companion reference updated** to v2 of [11], with Remark 7.4 clarifying that the Petz-related correction there does not affect anything used here.
7. **Numerical verification** added (Appendix C) with scripts distributed alongside the paper.

References

- [1] E. B. Davies, *Markovian master equations*, Commun. Math. Phys. **39** (1974), 91–110.
- [2] H. Spohn, *Entropy production for quantum dynamical semigroups*, J. Math. Phys. **19** (1978), 1227–1230.
- [3] D. Petz, *Sufficient subalgebras and relative entropy*, Commun. Math. Phys. **105** (1986), 123–131.
- [4] M. Takesaki, *Theory of Operator Algebras I–III*, Springer, 2003.
- [5] T. Baumgratz, M. Cramer, and M. B. Plenio, *Quantifying coherence*, Phys. Rev. Lett. **113** (2014), 140401.
- [6] M. Lostaglio, D. Jennings, and T. Rudolph, *Description of quantum coherence in thermodynamic processes requires constraints beyond free energy*, Nat. Commun. **6** (2015), 6383.
- [7] F. G. S. L. Brandão, M. Horodecki, N. Ng, J. Oppenheim, and S. Wehner, *The second laws of quantum thermodynamics*, Proc. Natl. Acad. Sci. **112** (2015), 3275–3281.
- [8] S. Doplicher and R. Longo, *Standard and split inclusions of von Neumann algebras*, Invent. Math. **75** (1984), 493–536.
- [9] D. Buchholz and E. H. Wichmann, *Causal independence and the energy-level density of states in local quantum field theory*, Commun. Math. Phys. **106** (1986), 321–344.
- [10] D. Buchholz, C. D’Antoni, and R. Longo, *Nuclear maps and modular structures I: General properties*, J. Funct. Anal. **88** (1990), 233–250.
- [11] L. Eriksson, *Clustering, Recovery, and Locality in Algebraic Quantum Field Theory: Quantitative Bounds via Split Inclusions and Modular Theory*, ai.viXra:2512.0060, v2 (2026; v1: 2025). Source, PDF, and exact verification suite in the companion repository ([papers/0060-clustering-recovery](#)).
- [12] L. Eriksson, *The Heisenberg Cut as a Resource Boundary: An Operational Outlook from Coherence Maintenance Costs*, ai.viXra:2512.0064, v2 (2026; v1: 2025). Source, PDF, and verification suite in the companion repository ([papers/0064-heisenberg-cut](#)).