

Maintenance Is Not Restoration

Endpoint No-Go Theorems, Exact SSH Resource Horizons, and Quantum Rapid-Replacement Equality

Working paper - major revision, 12 July 2026

Lluis Eriksson - Independent Researcher

Abstract

Lower bounds based on the instantaneous free-energy loss of a target state are often interpreted as lower bounds on the power required to maintain that state. We show that this interpretation depends decisively on the control task. If maintenance only requires exact restoration at periodically sampled endpoints and the intervention period is unrestricted, the infimal average power can vanish even when the target has strictly positive instantaneous entropy production. A full-rank dephasing qubit gives an exact energy-conserving SWAP counterexample. We repair the formulation by separating endpoint restoration, bounded-latency maintenance, and continuous holding.

For continuous-time finite Markov chains we adopt the established trajectory-relative-entropy holding cost and review the known reversible optimizer and Dirichlet-form representation. Our first composition result is then an exact geometric additivity theorem for suppressible and persistent generators sharing a detailed-balance reference. A two-state counterexample shows that this hypothesis is essential: channels with incompatible equilibria can cancel at a finite membrane width. A dual-rail pair of odd fermionic SSH chains supplies the microscopic layer: exact edge modes, a uniform bulk gap, and an explicitly filtered number-conserving Davies coupling produce two-sided holding-cost bounds without an assumed rate-inheritance bridge. The logical basis has fixed particle number and parity, so arbitrary logical coherence is physical under fermionic superselection. Finally, under a fully axiomatized resource-cell ledger, fresh target-state cells and energy-conserving SWAPs saturate the fixed-period free-energy bound; the correctly ordered rapid-control limit closes the SSH theorem for coherent logical targets. The fresh-copy construction is related to earlier collision-model stabilization work and is not claimed as a work-only controller. The remaining frontier is autonomous work-only control without preloaded target copies.

1. Three inequivalent control tasks

Let $(T_t)_{t \geq 0}$ be an uncontrolled Markov semigroup and let ρ_* be a target state.

- **Endpoint restoration with period τ :** the controlled state equals ρ_* at times $n \tau$, with no constraint between endpoints.
- **Bounded-latency maintenance:** endpoint restoration is required and every waiting interval is at most $\tau_{\max}$.
- **Continuous holding:** the controlled dynamics keeps the target invariant, or keeps the state inside a prescribed corridor at all times.

Only the third task literally holds a state continuously. The first can spend almost all of each cycle far from the target.

Work must be recorded by an additive ledger $A(t)$ or by a thermodynamic limit of fresh work cells. A fixed finite-dimensional battery has bounded free energy, so the expression

$$\limsup_{t \rightarrow \infty} [F_W(0) - F_W(t)]/t$$

is zero and cannot represent sustained nonzero power.

1.1 Relation to prior work and novelty boundary

The KL-rate minimization under a stationary-distribution constraint, its large-deviation interpretation, and explicit reversible/two-state optimizers are established results, particularly in Pavlichin, Quek, and Weissman. Theorem 5.1 below is included for self-contained composition and is **not claimed as new**.

Fresh-copy collision models also precede this work. Vacanti, Elouard, and Auffeves distinguish restoring from stabilizing maps, use target-state copies and SWAPs, and prove that stationary resource states plus energy-conserving interactions cannot stabilize energy-basis coherence. They then price the extra external work required when energy conservation of the interaction is relaxed. Our Theorem 11.2 instead allows coherent, nonstationary resource cells with the same Hamiltonian as the system;

the SWAP is energy conserving, but the cells themselves supply time-translation asymmetry. This is a different resource theory, not a contradiction or replacement of their result.

The claims specific to the present paper are:

- the zero-power degeneracy of endpoint-only maintenance when the intervention period is optimized without a latency/corridor constraint;
- the corrected order fixed-period optimization $\rightarrow$ rapid-control limit and the associated additive resource ledger;
- the common-reference composition rule interpreted as a geometric floor/ceiling theorem, together with the incompatible-equilibrium cancellation counterexample;
- the fixed-parity dual-rail SSH microscopic closure with width-uniform gap, exact remote transition weight, and a spectrally filtered logical Davies generator;
- the end-to-end composition of that microscopic model with both classical KL holding and the explicitly scoped quantum resource-cell protocol.

2. Endpoint restoration has a zero-power degeneracy

Theorem 2.1 - Infrequent-restoration degeneracy

Let $\rho_t = T_t(\rho_*)$. Suppose that for every $t > 0$ there exists an admissible exact reset $R_t: \rho_t \rightarrow \rho_*$ whose work cost $W(t)$ satisfies

$$\sup_{t > 0} W(t) < \infty.$$

If the endpoint-restoration task allows arbitrary periods, then its infimal average power is zero.

Proof

Use the reset R_τ once every period τ . Its average power is $W(\tau)/\tau$. Hence

$$0 \leq P_{\text{endpoint}} \leq \inf_{\tau > 0} W(\tau)/\tau = 0,$$

because $W(\tau)$ is uniformly bounded. QED.

Corollary 2.2 - What the free-energy inequality actually proves

If every reset satisfies

$$W(\tau) \geq F(\rho_*) - F(T_\tau \rho_*),$$

then for a *fixed* period

$$P_\tau \geq [F(\rho_*) - F(T_\tau \rho_*)]/\tau.$$

Consequently,

$$\liminf_{\tau \downarrow 0} P_\tau \geq -d/dt F(T_t \rho_*)|_{t=0}.$$

The instantaneous entropy-production rate is therefore a rapid-control bound. It is not a lower bound on the infimum over unrestricted periods.

3. Exact dephasing-qubit counterexample

Take a degenerate qubit Hamiltonian $H=0$, temperature T , and the faithful target

$$\rho_c = \begin{bmatrix} 1/2 & c \\ c & 1/2 \end{bmatrix}, \text{ with } 0 < c < 1/2.$$

Under pure dephasing, $c(t) = c \exp(-\Gamma t)$. Prepare a fresh, identical work cell in ρ_c . Since both Hamiltonians vanish, SWAP is energy conserving. Swapping the system and cell after time τ restores the system exactly and leaves the cell in $\rho_{\{c \exp(-\Gamma \tau)\}}$. Its free-energy loss is

$$W(\tau) = k_B T [D(\rho_c || I/2) - D(\rho_{\{c \exp(-\Gamma \tau)\}} || I/2)].$$

Thus

$$P_{\text{swap}}(\tau) = W(\tau)/\tau \rightarrow 0$$

as $\tau \rightarrow \infty$. In contrast, the instantaneous rate is

$$k_B T \sigma(\rho_c) = k_B T \Gamma c \log[(1/2+c)/(1/2-c)] > 0.$$

The counterexample does not contradict a continuous-holding theorem. It shows that endpoint restoration and continuous maintenance cannot share the same power definition.

4. Continuous holding of a Markov distribution

Let $Q=(q_{ij})$ be an irreducible continuous-time Markov generator on a finite state space. Let μ be a faithful target distribution. A controlled generator $P=(p_{ij})$ holds μ when $\mu P=0$.

We use the trajectory-relative-entropy rate

$$K_\mu(P||Q) = \sum_i \mu_i \sum_{j \neq i} [p_{ij} \log(p_{ij}/q_{ij}) - p_{ij} + q_{ij}].$$

This is the standard continuous-time KL control cost and is a thermodynamic lower bound in the stochastic-control realization.

Define

$$C_Q(\mu) = \inf\{K_\mu(P||Q) : \mu P=0\}.$$

5. Known reversible optimizer, reviewed for composition

This section records the continuous-time reversible result used later. It is a specialization of the minimum-holding-cost theory of Pavlichin, Quek, and Weissman and of the reversible Donsker-Varadhan formula.

Theorem 5.1 - Reversible holding-cost formula (known input)

Assume that Q is reversible with respect to a faithful distribution π :

$$\pi_i q_{ij} = \pi_j q_{ji}.$$

Set

$$h_i = \sqrt{\mu_i / \pi_i}.$$

Then the unique optimal off-diagonal controlled rates are

$$p_{ij}^* = q_{ij} h_j / h_i,$$

and

$$C_Q(\mu) = (1/2) \sum_{i,j} \pi_i q_{ij} (h_i - h_j)^2.$$

In particular, $C_Q(\mu)=0$ if and only if h is constant on every communicating component of Q .

Proof

First, detailed balance implies

$$\mu_i p_{ij}^* = \mu_j p_{ji}^*.$$

Therefore $\mu P^*=0$.

The objective is strictly convex in all positive off-diagonal rates, and the stationarity constraints are linear. Introduce Lagrange multipliers λ_i . The stationarity equation for p_{ij} is

$$\log(p_{ij}/q_{ij}) + \lambda_j - \lambda_i = 0.$$

Choosing $\lambda_i = -\log h_i$ gives $p_{ij} = q_{ij} h_j / h_i$, so the feasible candidate satisfies the necessary and sufficient convex optimality conditions.

At this candidate the logarithmic contributions cancel in reversible pairs because the controlled stationary flows are symmetric. The remaining terms are

$$\sum_i \mu_i \sum_{j \neq i} (q_{ij} - p_{ij}^*),$$

which expand, using $\mu_i = \pi_i h_i^2$ and detailed balance, to

$$(1/2) \sum_{i,j} p_{i,i} q_{i,j} (h_i - h_j)^2.$$

The zero-cost characterization follows from the sum-of-squares representation. QED.

6. Exact geometric resource-horizon theorem

Let Q describe a leakage mechanism suppressed by membrane width e_{ll} , let R describe a persistent channel, and assume both are reversible with respect to the same faithful reference distribution p_i . Define

$$L_{e_{ll}} = a_{e_{ll}} Q + g R, \text{ with } a_{e_{ll}} \geq 0 \text{ and } g \geq 0.$$

Theorem 6.1 - Exact additivity and achievability

For every faithful target μ ,

$$C_{\{L_{e_{ll}}\}}(\mu) = a_{e_{ll}} C_Q(\mu) + g C_R(\mu).$$

The infimum is attained by

$$p_{ij}^{*}(e_{ll}) = [a_{e_{ll}} q_{ij} + g r_{ij}] h_j / h_i.$$

Proof

The sum $L_{e_{ll}}$ remains reversible with respect to p_i . Apply Theorem 5.1. Its Dirichlet form is linear in the generator, giving the equality. The displayed Doob controller is the optimizer supplied by that theorem. QED.

Corollary 6.2 - Bidirectional membrane scaling

If

$$c_- \exp(-m e_{ll}) \leq a_{e_{ll}} \leq c_+ \exp(-m e_{ll}),$$

then

$g C_R(\mu) + c_- C_Q(\mu) \exp(-m e_{ll})$
$\leq C_{\{L_{e_{ll}}\}}(\mu) \leq$
$g C_R(\mu) + c_+ C_Q(\mu) \exp(-m e_{ll}).$

Hence:

- If $g=0$, the optimal achievable holding cost decays exponentially.
- If $g>0$ and $C_R(\mu)>0$, the exact resource floor is $g C_R(\mu)$.
- If $C_R(\mu)=0$, the persistent channel lies in the target's fixed-point sector and creates no floor.

No separate achievability assumption is required.

7. Why a channel witness is insufficient

The common-reference hypothesis is essential. Consider two states with target $\mu=(\alpha, 1-\alpha)$ and rates

$$q_{12}(e_{ll})=a_{e_{ll}} u+g r,$$

$$q_{21}(e_{ll})=a_{e_{ll}} v+g s.$$

Direct minimization gives

$$C_{e_{ll}}(\mu)=[\sqrt{\alpha q_{12}(e_{ll})}-\sqrt{(1-\alpha)q_{21}(e_{ll})}]^2.$$

The cost vanishes whenever

$$a_{e_{ll}}[\alpha u-(1-\alpha)v] + g[\alpha r-(1-\alpha)s]=0.$$

If the two brackets have opposite signs, a finite cancellation width exists. Thus two individually dissipative mechanisms may jointly make the target stationary. A persistent local witness does not imply a uniform family-wide floor. What matters is the fixed-point structure of the *combined* generator.

For the concrete values

$\alpha=0.7$, $(u,v)=(4,1)$, $(r,s)=(0.2,3)$, $g=1$,

the cancellation occurs at $a_{\text{ell}}=0.304$, where the exact cost is zero, while the asymptotic persistent-channel cost is approximately 0.330070426 .

8. Microscopic closure with a fixed-parity dual-rail SSH membrane

We now derive the geometric attenuation factor without using an unphysical coherence between even and odd fermion parity. Take two identical odd SSH rails $r=0,1$, each with sites

$A_{\{r,0\}}, \dots, A_{\{r,\text{ell}\}}$ and $B_{\{r,0\}}, \dots, B_{\{r,\text{ell}-1\}}$,

and canonical anticommutation relations across the combined fermionic Fock space. Let

$H_{\text{ell}} = \sum_{r=0,1} \{ \Delta_r N_r + \sum_{j=0}^{\text{ell}-1} [t_1 a_{\{r,j\}}^\dagger b_{\{r,j\}} + t_2 a_{\{r,j+1\}}^\dagger b_{\{r,j\}} + \text{h.c.}] \}$,

where $0 < t_1 < t_2$, $\zeta = t_1/t_2 < 1$, each $\Delta_r > t_1 + t_2$, and $\delta = \Delta_1 - \Delta_0 > 0$. The logical code is the one-particle, no-bulk sector

$H_{\text{log}} = \text{span}\{|0_L\rangle, |1_L\rangle\}$,

$|0_L\rangle = f_{\{0,0\}}^\dagger |\text{vac}\rangle$, $|1_L\rangle = f_{\{0,1\}}^\dagger |\text{vac}\rangle$.

Both codewords have total particle number one and odd total parity. Therefore every density matrix on H_{log} , including off-diagonal coherence, commutes with total parity and is an admissible fermionic state. This fixed-parity encoding is essential: the discarded single-mode encoding $\text{span}\{|\text{vac}\rangle, f_{\{0\}}^\dagger |\text{vac}\rangle\}$ would allow forbidden even-odd coherences.

Proposition 8.1 - Exact logical modes and uniform bulk gap

Each rail has an exact eigenmode at energy Δ_r ,

$f_{\{0,r\}}^\dagger = \sum_{j=0}^{\text{ell}} \psi_j a_{\{r,j\}}^\dagger$,

with

$\psi_j = N_{\text{ell}}(-\zeta)^j$,

$N_{\text{ell}}^2 = (1 - \zeta^2) / (1 - \zeta^{2\text{ell}+2})$.

All other single-particle energies on rail r differ from Δ_r by at least $t_2 - t_1$, uniformly in ell .

Proof

At energy Δ_r , the equations on the $B_{\{r,j\}}$ sites reduce to

$t_1 \psi_j + t_2 \psi_{j+1} = 0$,

which gives the displayed geometric mode. Normalizing the finite geometric series gives N_{ell} .

Writing either rail's chiral hopping matrix in block form, its nonzero singular values are the square roots of the eigenvalues of the $\text{ell} \times \text{ell}$ tridiagonal matrix with diagonal $t_1^2 + t_2^2$ and off-diagonal $t_1 t_2$. They are

$s_k^2 = t_1^2 + t_2^2 + 2t_1 t_2 \cos[k\pi/(\text{ell}+1)]$,

for $k=1, \dots, \text{ell}$. Hence $s_k > t_2 - t_1$. QED.

Corollary 8.2 - Exact dual-rail boundary attenuation

The single-rail zero-mode weights are

$w_{\text{near}}(\text{ell}) = (1 - \zeta^2) / (1 - \zeta^{2\text{ell}+2})$,

$w_{\text{far}}(\text{ell}) = (1 - \zeta^2) \zeta^{2\text{ell}} / (1 - \zeta^{2\text{ell}+2})$.

A local transfer between the two rails has logical matrix element $\text{conj}(\psi_x) \psi_x = w_x$, so its Davies transition weight is

$v_x(\text{ell}) = w_x(\text{ell})^2$.

Consequently,

$(1 - \zeta^2)^2 \leq v_{\text{near}}(\text{ell}) \leq 1$,

$$(1-\text{zeta}^2)^2 \text{zeta}^{4\text{ell}} \leq v_{\text{far}}(\text{ell}) \leq \text{zeta}^{4\text{ell}}.$$

Thus fixed-parity dual-rail encoding preserves explicit volume-independent two-sided bounds, with exponent 4ell rather than the single-mode exponent 2ell .

9. Number-conserving Davies reduction to the logical process

At corresponding boundary sites x of the two rails, define the even, number-conserving transfer operator

$$S_x = a_{\{1,x\}}^\dagger a_{\{0,x\}}$$

and couple it to an independent thermal reservoir through

$$V_x = \lambda (S_x \otimes B_x^\dagger + S_x^\dagger \otimes B_x).$$

The reservoir field B_x may be bosonic and is even under fermion parity. Since S_x is a product of two odd fermionic fields, V_x commutes with total fermion parity; moreover $[S_x, N_0 + N_1] = 0$. The interaction therefore never leaves the fixed-number, fixed-parity superselection sector by particle exchange.

In the normal-mode basis, the component acting inside the code is

$$P_{\text{log}} S_x P_{\text{log}} = \text{conj}(\psi_x) \psi_x \quad |1_L\rangle \langle 0_L|.$$

Transfers from a logical zero mode to a bulk mode have frequencies $\delta \pm s_k$, separated from the logical splitting δ by more than $t_2 - t_1$. Choose $0 < \delta_f < (t_2 - t_1)/2$ and assume the bath correlation spectrum vanishes at every zero-mode-to-bulk Bohr frequency while remaining positive in the KMS-related windows around $\pm \delta$. Bulk-to-bulk terms that happen to share those windows annihilate the no-bulk code sector. Thus every retained Fourier component A_ω obeys $(1 - P_{\text{log}}) A_\omega P_{\text{log}} = 0$, and the same holds for its adjoint; this is the standard reducing-subspace condition for the dissipator. The Davies restriction to H_{log} therefore has only the jump operators $|1_L\rangle \langle 0_L|$ and $|0_L\rangle \langle 1_L|$; the Lamb shift is diagonal on the code.

Let γ_+ and γ_- be the logical excitation and relaxation coefficients, satisfying

$$\gamma_+ / \gamma_- = \exp(-\beta \delta).$$

There is no missing chemical-potential term. The microscopic transfer conserves total fermion number, and both codewords have $N=1$; hence a common grand-canonical reference $\exp[-\beta(H - \mu N)]$ gives the same ratio because $N_{\{|1_L\rangle\}} - N_{\{|0_L\rangle\}} = 0$ and the μ contribution cancels. Equivalently one may set $\mu=0$ and use the canonical bath throughout.

Proposition 9.1 - Exact inherited dual-rail rates

On the logical basis $\{|0_L\rangle, |1_L\rangle\}$, a bath coupled at x induces

$$q_{\{0 \rightarrow 1\}} = v_x \gamma_+,$$

$$q_{\{1 \rightarrow 0\}} = v_x \gamma_-,$$

where $v_x = w_x^2$.

Proof

The logical transition matrix element of S_x is $\text{conj}(\psi_x) \psi_x = w_x$. Davies rates are the bath coefficient times its squared modulus, namely $v_x = w_x^2$. The uniform zero-to-bulk frequency separation and filter assumption make H_{log} invariant. Both jumps exchange the occupied rail while preserving total number and parity. QED.

9.1 Order of limits and scope of exactness

Fix $t_1, t_2, \Delta_0, \Delta_1$, the bath spectra, and a finite ell . The Davies generator is obtained in the standard weak-coupling limit $\lambda \rightarrow 0$ with rescaled time t/λ^2 . The uniform zero-to-bulk gap permits the same filter width δ_f for every ell . The theorem is exact for the restricted family of Davies generators. We do not claim that the microscopic weak-coupling convergence estimate is uniform in ell , nor do we interchange $\text{ell} \rightarrow \infty$ with $\lambda \rightarrow 0$.

The filtered generator need not be primitive on the full Fock space because other number and bulk sectors can be dynamically disconnected. Every subsequent SSH claim is made on the invariant code H_{log} ; there, positive γ_+ and γ_- give a primitive two-level Davies process with Gibbs ratio $\exp(-\beta \delta)$.

9.2 Imperfect filters and residual floors

The exact theorem assumes zero bath weight at every unwanted Bohr frequency. In this subsection and in Section 10, L_{ell} denotes the dissipative Davies generator in the interaction picture with respect to the bare logical Hamiltonian. Two distinct imperfections should be separated.

First, suppose the effective logical generator acquires an additional KMS-symmetric component $\epsilon_{\text{ell}} R$ sharing the logical Gibbs state:

$$L_{\text{ell}}^{\text{imp}} = L_{\text{ell}}^{\text{ideal}} + \epsilon_{\text{ell}} R.$$

The reversible/rapid-control formulas are generator-linear, so the holding cost or fast resource-cell power gains the exact nonnegative term $\epsilon_{\text{ell}} C_R$ or $k_B T \epsilon_{\text{ell}} \sigma_R(\rho)$. If $\epsilon_{\text{ell}} = \epsilon_{\text{ell},0} > 0$, imperfect filtering creates an additional resource floor; if $\epsilon_{\text{ell}} \rightarrow 0$, it contributes a controlled correction.

Second, genuine spectral weight at bulk frequencies breaks exact invariance of H_{log} . Write the unwanted GKLS part as

$$E_{\text{tail}} = \sum_j \gamma_j D[L_j], \quad \nu_{\text{tail}} = \sum_j \gamma_j \|L_j\|^2.$$

Since $\|D[L_j](X)\|_1 \leq 2\|L_j\|^2\|X\|_1$, Duhamel's formula gives, for finite times,

$$\|\exp[t(L_{\text{ideal}} + E_{\text{tail}})] - \exp(t L_{\text{ideal}})\|_1 \leq 2 \nu_{\text{tail}} t + O((\nu_{\text{tail}} t)^2).$$

For a faithful full-Fock target, the corresponding entropy-production correction obeys

$$|\sigma_{\text{tail}}(\rho)| \leq 2 \nu_{\text{tail}} \|\log \rho - \log \gamma\|_{\infty}.$$

For the strict dual-rail target with the bulk exactly in vacuum, this logarithmic bound is singular because the full-Fock state is not faithful. In that case the trace-norm leakage estimate remains valid, but a power theorem requires either a faithful bulk regularization or an enlarged maintenance task. Thus raw bulk-frequency tails are not silently identified with a logical floor.

If far and near channels have nonnegative strengths η_{far} and η_{near} , share the same temperature and, when used, the same chemical potential, and are otherwise identical, their logical generator is

$$Q_{\text{ell}} = s_{\text{ell}} Q_0,$$

$$s_{\text{ell}} = \eta_{\text{far}} v_{\text{far}}(\text{ell}) + \eta_{\text{near}} v_{\text{near}}(\text{ell}),$$

where Q_0 has rates γ_+ , γ_- .

10. End-to-end SSH maintenance-horizon theorem

Let α be the target probability of logical rail $|1_L\rangle$ and define

$$c(\alpha) = [\sqrt{(1-\alpha)\gamma_+} - \sqrt{\alpha\gamma_-}]^2.$$

This vanishes exactly when the target equals the thermal logical occupation.

Theorem 10.1 - Microscopic geometry-to-holding-cost equality

For the spectrally filtered fixed-parity dual-rail SSH code, the minimum continuous KL-control cost holding the logical populations is

$$C_{\text{ell}}(\alpha) = s_{\text{ell}} c(\alpha).$$

The optimum is attained by the explicit controlled logical rates

$$p^*_{\{0 \rightarrow 1\}} = s_{\text{ell}} \sqrt{\alpha \gamma_+ \gamma_- / (1-\alpha)},$$

$$p^*_{\{1 \rightarrow 0\}} = s_{\text{ell}} \sqrt{(1-\alpha) \gamma_+ \gamma_- / \alpha}.$$

Moreover,

$$\begin{aligned} c(\alpha) [\eta_{\text{near}} (1-\zeta^2)^2 + \eta_{\text{far}} (1-\zeta^2)^2 \zeta^{4\text{ell}}] \\ \leq C_{\text{ell}}(\alpha) \leq \end{aligned}$$

$$c(\alpha) [\eta_{\text{near}} + \eta_{\text{far}} \zeta^{4\text{ell}}].$$

Proof

Proposition 9.1 reduces the filtered logical sector exactly to the reversible two-state process $s_{e11} \rightarrow Q_0$. The reversible holding-cost formula is homogeneous in the generator, giving $C_{e11} = s_{e11} c(\alpha)$ and the displayed Doob rates. The two-sided bounds follow from Corollary 8.2. QED.

Consequences

- **Achievable isolation:** if $\eta_{\text{near}}=0$, the optimum has matching upper and lower exponential scaling $\Theta(\zeta^{4e11})$.
- **Persistent horizon:** if $\eta_{\text{near}}>0$ and $c(\alpha)>0$, then

$$C_{e11}(\alpha) \geq \eta_{\text{near}}(1-\zeta^2)^2 c(\alpha) > 0$$

for every membrane width.

- **Fixed-point exception:** if the target is thermal, $c(\alpha)=0$; neither channel creates a holding cost.
- **Uniform secular validity:** the spectral-filter window can be chosen independently of $e11$ because the bulk gap is bounded below by $t_2 - t_1$.

Multiplying C_{e11} by $k_B T$ gives the optimal power in the trajectory-KL control convention. This is an exact operational statement for that control model; it is not asserted to equal the battery-assisted thermal-operation power of every implementation.

The theorem closes, in a single specified model,

fixed-parity membrane code $\rightarrow$ boundary transfer weight $\rightarrow$ Davies rates $\rightarrow$ optimal controller $\rightarrow$ two-sided holding cost.




Figure 1. Left: exact fixed-parity dual-rail SSH costs with $\zeta=0.72$ and transfer weights $v_x=w_x^2$; blue uses $\eta_{\text{far}}=1$ and $\eta_{\text{near}}=0$, while orange uses $\eta_{\text{far}}=0.77$ and $\eta_{\text{near}}=0.23$. The dashed line is the rigorous floor $0.23(1-\zeta^2)^2$. Right: incompatible-equilibrium cancellation for $\alpha=0.7$, $(u,v)=(4,1)$, $(r_{\text{rate}},s)=(0.2,3)$, $g=1$, and $a_{\text{ell}}=\exp(-m \text{ell})$, with $m=-\log(0.304)/6$, placing the zero at $\text{ell}=6$.

11. Quantum rapid-replacement equality

The classical KL controller is not the only available achievability mechanism. A fully quantum equality follows once the maintenance protocol and its resource ledger are specified precisely.

Fix a faithful quantum target ρ , system Hamiltonian H , and uncontrolled thermal semigroup $T_t=\exp(tL)$ with stationary Gibbs state γ . For a fixed period τ , let a cycle consist of free evolution $\rho \rightarrow \rho_\tau=T_\tau(\rho)$ followed by an admissible operation restoring ρ .

For subsystem i , write the absolute nonequilibrium free energy as $F_i(\omega)=\text{Tr}(H_i \omega)-k_B T S(\omega)$ and its excess above equilibrium as

$$F_i^{\text{ex}}(\omega)=F_i(\omega)-F_i(\gamma_i)=k_B T D(\omega||\gamma_i).$$

Only free-energy differences enter the resource ledger. When a common chemical potential is included, replace H_i by $H_i - \mu N_i$ and require the allowed unitary to conserve both total energy and total particle number.

Operational axioms for the resource-cell model

- **Fresh product inputs.** Cycle n begins with the system marginal ρ_{τ} , a fresh bath in its Gibbs state, and fresh resource systems in a product state uncorrelated with the system and all earlier cycles.
- **Allowed operation.** The restoring operation is induced by a unitary commuting with the sum of the fixed system, bath, and resource Hamiltonians, followed by discarding the bath and outgoing resources. No time-dependent Hamiltonian or unrecorded work source is allowed.
- **Endpoint condition.** The final system marginal is exactly ρ . Outgoing system-resource correlations are permitted but discarded; no memory is carried between cycles.
- **Ledger.** The cost of a cycle is the sum of the decreases $F(\omega_{\text{in}}) - F(\omega_{\text{out}})$ of every nonthermal resource marginal. Bath athermalities and final correlations are treated as discarded waste and are not credited back to the controller. Any catalyst must be returned exactly and uncorrelated at the end of the cycle; otherwise it is included among the consumed resources.
- **Asymptotic power.** For N cycles, A_N is the additive ledger and $P_{\tau} = \limsup_{N \rightarrow \infty} A_N / (N \tau)$. The optimum P_{τ}^* is the infimum over protocols obeying these axioms.

This model deliberately allows nonstationary coherent resource cells. Their time-translation asymmetry is a consumed resource even though the scalar ledger records only their nonequilibrium free-energy decrease. Optimality below is therefore relative to this declared scalar ledger, not to the full preorder of asymmetry resources.

Lemma 11.1 - Multipartite free-energy ledger

Let the final state of system, bath, and resource cells be Ω_f and let

$$\Gamma_{\text{tot}} = \Gamma_S \otimes \Gamma_B \otimes \prod_j \Gamma_{R_j}.$$

For any multipartite state,

$$D(\Omega_f || \Gamma_{\text{tot}}) = \sum_i D(\Omega_{\{f, i\}} || \Gamma_i) + I_{\text{tot}}(\Omega_f),$$

where

$$I_{\text{tot}}(\Omega_f) = \sum_i S(\Omega_{\{f, i\}}) - S(\Omega_f) \geq 0.$$

If the initial state is the product $\rho_{\tau} \otimes \Gamma_B \otimes \prod_j \Omega_j$, the global unitary commutes with the total Hamiltonian, and the final system marginal is ρ , then

$$\begin{aligned} & \sum_j [F(\omega_j) - F(\Omega_{\{f, R_j\}})] \\ &= F_S(\rho) - F_S(\rho_{\tau}) + F_B^{\text{ex}}(\Omega_{\{f, B\}}) + k_B T I_{\text{tot}}(\Omega_f) \end{aligned}$$

$$\geq F_S(\rho) - F_S(\rho_{\tau}).$$

Here $F_B^{\text{ex}}(\Omega_{\{f, B\}}) = F_B(\Omega_{\{f, B\}}) - F_B(\Gamma_B) = k_B T D(\Omega_{\{f, B\}} || \Gamma_B) \geq 0$; it is the bath's excess free energy, not its absolute free energy.

Proof

Expand the relative entropy against the product Gibbs reference, add and subtract the marginal entropies, and identify I_{tot} . The initial multi-information and bath relative entropy vanish. Energy conservation implies that the global Gibbs reference commutes with the unitary, so global relative entropy is invariant. Equating the initial and final decompositions and multiplying by $k_B T$ gives the identity. QED.

Theorem 11.2 - Exact fixed-period resource-cell power

Within the operational class above, assume that the controller may use a fresh cell with Hamiltonian $H_W = H$ prepared in ρ , and may apply the system-cell SWAP. Then the optimal period- τ resource-cell power is

$$P_{\tau}^*(\rho, L) = [F(\rho) - F(T_{\tau} \rho)] / \tau$$

$$= (k_B T / \tau) [D(\rho || \Gamma) - D(T_{\tau} \rho || \Gamma)].$$

Proof

Every restoring cycle maps ρ_{τ} back to ρ . Lemma 11.1 shows that its resource ledger decreases by at least $F(\rho) - F(\rho_{\tau})$. Summing independent cycles gives the same lower bound on P_{τ} .

For the reverse inequality, take an identical fresh cell in state ρ and SWAP it with the system. Since the two Hamiltonians are identical, SWAP commutes with $H_S + H_W$. The system is restored exactly and the cell changes from ρ to ρ_{τ} , consuming exactly $F(\rho) - F(\rho_{\tau})$. Thus the lower bound is attained. QED.

For the fermionic dual-rail application, SWAP means the canonical graded, second-quantized permutation of corresponding system and cell modes. It is an even, number-conserving fermionic unitary and can be synthesized from even bilinear mode permutations and number-phase corrections. Because the system and cell each remain in an odd-parity one-particle code, graded SWAP differs from the ordinary tensor-factor SWAP only by the constant sign $(-1)^{(1*1)} = -1$, a global phase with no physical effect.

This construction is close in spirit to the target-copy SWAP stabilization studied by Vacanti, Elouard, and Auffeves. Their resource system is initially stationary with respect to its own engineered Hamiltonian and the mismatch with the system Hamiltonian produces an external work cost. Here the cell Hamiltonian equals H , so SWAP conserves energy, but a coherent ρ makes the cell nonstationary and explicitly imports asymmetry. The operational resources are therefore different.

Corollary 11.3 - Exact fast-maintenance power

Define rapid maintenance by taking the fixed-period optimum before the limit:

$$P_{\text{fast}}^*(\rho, L) = \lim_{\tau \downarrow 0} P_{\tau}^*(\rho, L).$$

Then

$$P_{\text{fast}}^*(\rho, L) = k_B T \sigma_L(\rho),$$

where

$$\sigma_L(\rho) = -\frac{d}{dt} D(\exp(tL)\rho) \big|_{t=0}.$$

The maximum deviation from the target during a cycle also tends to zero with τ , so this limit realizes continuous corridor maintenance. The order of operations is essential: optimize at fixed τ , then take $\tau \rightarrow 0$; do not take an infimum over unrestricted periods.

The theorem is noncommutative: ρ need not commute with H . In that case the fresh cells supply time-translation asymmetry as well as free energy. The equality is therefore a statement about **resource-cell consumption under the stated ledger**, not a claim that ordinary incoherent mechanical work can prepare coherent cells at the same cost.

Theorem 11.4 - Fixed-parity quantum SSH maintenance horizon

Let the physical Schrödinger-picture generator on the spectrally filtered dual-rail code be

$$G_{\text{ell}}(\rho) = -i[H_{\text{log}}^{\text{bare}} + s_{\text{ell}} H_{\text{LS}}, \rho] + s_{\text{ell}} D_0(\rho),$$

where H_{LS} commutes with the bare logical Hamiltonian $H_{\text{log}}^{\text{bare}}$, D_0 is the unit-strength primitive Davies dissipator, and $s_{\text{ell}} = \eta_{\text{far}} v_{\text{far}} + \eta_{\text{near}} v_{\text{near}}$. Equivalently, in the interaction picture of $H_{\text{log}}^{\text{bare}}$,

$$\tilde{L}_{\text{ell}} = s_{\text{ell}} L_0, L_0(\rho) = -i[H_{\text{LS}}, \rho] + D_0(\rho).$$

Restrict this dynamics to

$$H_{\text{log}} = \text{span}\{f_{\{0,0\}}^{\dagger} | \text{vac}\rangle, f_{\{0,1\}}^{\dagger} | \text{vac}\rangle\}.$$

Let ρ be any faithful density matrix on this code, including a state with coherence between the two rails. Both basis states have $N=1$ and odd parity, so ρ is a physical fermionic state under parity superselection. Then

$$P_{\text{fast}}^*(\rho, G_{\text{ell}}) = k_B T s_{\text{ell}} \sigma_{\{D_0\}}(\rho).$$

Consequently,

$$\begin{aligned} k_B T \sigma_{\{D_0\}}(\rho) [\eta_{\text{near}}(1-\zeta^2)^2 + \eta_{\text{far}}(1-\zeta^2)^2 \zeta^4] \\ \leq P_{\text{fast}}^*(\rho, G_{\text{ell}}) \leq \end{aligned}$$

$$k_B T \sigma_{\{D_0\}}(\rho) [\eta_{\text{near}} + \eta_{\text{far}} \zeta^4].$$

If D_0 is primitive, $\sigma_{\{D_0\}}(\rho) > 0$ for every nonthermal target. Thus the theorem gives achievable $\Theta(\zeta^4)$ isolation when $\eta_{\text{near}} = 0$ and an achievable, strictly positive resource horizon when $\eta_{\text{near}} > 0$,

for physical coherent as well as diagonal logical targets. No parity reference frame is assumed or consumed.

Proof

Proposition 9.1 gives the interaction-picture factorization $\tilde{L}_{\text{ell}} = s_{\text{ell}} L_0$. The logical Gibbs state γ_{log} commutes with both $H_{\text{log}}^{\text{bare}}$ and the diagonal Lamb shift. For any Hamiltonian K commuting with γ_{log} , cyclicity of the trace and $[\rho, \log \rho] = 0$ give

$$-\text{Tr}\{-i[K, \rho](\log \rho - \log \gamma_{\text{log}})\} = 0.$$

Hence the unattenuated bare Hamiltonian and the Lamb shift make no contribution to the relative-entropy derivative, and

$$\sigma_{\{G_{\text{ell}}\}}(\rho) = s_{\text{ell}} \sigma_{\{D_0\}}(\rho).$$

Corollary 11.3 gives the power equality, while Corollary 8.2 bounds v_{far} and v_{near} . Primitivity makes the dissipative entropy production strictly positive away from the logical Gibbs state. Because the code projector lies wholly in the $N=1$, odd-parity sector and both the free evolution and fermionic SWAP preserve that sector, the protocol never invokes an inadmissible even-odd coherence. QED.

12. Uniform families and the remaining quantum frontier

For a family indexed by volume N , Theorem 6.1 is algebraically uniform. Nontrivial uniform constants require target families satisfying

$$0 < s_- \leq C_{\{Q_N\}}(\mu_N) \leq s_+ < \infty$$

and, in the floor regime,

$$C_{\{R_N\}}(\mu_N) \geq r_- > 0.$$

The general many-body quantum target is now precise. Beyond the exactly closed logical-qubit replacement model, one needs a Davies or KMS-symmetric theory in which:

- the holding cost is operationally defined by a cumulative work ledger;
- the target is invariant under a physically restricted autonomous controlled Lindbladian;
- a quantum Doob transform or another controller attains the optimum without consuming preloaded target copies;
- the optimal cost is a generator-linear KMS Dirichlet functional, at least on a typed family;
- a membrane model supplies two-sided, volume-uniform bounds on a_{ell} .

The classical theorem shows exactly where a quantum result can fail: noncommutativity, coherence resources not priced by free energy alone, incompatible fixed-point states, and a loss of additivity of the optimal control functional.

13. Status and limitations

Proved in this draft:

- endpoint-restoration degeneracy under uniformly bounded reset cost;
- an exact dephasing-qubit SWAP witness;
- a self-contained proof of the reversible finite-chain optimizer and Dirichlet formula (known input, not a novelty claim);
- exact additivity, two-sided scaling, and controller achievability;
- a two-state cancellation counterexample outside the common-reference class.

Not yet proved:

- conversion of KL rate into hardware work without the stochastic-control realization;
- volume-uniform lower bounds for a nontrivial interacting target family.
- uniform-in-ell error estimates for the prelimit microscopic weak-coupling approximation;
- autonomous coherent maintenance using only stationary batteries and ordinary work.

Additional results proved by the microscopic closure:

- exact odd-chain SSH modes and a bulk gap uniform in membrane width;
- a fixed-number, fixed-parity dual-rail code supporting physical logical coherence;
- two-sided exponential bounds on the remote boundary transfer weight;
- an exact filtered, number-conserving Davies reduction to the logical birth-death process;
- a constructed optimal controller and a persistent-floor/cheap-isolation dichotomy with no interface hypothesis.
- an exact quantum fixed-period and rapid-maintenance equality using energy-conserving SWAPs with fresh resource cells;
- an end-to-end coherent fixed-parity logical-qubit SSH theorem with explicit upper and lower power bounds.

The autonomous Lindbladian-control problem remains open. The resource-cell theorem does maintain noncommuting targets, but it prices a stream of preloaded target copies and should not be conflated with ordinary work-only control.

Appendix A. Reproducibility and numerical scope

The numerical suite is diagnostic, not a substitute for the analytic proofs.

- **Endpoint witness:** $H=0$, coherence $c=0.3$, and dephasing rate $\Gamma=1$. The periodic SWAP power approaches the instantaneous rate as $\tau \rightarrow 0$ and falls to 0.00463 of that rate at $\tau=100$.
- **Reversible additivity:** 100 complete reversible generators of dimensions 2 through 8 were generated from random symmetric conductances and random faithful ρ , μ . The worst stationarity residual was $3.7e-15$; the worst additivity and KL/Dirichlet discrepancies were $1.8e-15$.
- **SSH closure:** $t_1=0.37$, $t_2=1$, and widths $e_{ll}=1, \dots, 40$ on each rail. The largest zero-mode residual was $2.9e-17$, the largest boundary-weight discrepancy $1.4e-17$, the smallest gap margin above t_2-t_1 was $1.7e-3$, the factorization error was $1.1e-17$, and the largest revised-bound residual was $5.2e-18$.
- **Coherent logical target:** amplitude-damping coefficients $\gamma_+=0.4$, $\gamma_-=1.2$ and dual-rail target $\rho = [[0.65, 0.18+0.07i], [0.18-0.07i, 0.35]]$ in the basis $\{|10\rangle, |01\rangle\}$. The largest direct Hamiltonian contribution to the relative-entropy derivative was $3.5e-18$. At $\tau=1e-6$, the worst rapid-limit discrepancy over $e_{ll}=1, \dots, 30$ was $8.1e-9$; the revised SSH bounds had no violation.
- **Ledger, parity, and tails:** 100 random tripartite zero-Hamiltonian tests reproduced Lemma 11.1 with worst residual $1.7e-15$; the Jordan-Wigner dual-rail transfer commuted exactly with total number and parity, the embedded coherent state lay exactly in the $N=1$ parity sector, the dual-rail dissipator had zero code leakage, and 1000 random dissipator tests produced no violation of the $2\|L\|^2$ trace-norm bound.

All tolerances, parameters, and formulas used in Figure 1 are stated above so the plots can be reconstructed independently.

References

- J. M. Horowitz, K. Zhou, and J. L. England, *Minimum energetic cost to maintain a target nonequilibrium state*, arXiv:1703.04722 (2017).
- D. S. Pavlichin, Y. Quek, and T. Weissman, *Minimum Power to Maintain a Nonequilibrium Distribution of a Markov Chain*, arXiv:1907.01582 (2019).
- J. Fuchs, S. Goldt, and U. Seifert, *Stochastic thermodynamics of resetting*, arXiv:1603.01141 (2016).
- E. Todorov, *Linearly-solvable Markov decision problems*, Advances in Neural Information Processing Systems 19 (2007).
- M. Zaimi et al., *Detecting topological edge states with the dynamics of a qubit*, arXiv:1910.09926 (2019).
- S. Cusumano and G. De Chiara, *Structured quantum collision models: generating coherence with thermal resources*, arXiv:2307.07463 (2023).
- F. Carollo, J. P. Garrahan, I. Lesanovsky, and C. Perez-Espigares, *Making rare events typical in Markovian open quantum systems*, arXiv:1711.10951 (2017).
- E. A. Carlen and J. Maas, *Gradient flow and entropy inequalities for quantum Markov semigroups with detailed balance*, arXiv:1609.01254 (2016).
- G. Vacanti, C. Elouard, and A. Auffeves, *The work cost of keeping states with coherences out of thermal equilibrium*, arXiv:1503.01974 (2015).
- J. M. Horowitz and K. Jacobs, *The Energy Cost of Controlling Mesoscopic Quantum Systems*, arXiv:1509.00914 (2015).
- E. B. Davies, *Markovian master equations*, Commun. Math. Phys. 39, 91-110 (1974).
- H. Spohn, *Entropy production for quantum dynamical semigroups*, J. Math. Phys. 19, 1227-1230 (1978).
- W. P. Su, J. R. Schrieffer, and A. J. Heeger, *Solitons in polyacetylene*, Phys. Rev. Lett. 42, 1698-1701 (1979).
- S. Szalay, Z. Zimboras, M. Mate, G. Barcza, C. Schilling, and O. Legeza, *Fermionic systems for quantum information people*, arXiv:2006.03087 (2020).
- N. T. Vidal, M. L. Bera, A. Riera, M. Lewenstein, and M. N. Bera, *Quantum Operations in an Information Theory for Fermions*, Phys. Rev. A 104, 032411 (2021), arXiv:2102.09074.