The heat price of precision

Seth Douglas Nidhal Mghirbi

Version 1.0.0, October 2026

Abstract

Does playing a nonlocal game more precisely require more heat? We study devices that start and finish in pure product blanks, receive thermal answer carriers, and release classical answers. The answer depends on when the devices may communicate again. If they reset locally first, their minimum cyclic work differs from the game's minimum mutual information by a bounded alphabet cost and is at least twice its one-way distillable entanglement demand. Published entanglement bounds therefore give fixed finite games with local work $2\alpha \log_2(1/\varepsilon) + O(\log\log(1/\varepsilon))$, for a dense set of coefficients α, including transcendental values. Quantum reunion before reset keeps the work bounded.

With coherent block recovery allowing q qubits and c dephased bits per play, the leading cost is $(\alpha L - q)_+ + (\alpha L - q - c)_+$, where $L = \log_2(1/\varepsilon)$. A relative-entropy potential proves the lower bound; compression proves the upper bound. An exact entropy-production identity records arbitrary schedules without solving their optimization. For exactly matched controls, a sector allocation theorem evaluates the entanglement and correlation costs. A W-state example shows why knowledge of accessible marginals cannot replace operational access. Perfect magic-square play has cyclic local work exactly four units. Work and bath heat agree at the zero-energy endpoints used here; one unit is $k_B T \ln 2$.

The authors developed this work with current frontier AI models from more than one developer and take responsibility for every claim. How the work was done and checked, including a proof audit, two blind audits and finite computational checks, is described before the appendices.

Contents

1 The question and the operating cycle **2**

2 Local reset and quantum reunion **3**

3 What an arbitrary schedule records **4**

4 Quantum and classical return **5**

5 Exactly matched controls and sector allocation **8**

6 Calibration and a finite game floor **10**

7 Operational access and prior work **13**

How this work was done **14**

A Weighted Weyl tests for several observers **15**

B Several phase controls sharing one shift **16**

C Approximate invariance and its limit **18**

1 The question and the operating cycle

How much heat must a pair of devices discharge to play one fixed nonlocal game almost perfectly? Some finite games need an increasing amount of usable entanglement as their loss approaches zero. That fact alone does not set a heat cost. The devices may keep their correlations until they can act jointly, or they may have to erase each half while its partner is inaccessible. This paper asks what that reset deadline changes.

There are three takeaways. Local reset converts a sharp entanglement demand into a sharp work law. Quantum reunion before reset removes its divergence, whereas classical recovery removes only one of the two leading units. Exact matching provides a sector formula for the correlations that are priced; approximate matching needs additional stability estimates.

The entanglement laws are inputs, not results reproved here. Douglas and Mghirbi [13, Theorem 11.9 and Corollary 11.11] give fixed games with pure entanglement entropy at most $\alpha \log_2(1/\varepsilon) + O_G(1)$ and one-way distillable entanglement demand at least $\alpha \log_2(1/\varepsilon) - O_G(\log\log(1/\varepsilon))$. Their explicit companion [24] supplies an exponential-dimension example. We use the schedule contract of the published time-order papers [14, 11, 15, 12], whose general work law for schedules is [14, Theorem 1 and Appendix B, Theorem 7]. The entropy arguments below establish the bounds used here directly and do not use that law.

Thermodynamic contract. All entropies and relative entropies are in bits. Write $h_2(t) = -t\log_2 t - (1-t)\log_2(1-t)$, with $0\log_2 0 = 0$, and $w = W/(k_B T \ln 2)$, positive for work supplied. Every logical carrier has zero Hamiltonian at its endpoints. Controllers return, and batteries finish without entropy or correlations that could store information. Baths initially are independent Gibbs stocks at one temperature. Every discarded record belongs to the accounting. A reset means return to the specified pure blank, decoupled from all other systems.

A round begins with pure product working blanks. A station with joint access prepares a finite state ρ_{AB} before questions (X, Y), sampled from the game's fixed distribution μ independently of the preparation and the baths, arrive. Its environment then becomes inaccessible. Every register the players can use belongs to A or B. There is no uncharged shared catalyst. Alice and Bob receive independent thermal answer carriers I_m/m and I_n/n. Let $K = mn$. Question labels remain available as unchanged local classical controls through reset. Classical answers leave permanently; the referee gives no reset assistance.

In the *local* schedule no communication is allowed until both working devices are blank. In the *joint* schedule the residual devices can reunite quantum mechanically after answer release and before reset. Recovery communication between these endpoints is treated in Section 4.

Limits and resources. Fix a finite strategy and ε first. Its thermodynamic work is the reversible rate on independent copies, with whole-block error tending to zero, all entropy and energy return errors $o(N)$, and declared hardware. Then optimize over strategies, and only then let $\varepsilon \downarrow 0$. Local compression across copies within a laboratory is allowed. The ban on communication persists across the entire local reset block. This convention does not assert additivity of entanglement demand for arbitrary correlated multiround strategies. Initial blank hardware is startup capital; its preparation is not charged again each round. The full operating cycle restores it.

For a game G and a loss that is attainable in finite dimension, define

$$d_G(\varepsilon) = \inf_{\text{loss}\le\varepsilon} E_{\mathrm{D}}^{\rightarrow}(\rho_{AB}), \qquad E_G(\varepsilon) = \inf_{\substack{\text{loss}\le\varepsilon \\ \rho \text{ pure}}} S(\rho_A), \tag{1.1}$$

$$C_G(\varepsilon) = \inf_{\text{loss}\le\varepsilon} I(A:B)_\rho, \qquad I(A:B) = S(A) + S(B) - S(AB). \tag{1.2}$$

The infima include arbitrary finite dimensions and mixed strategies where indicated. Let $D_G(\varepsilon)$ be the minimum possible maximum local dimension. All mathematical assertions labelled Proved have a proof below or an exact imported theorem citation. Experimental implementations are labelled Sketch; the missing stability extension is labelled Conjecture.

Work equals the total heat delivered to the baths in this zero-endpoint-energy model, since energy conservation gives $W = \Delta U + Q_{\mathrm{bath}}$ and $\Delta U = 0$. This equality does not extend to arbitrary nondegenerate devices.

2 Local reset and quantum reunion

Theorem 2.1 (Proved: the local correlation price)**.** *For every finite game and attainable loss,*

$$C_G(\varepsilon) \le w_{\mathrm{local,G}}(\varepsilon) \le C_G(\varepsilon) + 2\log_2 K, \tag{2.1}$$

$$2d_G(\varepsilon) \le C_G(\varepsilon) \le 2E_G(\varepsilon) \le 2\log_2 D_G(\varepsilon). \tag{2.2}$$

Moreover $0 \le w_{\mathrm{joint,G}}(\varepsilon) \le 2\log_2 K$*.*

Proof. For degenerate endpoints the entropy balance, Landauer's bound [21] in general form, is $w \ge S_{\mathrm{in}} - S_{\mathrm{out}}$. Reversible formation of a known state from a pure blank costs $-S(\rho)$; reversible reset costs $S(\rho)$. These are the asymptotic erasure primitives of [8, Corollary 2 and Appendix D], with their lower bounds supplied by the entropy balance.

Formation gives $w_{\mathrm{prep}} \ge -S(AB)$. At fixed Alice question x, her device and thermal answer port start with entropy $S(A) + \log_2 m$ and end with a pure device and a classical answer of entropy $H(a|x)$. The answer can remain in the entropy ledger after it leaves; later maps act trivially on it. Alice's bath is independent of her marginal, so

$$w_A(x) \ge S(A) + \log_2 m - H(a|x).$$

Bob obeys the corresponding inequality. Their sum with formation is at least

$$I(A:B) + \log_2 K - H(a|X) - H(b|Y) \ge I(A:B).$$

The same calculation applies to blocks. This proves the first lower bound.

For the upper bound, reversibly form ρ_{AB} and pay $\log_2 K$ for pure answer blanks. For a POVM M^x, use the isometry $V_x = \sum_a \sqrt{M_a^x} \otimes |a\rangle_P$ with a retained pointer P. Copy P to the answer port and release the copy. The retained state is $V_x \rho_A V_x^*$ dephased in m pointer values. Pinching is an average of m phase unitaries, so its entropy is at most $S(A) + \log_2 m$. Erase that state locally. Bob's cost is at most $S(B) + \log_2 n$. The total is at most $I(A:B) + 2\log_2 K$. All isometries extend to square unitaries on pure working ancillas, included in the reset.

Here is a proof of $2E_{\mathrm{D}}^{\rightarrow}(\rho) \le I(A:B)_\rho$. On N copies, combine Alice's part of a one-way distillation protocol into an instrument with output $A'Z$, where Z is the classical message. Data processing and the chain rule give

$$NI(A:B) \ge I(A'Z:B^N) \ge \sum_z p_z I(A':B^N)_z \ge 2\sum_z p_z [S(B^N)_z - S(A'B^N)_z].$$

The last inequality is Araki–Lieb. The sum is the coherent information $J(A'\rangle B^N Z)$. Bob's controlled channel cannot increase coherent information. For an output tending to M_N Bell pairs, entropy continuity gives $J(A'\rangle B') \ge M_N - o(N)$. Thus $I(A:B)/2$ bounds every one-way distillation rate. For a pure strategy $I = 2S(A)$. If a mixed state of local dimensions at most D succeeds, some pure component succeeds with the same measurements, since score is affine in the state. Its entropy is at most $\log_2 D$. These observations prove (2.2).

For joint reset choose a successful pure component, form it by a global unitary, and use the coherent measurement implementation above. At fixed x, y the residual is an orthogonally flagged mixture

$$\sum_{ab} p(ab|xy)|\phi_{ab}\rangle\langle\phi_{ab}| \otimes |ab\rangle\langle ab|_P.$$

After reunion, a unitary controlled by the joint pointer maps each normalized ϕ_{ab} to the same product blank. Only the pointer remains. Its erasure costs at most $\log_2 K$; answer blanks cost the other $\log_2 K$. Question controls permit the required choice of unitary. Nonnegative total work follows from the entropy balance for the complete cycle: its boundary term is $\log_2 K - H(ab|XY) \geq 0$. □

Corollary 2.2 (PROVED: the sharp precision law)**.** *Put $L = \log_2(1/\varepsilon)$. Suppose, for small enough ε,*

$$E_G(\varepsilon) \leq \alpha L + a, \qquad d_G(\varepsilon) \geq \alpha L - b \log_2 L - c.$$

Then

$$2\alpha L - 2b \log_2 L - 2c \leq w_{\mathrm{local,G}}(\varepsilon) \leq 2\alpha L + 2a + 2\log_2 K. \tag{2.3}$$

Consequently $w_{\mathrm{local,G}} = 2\alpha L + O_G(\log L)$, whereas $w_{\mathrm{joint,G}} = O_G(1)$. The set

$$\mathcal{S} = \left\{ \frac{1}{2} \sum_{m \text{ even}} \frac{n_m}{\log_2 m} : n_m \in \mathbb{Z}_{\geq 0},\ 0 < \sum_m n_m < \infty \right\}$$

is realized by fixed finite games. The work coefficients $2\mathcal{S}$ are dense in $(0, \infty)$ and include transcendental numbers.

Proof. Substitute the two inequalities into Theorem 2.1. The existence of games with those inequalities for every $\alpha \in \mathcal{S}$ is [13, Theorem 11.9(a,b)]; density and transcendence are [13, Corollary 11.11]. Multiplication by two preserves both properties. In particular the upper correction is bounded; the lower correction need only be $O(\log L)$. No inference from a formation bound to a distillable-entanglement bound is used. □

The same d_G and D_G appear in both reset schedules. If d_G diverges, one work profile diverges and the other stays bounded. Hence these game resources alone cannot determine operating work without a reset schedule.

3 What an arbitrary schedule records

The next identity is standard entropy production [16, 27]. Its use here fixes the complete boundary of a game cycle. It does not evaluate an optimization over causal implementations.

Proposition 3.1 (PROVED: the complete entropy ledger)**.** *Let R contain every nonbath carrier and reference being tracked, apart from an ideal entropy-preserving work battery. Independent reservoirs start in $\Gamma_E = \bigotimes_j \gamma_{E_j}$, initially independent of R. Assume unitary global evolution, a final decoupled battery of unchanged entropy, and vanishing endpoint interactions. For final state η_{RE} define*

$$A(\Pi) = D_2(\eta_{RE} \| \eta_R \otimes \Gamma_E).$$

Then every physical implementation satisfies

$$w(\Pi) = \frac{\Delta U_R}{k_B T \ln 2} + S(R_i) - S(R_f) + A(\Pi), \tag{3.1}$$

$$A(\Pi) = S(R_f) + \sum_j S(E_j)_f - S(RE)_f + \sum_j D_2(\eta_{E_j} \| \gamma_{E_j}). \tag{3.2}$$

For r game plays with independent thermal ports, a classical question transcript Q, no additional input purifiers, and fully restored hardware,

$$w(\Pi) = r \log_2 K - H(O|Q)_\Pi + A(\Pi). \tag{3.3}$$

Proof. Unitarity and the battery assumptions give $S(RE)_f = S(R_i) + S(\Gamma_E)$. Expanding the relative entropy and using the Gibbs formula gives

$$A = -S(RE)_f + S(R_f) + S(\Gamma_E) + \frac{Q_{\text{bath}}}{k_B T \ln 2} = S(R_f) - S(R_i) + \frac{Q_{\text{bath}}}{k_B T \ln 2}.$$

Energy conservation $W = \Delta U_R + Q_{\text{bath}}$ proves (3.1). Adding and subtracting $\sum_j S(E_j)_f$ proves (3.2). Initially the classical questions and thermal answer ports have entropy $H(Q) + r \log_2 K$; finally the questions and classical answers have entropy $H(QO)$. Hardware is pure at both ends. This proves (3.3). For approximate return, include the controller energy, entropy and battery errors explicitly; their $o(r)$ budgets vanish in the stipulated rate. An arbitrary purifying observer instead requires the full boundary term in (3.1). □

For a specified finite release and return schedule Σ, minimize (3.3) over its admissible implementations and take precisely the same rate closure as for physical work. This yields the exact operational formula

$$w_{\Sigma,G}(\varepsilon) = \inf_{\Pi \text{ admissible}} \frac{r \log_2 K - H(O|Q)_\Pi + A(\Pi)}{r}. \tag{3.4}$$

Indeed the identity holds for every implementation, so it holds along every approximating sequence for either infimum. With $R_{\Sigma,G} = \inf_\Pi A(\Pi)/r$ and the same closure,

$$R_{\Sigma,G} \le w_{\Sigma,G} \le R_{\Sigma,G} + \log_2 K. \tag{3.5}$$

This follows from $0 \le r \log_2 K - H(O|Q) \le r \log_2 K$, before taking infima.

Reservoirs in (3.2) can hold correlations with each other even when the final devices are blank. Their return is an information return and must be allowed by the schedule. Calling the optimization an exact formula does not supply a closed entropy expression for every schedule. In particular, the serial Werner–Holevo example of [11, Theorem 3 and Section 5] rules out replacing a constrained implementation by a reversible batch implementation at each stage. The identity itself is credited to the entropy-production literature, rather than claimed as new.

4 Quantum and classical return

After answer release, allow coherent local processing on N independent copies of the chosen strategy. Recovery may transfer at most $Nq + o(N)$ qubits and $Nc + o(N)$ fully dephased bits in total, with an Alice-to-Bob route available for the construction. A quantum carrier of dimension d uses $\log_2 d$ qubits of capacity. Classical carriers are physically moved dephased registers. Blank communication hardware, flags, failed branches and copied messages all enter the work ledger; no pure message register or entangled pair is free. Questions are local controls as in Section 1. This model allows compression before transfer and imposes the blank deadline after the recovery block.

Let

$$F(e; q, c) = (e - q)_+ + (e - q - c)_+, \qquad B_G = \log_2(|X||Y|K).$$

Theorem 4.1 (PROVED: the recovery tradeoff)**.** *In this block model, for $q, c \geq 0$,*

$$F(d_G(\varepsilon); q, c) \leq w_{q,c,G}(\varepsilon) \leq F(E_G(\varepsilon); q, c) + 4B_G. \tag{4.1}$$

For the games of Corollary 2.2, uniformly over $q, c \geq 0$,

$$w_{q,c,G}(\varepsilon) = (\alpha L - q)_+ + (\alpha L - q - c)_+ + O_G(\log L). \tag{4.2}$$

At fixed total dimension, the potential below is the complement of $S + E_{\mathrm{R}}$. Horodecki and collaborators show that $S + E_{\mathrm{R}}$ cannot decrease under closed local operations, including transfer of dephased registers, and use this to bound the work deficit and the cost of erasing entanglement [18, Theorems 2 and 3 and the proof of Theorem 3]. The proof below adds the disposal of independent local Gibbs baths and a charged quantum-transfer bound for the present cycle.

Lemma 4.2 (PROVED: a work and entanglement potential)**.** *For degenerate bipartite hardware set*

$$M(\rho) = \log_2 \dim(AB) - S(\rho) - E_{\mathrm{R}}(\rho), \qquad E_{\mathrm{R}}(\rho) = \inf_{\sigma \text{ separable}} D_2(\rho \| \sigma).$$

A local thermal step supplied with w work units increases M by at most w. Moving an already classical register leaves E_{R} unchanged and increases M by zero. Moving a d-dimensional quantum register can increase M by at most $\log_2 d$.

Proof. Append a local Gibbs bath E, and take the entanglement cut to be $AE : B$. For its final state η, the separable competitor $\sigma_{AB} \otimes \eta_E$ gives

$$E_{\mathrm{R}}(AE : B)_\eta - E_{\mathrm{R}}(A : B)_\eta \leq I(AB : E)_\eta. \tag{4.3}$$

This follows from $D_2(\eta_{ABE} \| \sigma_{AB} \otimes \eta_E) = D_2(\eta_{AB} \| \sigma_{AB}) + I(AB : E)$, followed by taking the infimum. Write F_2 for relative entropy to the product Gibbs state; on AB this is $\log_2 \dim(AB) - S(AB)$. Then

$$F_2(ABE) = F_2(AB) + D_2(\eta_E \| \gamma_E) + I(AB : E).$$

Combining this equality with (4.3) shows that discarding a local bath cannot increase $F_2 - E_{\mathrm{R}}$. Appending that bath adds zero. Local unitary evolution preserves entanglement across $AE : B$; with an ideal local battery the change in relative free energy is the supplied work. This proves the thermal-step bound. Pure ancillas have their relative free energy charged, or are returned; they cannot circumvent it.

For a dephased flag, the direct-sum formula is

$$E_{\mathrm{R}}\left(\sum_z p_z |z\rangle\langle z| \otimes \rho_z\right) = \sum_z p_z E_{\mathrm{R}}(\rho_z).$$

To prove it, dephase any separable competitor in the flag basis. Its blocks remain separable. Relative entropy splits into $D(p\|r) + \sum_z p_z D(\rho_z \| \sigma_z)$, which is at least the right side; competitors with $r = p$ attain the infimum. This proof works with the flag at either party. Transfer changes neither global entropy nor total dimension, so it leaves M unchanged. Dephasing a coherent message is a local physical step and is already subject to the work bound.

Finally simulate the *reverse* transfer of a d-dimensional quantum carrier by teleportation, using Φ_d and free classical communication for this mathematical comparison. LOCC monotonicity and a product separable competitor imply

$$E_{\mathrm{R}}(\text{before}) \leq E_{\mathrm{R}}(\text{after} \otimes \Phi_d) \leq E_{\mathrm{R}}(\text{after}) + \log_2 d.$$

Here $E_{\mathrm{R}}(\Phi_d) = \log_2 d$, by a separable diagonal competitor and the test onto Φ_d, whose overlap with a separable state is at most $1/d$. Hence transfer can decrease E_{R} by at most $\log_2 d$, proving the potential bound. Teleportation in this argument is a comparison of cuts; it supplies no free ebits to the operating protocol. □

Proof of Theorem 4.1. First sum local entropy balances, including formation and thermal answer ports. Moving C from Alice to Bob changes the sum of local entropies by

$$S(A) + S(BC) - S(AC) - S(B) = S(C|B) - S(C|A).$$

Each conditional entropy lies in $[-\log_2 \dim C, \log_2 \dim C]$. For dephased C both lie in $[0, \log_2 \dim C]$. Communication can therefore remove at most $2Nq + Nc + o(N)$ of the local entropy burden. The nonnegative answer-port terms in the proof of Theorem 2.1 give

$$w \geq I(A:B)_\rho - 2q - c \geq 2E_{\mathrm{D}}^{\rightarrow}(\rho) - 2q - c. \tag{4.4}$$

For the second bound use Lemma 4.2 from the prepared state through measurement and recovery. Keep the classical outgoing answer registers on their original sides of the mathematical cut even after release. Subsequent operations act trivially on them. Their dephasing is a local step; their being inaccessible cannot improve a lower bound. The final state on this cut is separable: only pure blanks and classical outputs remain. The initial answer ports are thermal and independent. At fixed question transcript the increase in M is

$$NS(\rho) + E_{\mathrm{R}}(\rho^{\otimes N}) + N\log_2 K - H(O|Q) \geq NS(\rho) + E_{\mathrm{R}}(\rho^{\otimes N}).$$

The quantum transfers contribute at most $Nq + o(N)$. Formation contributes at least $-NS(\rho)$. Thus total work satisfies

$$w \geq E_{\mathrm{R}}^{\infty}(\rho) - q \geq E_{\mathrm{D}}^{\rightarrow}(\rho) - q. \tag{4.5}$$

For clarity, the second inequality needs no unproved additivity. Apply any one-way distillation map to $\rho^{\otimes N}$ and a separable competitor. If the output is within half trace distance t of Φ_M, its projection onto Φ_M succeeds with probability at least $1 - t$; every separable output succeeds with probability at most $1/M$. Binary relative entropy and data processing give $E_{\mathrm{R}}(\rho^{\otimes N}) \geq (1-t)\log_2 M - h_2(t)$. Divide by N, let $t \to 0$, and optimize the rate. The limit defining E_{R}^{∞} exists by subadditivity. Approximate hardware returns are taken at fixed strategy dimension with the entropy budgets of Section 1.

Together with $w \geq 0$, (4.4) and (4.5) give

$$w \geq \max\{0, 2(d_G - q) - c, d_G - q\} = F(d_G; q, c).$$

For achievability choose a successful pure strategy of entropy E. Form it unitarily and use retained measurement pointers and pure answer blanks. Copy local question labels into working flags if desired. The retained average state σ is a mixture of at most $|X||Y|K$ pure branches. Pinching and finite question mixing give

$$S(\sigma) \leq B_G, \qquad s := S(\sigma_A) \leq E + B_G.$$

Coherently compress Alice's N copies to $N(s+\delta)$ qubits while preserving correlations with Bob and the referee. Here is the needed compression argument [28]. In an eigenbasis of σ_A, the i.i.d. typical strings have probability tending to one and number at most $2^{N(s+\delta)}$. A unitary moves their span into that many qubits with the remaining hardware blank. The gentle projection estimate makes the joint error tend to zero. The atypical branch can be retained and erased with $o(N)$ entropy cost at fixed dimension. Let $N \to \infty$ and then $\delta \downarrow 0$.

Send $N\min(q,s)+o(N)$ compressed qubits. Dephase and send a further $Nt+o(N)$ bits, where $t=\min(c,(s-q)_+)$. Alice's remaining register L has logarithmic dimension at most $N(s-q-t)_+ + o(N)$. Dephasing adds at most $Nt+o(N)$ to global entropy. Araki–Lieb gives

$$S(L)+S(B_{\text{final}}) \le 2S(L)+S(LB_{\text{final}}).$$

Both parties can erase their own marginals. Their work per play is at most

$$2(s-q-t)_+ + S(\sigma) + t = F(s;q,c)+S(\sigma).$$

Add $\log_2 K$ for the answer blanks. Because F is nondecreasing and 2-Lipschitz in s, the total is at most $F(E;q,c)+4B_G$. Transfers are ownership transfers of registers; empty slots are pure blanks. A fixed inventory can instead be maintained by swapping into blank receiving slots and charging and returning those slots. Take the infimum over E. The same Lipschitz bound and the two resource estimates in Corollary 2.2 prove (4.2). □

Consequences of the deadline. With $q=0$, unlimited classical communication leaves a cost at least d_G. Returning Bob's device to another processor on Bob's side of the cut changes neither bound. A finite reunion delay also leaves the joint upper bound intact if independent strategy buffers may survive until reunion: a block of r pure strategies leaves at most K^r flagged pure branches, so joint erasure and answer preparation together cost at most $2r\log_2 K$. Storage increases with the delay. If reset is required first, that construction is inadmissible.

For partial return, bounded work requires $(d_G-q)_+ = O(1)$, even with unlimited c. In this block model $(E_G-q)_+ = O(1)$ suffices. The sharp family leaves an $O_G(\log L)$ uncertainty window between these conditions; (4.2) alone does not turn its bounded leading term into bounded actual work. A fixed classical answer carrier can save at most its entropy, and a coherent carrier at most twice its logarithmic dimension, by mutual-information bounds. Thus fixed output alphabets cannot supply diverging quantum return capacity. Serial schedules that forbid compression require their own upper-bound analysis.

5 Exactly matched controls and sector allocation

The schedule states what can be controlled jointly and when. The tests state which correlations a successful preparation must carry. Exact matching makes the latter requirement explicit. It is more restrictive than small game loss.

Let $\mathcal{A}$ be a unital $*$-algebra of finitely many bounded controls, with any exact projective or unitary relations declared. Let $h=\sum_a w_a R_a^* R_a$, $w_a>0$, be a positive defect element. For a finite-dimensional irreducible representation π write

$$\ell_\pi = \frac{\operatorname{Tr}\pi(h)}{\dim\pi}, \qquad c_\pi = \log_2 \dim\pi.$$

The two parties carry opposite representations. For self-adjoint generators, *exact matching* means

$$(A_s\otimes I - I\otimes B_s)\rho_{AB}^{1/2} = 0. \tag{5.1}$$

For unitaries impose the equations for the generators and their adjoints, with Bob's transpose supplying the opposite algebra. Define, over finite allocations,

$$f(\varepsilon) = \inf\left\{ \sum_j p_j c_{\pi_j} :\ p_j\ge 0,\ \sum_j p_j = 1,\ \sum_j p_j \ell_{\pi_j} \le \varepsilon \right\}. \tag{5.2}$$

Theorem 5.1 (PROVED: allocation of finite sectors)**.** *Among finite exactly matched states with defect at most ε,*

$$\inf E_{\mathrm{D}}^{\rightarrow}(\rho) = f(\varepsilon), \tag{5.3}$$

$$\inf I(A:B)_\rho = \inf_{p,\pi} \left[H(p) + 2\sum_j p_j c_{\pi_j} \right] =: C_{\mathrm{ex}}(\varepsilon). \tag{5.4}$$

Both infima can be approached with at most two sectors, and

$$2f(\varepsilon) \le C_{\mathrm{ex}}(\varepsilon) \le 2f(\varepsilon) + 1. \tag{5.5}$$

If arbitrarily small positive losses are feasible, then, for $\varepsilon > 0$,

$$f(\varepsilon) = \sup_{\lambda \ge 0} \left[\inf_{\pi} (c_\pi + \lambda \ell_\pi) - \lambda \varepsilon \right]. \tag{5.6}$$

Proof. For a support vector of ρ, write its coefficient matrix as $X : \overline{\mathcal{H}_B} \to \mathcal{H}_A$. Matching gives $A_s X = X B_s^T$ for all generators and adjoints. Finite-dimensional $*$-representations decompose into irreducibles. Schur's lemma gives

$$X = \bigoplus_j I_{V_j} \otimes X_j$$

on the common equivalent sectors, with multiplicity spaces on each side. Thus the support lies in $\bigoplus_j |\Phi_{n_j}\rangle \otimes K_j^A \otimes K_j^B$, where $n_j = \dim V_j$. Both parties can locally measure the common label and discard their multiplicities, obtaining

$$\omega = \bigoplus_j p_j |jj\rangle\langle jj| \otimes \Phi_{n_j}. \tag{5.7}$$

Every tested statistic is preserved: controls are block diagonal, act trivially on multiplicities, and on Φ_{n_j} their marginal is I_{n_j}/n_j. This is the finite trace structure familiar from synchronous correlations [26]; the resource comparison uses the explicit local reduction above.

The label is locally readable on both sides, so typical frequencies of labels and entanglement concentration give $E_{\mathrm{D}}^{\rightarrow}(\omega) \ge \sum_j p_j \log_2 n_j$. For the converse use the separable competitor obtained by replacing each Φ_n by $n^{-1}\sum_k |kk\rangle\langle kk|$. Its relative entropy from ω is exactly $\sum_j p_j \log_2 n_j$. On N copies the product competitor has N times this value. Data processing and the maximally entangled projection test used in the proof of (4.5) bound every distillation rate by it. Hence equality holds. Direct calculation gives

$$S(\omega_A) = S(\omega_B) = H(p) + \sum_j p_j \log_2 n_j, \quad S(\omega) = H(p),$$

and therefore $I(A:B)_\omega = H(p) + 2\sum_j p_j \log_2 n_j$. Local channels cannot increase either resource. This gives lower bounds on the original mixed state, rather than only on a replacement state. Conversely every finite allocation is realized by (5.7) with blockwise controls. Its defect is the tracial average $\sum_j p_j \ell_{\pi_j}$. This proves (5.3)–(5.4).

For a fixed finite menu, feasible weights form a compact polytope with normalization and one loss inequality. Every vertex has at most two nonzero coordinates. The entanglement objective is linear and the mutual-information objective is concave. Each has some vertex with value no greater than any prescribed feasible weight vector: express that vector as a convex combination of vertices and use linearity or concavity. Apply this to each finite menu and then take infima. Since a binary label has entropy at most one, (5.5) follows.

Every $\lambda \geq 0$ gives the lower bound in (5.6) by averaging $c_\pi + \lambda \ell_\pi$. For the reverse, f is finite, convex and nonincreasing on $(0, \infty)$: feasibility at smaller positive losses guarantees finiteness there, and mixing finite allocations proves convexity. A subgradient $s \leq 0$ at ε exists. Put $\lambda = -s$. Since $c_\pi \geq f(\ell_\pi)$, the supporting line gives $c_\pi + \lambda \ell_\pi \geq f(\varepsilon) + \lambda \varepsilon$. A zero-loss sector can be handled by taking the limit from positive losses. Infimize over π to obtain the reverse bound. □

For synchronous POVMs, perfect agreement on the same question implies (5.1). Indeed, for corresponding effects A, B, the agreement equalities give $\operatorname{Tr} \rho A = \operatorname{Tr} \rho B = \operatorname{Tr} \rho (A \otimes B)$. Since $A^2 \leq A$ and $B^2 \leq B$, the expectation of $(A \otimes I - I \otimes B)^2$ is at most zero. It is nonnegative, so vanishes; equality also makes the effects projections on their marginal supports. For a synchronous game with projections p_x^a, take

$$h_G = \sum_{\text{losing } x,y,a,b} \mu(x,y) p_y^b p_x^a p_y^b.$$

Its normalized trace is exactly the game loss. Theorem 2.1 restricted to this class and (5.5) give $w_{\text{local,ex}} = 2f + O_G(1)$ and $w_{\text{joint,ex}} = O_G(1)$.

Corollary 5.2 (Proved: finite-dimensional obstruction)**.** *Suppose the admissible bounded matrix tuples in each dimension form a closed set and no finite-dimensional sector has zero defect. For the minimum tracial defect δ_D over dimensions at most D,*

$$f(\varepsilon) \geq \sup_D (1 - \varepsilon/\delta_D)_+ \log_2(D+1), \qquad f(\varepsilon) \longrightarrow \infty.$$

Proof. The finite union of admissible matrix-tuple sets is compact. Its nonnegative continuous defect has a positive minimum, for zero defect would give a forbidden representation. A feasible allocation puts weight at most ε/δ_D in dimensions at most D. The remaining weight costs at least $\log_2(D+1)$. First send ε to zero for fixed D, then let D grow. □

If an enlarged algebra $\mathcal{B}$ has unital $*$-homomorphisms $i : \mathcal{A} \to \mathcal{B}$ and $r : \mathcal{B} \to \mathcal{A}$ with $ri = \mathrm{id}$ and defect $i(h)$, restriction along i and extension along r preserve states, matching and tested statistics. The two directions give equality of all resource infima. This proves that an untested spectator with such a retraction adds no price. Exact matching is essential; Appendix C explains the stability gap.

6 Calibration and a finite game floor

Proposition 6.1 (Proved: the erasure calibration)**.** *Let $0 \leq p \leq 1/2$, $E = h_2(p)$, and give three sources the same marginals $\rho_A = \rho_B = \operatorname{diag}(1-p, p)$:*

$$|\psi_p\rangle = \sqrt{1-p}|00\rangle + \sqrt{p}|11\rangle, \quad \rho_C = (1-p)|00\rangle\langle 00| + p|11\rangle\langle 11|, \quad \rho_P = \rho_A \otimes \rho_B.$$

Erasing A while preserving B costs, respectively, $-E, 0, E$ with joint access to B, and E, E, E after B is inaccessible. The release penalties are $2E : E : 0$, equal to their mutual informations. For the specified pure family its optimal CHSH value $\mathcal{B}$ determines the penalty as

$$\Delta w = 2h_2\left(\frac{1 - \sqrt{2 - \mathcal{B}^2/4}}{2}\right).$$

Proof. The final joint entropy is E and the three initial joint entropies are $0, E, 2E$, giving the held lower bounds. They are attained by, respectively, unitarily making ψ_p a product blank and reversibly mixing B, a CNOT from B to A, and ordinary erasure of the product A. If

the classical source has an external purifier, dephase B after the CNOT to restore its original reference marginal. For the quantum source no such reference is present, and the product operation preserves all external correlations of B. After release only the identical marginal of A is available, giving cost E in all cases. These are elementary instances of conditional erasure [8].

Set $s = 2\sqrt{p(1-p)}$. The correlation matrix of ψ_p is $T = \mathrm{diag}(s, -s, 1)$. The observables

$$A_0 = Z, \quad A_1 = X, \quad B_0 = (Z + sX)/\sqrt{1+s^2}, \quad B_1 = (Z - sX)/\sqrt{1+s^2}$$

give $\mathcal{B} = 2\sqrt{1+s^2}$. For optimality, write the sum and difference of Bob's two unit directions as $2\cos\phi\, u$ and $2\sin\phi\, v$ with $u \perp v$. Optimizing Alice and then ϕ gives at most $2\sqrt{\|Tu\|^2 + \|Tv\|^2} \le 2\sqrt{1+s^2}$; this is the argument of the Horodecki criterion [19]. Extremal nonconstant qubit observables have such directions; a constant observable gives no violation. Solving for p proves the formula. An unrestricted mixed device with the same score need not have this work cost. □

Proposition 6.2 (PROVED: a paired precision experiment)**.** *Use the same pure source of entropy E and the same coherent measurement and copy implementation in two arms. Release the same answers, then reset the residual devices locally in one arm and jointly in the other. Their optimal reversible reset-work difference is*

$$\Delta w = \sum_{xy} \mu(x,y) I(A' : B')_{\sigma^{xy}}, \qquad 2E - \log_2 K \le \Delta w \le 2E. \tag{6.1}$$

Thus the flat strategies supplied by [13, Theorem 11.9(b)] give $\Delta w = 2\alpha L + O_G(1)$.

Proof. At fixed questions local reset costs $S(A') + S(B')$, while joint reset costs $S(A'B')$. No signalling makes each local marginal independent of the other question. This proves the averaged equality. The retained joint state has at most K orthogonally flagged pure branches, so $S(A'B') \le \log_2 K$. Each local retained state is an isometric image of its source marginal followed by pinching, which cannot decrease entropy, so both local entropies are at least E. Conversely the retained state comes from the source by local channels; mutual-information data processing bounds it by $2E$. Preparation and answer-port work cancel between the arms. □

Theorem 6.3 (PROVED: the magic-square floor)**.** *For the standard uniform row-column binary magic-square game, with four parity-allowed answers for each question and thermal four-dimensional answer ports,*

$$w_{\mathrm{local,MS}}(0) = 4.$$

Proof. First dilate all local POVMs to projective measurements on finite local extensions initialized in pure product ancillas. This preserves the score and the one-way distillable entanglement in either direction, since appending or discarding a pure local ancilla is a local operation. It therefore suffices to prove the lower bound for these dilated measurements.

Perfect intersection agreement makes the corresponding row and column reflections agree on every support vector. The marginal supports are invariant under them, because moving a reflection to the other party in a purification leaves its marginal invariant. Row parity and row commutation hold by the joint projective measurement; the column relations transfer from Bob using the mirrors. Thus on Alice's support they represent the magic-square solution group with marked element $J = -I$. Its identification with the two-qubit Pauli group is [5, Proposition 6.10]. Its $J = -I$ representations are copies of the four-dimensional irreducible: one anticommuting reflection pair gives X, Z on one qubit by diagonalizing Z and pairing its eigenspaces with X, and the commuting second pair gives another such qubit on the

multiplicity space. The same holds on Bob. Matching then has the Schur form in the proof of Theorem 5.1; it locally extracts Φ_4 even from a mixed state. Hence $E_{\mathrm{D}}^{\rightarrow} \geq 2$, and Theorem 2.1 gives $w \geq 4$. This also recovers the established two-singlet rigidity [31, Appendix C, Lemma C.1 and its application].

For attainment form Φ_4 on two working qubits at each party by a global unitary. At each question rotate the local source into the required rank-one joint eigenbasis, swap its four-dimensional register into the thermal answer port, dephase that port in the answer basis and release it. The working register is now $I_4/4$ and can be erased at cost two. Both parties spend four units in total. The answer ports received the actual measured source; no pure answer blanks were supplied for free. All working registers return pure. □

Proposition 6.4 (Proved: finite-error conversion)**.** *Suppose local extraction from a strategy of loss ε gives a state ω on $\mathbb{C}^4 \otimes \mathbb{C}^4$ with $\frac{1}{2}\|\omega - \Phi_4\|_1 \leq t_{\mathrm{MS}}(\varepsilon) \leq 1/2$. Then*

$$w_{\mathrm{local}}(\varepsilon) \geq \max\{0, 4 - 2t_{\mathrm{MS}}(\varepsilon)\log_2 45 - 4h_2(t_{\mathrm{MS}}(\varepsilon))\}. \tag{6.2}$$

Proof. Writing $t = \frac{1}{2}\|\omega - \Phi_4\|_1$, finite-dimensional entropy continuity gives

$$S(\omega_B) \geq 2 - t\log_2 3 - h_2(t), \qquad S(\omega) \leq t\log_2 15 + h_2(t).$$

The hashing inequality [10, Theorem 10] lower-bounds one-way distillation by their difference. Local extraction cannot increase $E_{\mathrm{D}}^{\rightarrow}$ of the original state. Apply Theorem 2.1. On $[0, 1/2]$ the error expression is increasing, so replacing t by its upper bound is valid; also use $w \geq 0$. □

For the equation-variable binary magic-square game, Coladangelo–Stark [5, Definition 4.14 and Theorem 6.9, arXiv v2] supplies such an extraction modulus tending to zero, including mixed states. Divide their full trace-norm modulus by two before using (6.2). For the row-column game use [31, Appendix C, Lemma C.1 and equations (C18)–(C35), arXiv v2]: each context loss is at most the mean loss divided by the smallest question-pair probability, and for a mixed strategy the same local isometries, applied to a purification whose reference is left untouched, give an extraction channel once the junk is discarded. Each formulation uses its own modulus; no common numerical constant and no numerical score-to-heat threshold is asserted.

Sketch**: laboratory realization.** Interleave game-calibration and erasure trials with independently calibrated sources, thermal ports and reset reservoirs. Compare the work difference in (6.1), or the four-unit ideal magic-square floor. A hardware-specific reversible control sequence, measured bath heat and reset of every controller record are still needed. The finite computations in `verify/` establish the algebra and entropy values, not experimental attainability.

The coefficient α alone gives no accessible onset scale. Given G, test it with probability $p > 0$ and otherwise ask a distinct dummy question with deterministic perfect answers. For $\varepsilon < p$, restriction and extension of strategies give $d_{G_p}(\varepsilon) = d_G(\varepsilon/p)$ and $E_{G_p}(\varepsilon) = E_G(\varepsilon/p)$. The coefficient is unchanged. For $\varepsilon \geq p$, a product strategy wins all dummy rounds and has total loss at most p. Since p can be arbitrarily small with only a fixed increase in question types, an actual logarithmic experiment requires finite-error game constants. Also, (2.3) controls large relative precision changes: for fixed $r > 1$, $[w_{\mathrm{local}}(\varepsilon^r) - w_{\mathrm{local}}(\varepsilon)]/L \to 2\alpha(r-1)$. It does not control every halving increment.

7 Operational access and prior work

Proposition 7.1 (PROVED: marginals do not specify a reset price)**.** *The pure state* $|W\rangle = (|100\rangle+|010\rangle+|001\rangle)/\sqrt{3}$ *is uniquely determined among all states by its AB and BC marginals. Nevertheless erasing A using AB while preserving the full BC state, followed by erasure of BC, has excess work over joint erasure*

$$I(A:C|B)_W = h_2(1/3) > 0.$$

Proof. Both two-qubit marginals are $\frac{1}{3}|00\rangle\langle 00|+\frac{2}{3}|\psi^+\rangle\langle\psi^+|$. Positivity forces every compatible state into the intersection of their lifted supports. Direct coefficient comparison gives $\mathrm{span}\{|000\rangle, |W\rangle\}$. If $a = \langle W|\sigma|W\rangle$, its AB weight on ψ^+ is $2a/3$. Thus $a = 1$, and positivity and trace one force $\sigma = |W\rangle\langle W|$.

The first reset costs at least $S(AB)-S(B) = S(A|B)$ by the marginal AB entropy balance. Conditional erasure attains it while preserving BC [8, Corollary 2 and Appendix A]; this preservation includes the reference C. The second reset costs $S(BC)$. Subtracting the joint cost $S(ABC)$ gives $I(A:C|B)$. All one-qubit and complementary two-qubit entropies of W are $h_2(1/3)$. Hence the excess is as claimed. □

In particular, the maximum compatible global entropy minus the actual entropy is zero here, yet the schedule surcharge is positive. Inference from overlapping marginals cannot replace the access and preservation rules of a physical operation. The uniqueness phenomenon itself belongs to the marginal-reconstruction literature [22]; the example here applies it to a reset deadline. Related operational roles of lost correlations appear in the thermodynamics of modular processing [4] and in deconstruction and conditional erasure [2, 3]. The latter identifies conditional mutual information with the optimal rate of noise that decouples a system while negligibly disturbing a specified marginal; the reset schedule here is a different task.

What is already known. Restricted joint access has long had a thermodynamic price. Oppenheim and the Horodeckis introduced work deficit; its pure-state value is the entanglement entropy [25]. This is the sufficiently classical-assisted regime here. Conditional erasure, including negative work, is due to del Rio and collaborators [8]. The elementary $2:1:0$ comparison is an illustration of that theory. Entropy-production identities are due to Esposito, Lindenberg and Van den Broeck and to Reeb and Wolf [16, 27]. None of these ingredients is claimed as new.

Devetak's local purity distillation [9] and Krovi–Devetak's bounded classical communication [20] already study distributed purity and communication tradeoffs. Horodecki and collaborators bound the work deficit and the cost of erasing entanglement below by E_{R}, using the potential on which Lemma 4.2 is built [18]. The lemma accounts explicitly for local Gibbs baths and quantum transfer; the contribution claimed here is its application to the game-level recovery bounds. Groisman, Popescu and Winter identify mutual information with the minimum asymptotic rate of local noise needed to erase all correlations of a bipartite state [17]. These are correlation-cost tasks; the present task also includes preparation, answer release and cyclic reset of a device that must produce a successful game strategy. Our pure-state recovery construction sits in this framework; a communication-work framework by itself is not a novelty claim.

Alimuddin and collaborators study helper-assisted erasure and exclusive control, including Bell certification [1]. Their target is a specified memory and competing helpers. Here the target includes preparation and cyclic reset of the entire game apparatus, minimized over successful strategies. The required apparatus correlations can diverge with fixed answer

alphabets. The distinction concerns the thermodynamic boundary, rather than a new principle of assisted erasure.

Ćwikliński values supplied XOR side information in a Szilard engine [7], and routes a supplied excitation into a battery on game success [6]. Earlier, Maruyama, Morikoshi and Vedral used work extraction to detect entanglement [23]. These are downstream valuations of supplied correlations or behavior with explicit trusted transducers. They do not establish the upstream minimum cyclic work over game-producing devices studied here. Broader proposals connecting Bell behavior and Landauer cost include [29]; no equivalence with the present cyclic task is inferred from a shared vocabulary.

What this paper adds. Relative to those sources, the claims advanced here are the game-optimized correlation sandwich, its divergent sharp coefficient for a fixed finite game, the game-level recovery bounds, the paired precision-dependent gap, and the exact four-unit cyclic magic-square construction. The exact-matching sector theorem combines the known trace structure with original-state resource comparisons. The multiphase example in Appendix B supplies a further calibrated price. These are mathematical contribution claims under the stated contracts.

Conjecture**: stability preserving all resources.** A useful extension would turn approximately matched mixed strategies into sector allocations with quantitative control of loss, the original state's one-way distillable entanglement, and the schedule's work. No such unrestricted equality is proved here. Appendix C proves an entropy sandwich with a precision gap and rules out fixing the same controls by fiat. Almost-synchronous rounding [30] alone does not supply the original-state entanglement control this extension needs.

How this work was done

Roles. This paper is the work of its two authors together with current frontier AI models from more than one developer. The authors set the direction throughout: they chose the questions, fixed the scope and the standards of proof and review, decided which lines to pursue and which to stop, and take full responsibility for every claim. The models did most of the exploration, derivation, computation and drafting, in a sequence of short research sprints. For this paper the models first explored exchange rates, recovery schedules and operational access. A subsequent proof audit reconstructed the derivations and checked their scope and cited inputs. Two independent blind audits, by models of another developer that had not taken part in the work, then read the manuscript refutation first, reran the checks and looked up the references in primary records. They found no error in a main result. They asked for a missing hypothesis in Theorem A.1, credit to related work on the cost of erasing correlations, corrected equation and source pointers, and narrower descriptions of the checks; all are applied here.

Validation. The evidence for the results is of two kinds: complete written proofs and computations by the scripts in `verify/`, which the reader can check.

- *Proofs.* Every theorem is proved here or supplied with an exact result citation. The imported results are the published game-demand laws; the schedule bounds used here are proved directly. The potential lemma accounts for discarded bath records, dephased communication and quantum transfer. The finite-error result states its extraction input explicitly.

- *Proof audit.* The audit checked the input and reference boundary of the game entropy ledger, the order of limits, the recovery communication budgets, the domain for substituting an extraction-error bound, and the versions of the cited sources. The unrestricted sector stability extension remains open.
- *Computation.* `verify/run_all.py` checks the three erasure penalties, the CHSH family, every magic-square question and outcome, the working-register entropy budget of the swap construction, finite communication tradeoffs, sector allocation against linear programming and dual witnesses, weighted commutator constructions, multiphase clocks, and the two access obstructions. The swap channel itself is not simulated, the transfer-bound checks use pure states, and no finite test evaluates the relative entropy of entanglement of a mixed state. Five negative controls are rejected: a product source assigned a quantum penalty and a wrong magic-square parity, and three sanity assertions (two fixed arithmetic comparisons and the positivity of $h_2(1/3)$) that do not simulate a faulty protocol. Scripts run sequentially with one numerical thread, fixed seeds and small matrices. The recorded logs give the actual environment and resource use. Finite checks support the proofs and do not replace them.

No argument here has been formally verified; the proofs and computations are offered for checking as they stand.

Priority. Bounded searches of primary literature through 27 October 2026 covered work deficit, conditional erasure, the noise cost of erasing correlations, local purity distillation, assisted-erasure exclusivity, Bell-related thermodynamic valuations, and the mathematical tools cited at their use. The comparisons in Section 7 distinguish these established ingredients from the game-level deductions. The search does not establish exhaustive publication priority.

A Weighted Weyl tests for several observers

This appendix evaluates Weyl-relation defects against a family of weighted probes. Probe entropy alone is not identified with usable entanglement.

Theorem A.1 (Proved: the observer-family defect)**.** *Fix $q \in \mathbb{C}$ with $|q| = 1$ and an integer $d \geq 1$. Let $\mathcal{K}$ be a compact convex set of density matrices on $\mathbb{C}^d$, containing a positive-definite matrix. Define inverse trace as $+\infty$ on singular matrices and set*

$$S_{\mathcal{K}} = \min_{\sigma \in \mathcal{K}} \operatorname{Tr} \sigma^{-1}, \qquad \delta = \operatorname{Arg}(q^d) \in (-\pi, \pi].$$

For unitary gates U, V, put $C = qV^*U^*VU$ and $H = -i \operatorname{Log} C$, with principal eigenphases in $(-\pi, \pi]$. Then

$$\inf_{U,V} \max_{\rho \in \mathcal{K}} \operatorname{Tr} \rho H^2 = \frac{\delta^2}{S_{\mathcal{K}}}.$$

With $q = e^{2\pi i \alpha}$ for $\alpha \in \mathbb{R}$ and $\tau_d = \|d\alpha\|_{\mathbb{R}/\mathbb{Z}}$,

$$\frac{4\tau_d}{\sqrt{S_{\mathcal{K}}}} \leq \inf_{U,V} \max_{\rho \in \mathcal{K}} \|(UV - qVU)\sqrt{\rho}\|_2 \leq \frac{2\pi\tau_d}{\sqrt{S_{\mathcal{K}}}}. \tag{A.1}$$

Proof. Since $|q| = 1$, C is unitary and H is Hermitian. Let σ_* minimize inverse trace. It is positive definite, since inverse trace diverges at singular boundaries. The determinant forces $\operatorname{Tr} H = \delta + 2\pi k$. Weighted Hilbert–Schmidt Cauchy–Schwarz gives

$$|\operatorname{Tr} H|^2 \leq (\operatorname{Tr} \sigma_*^{-1})(\operatorname{Tr} \sigma_* H^2),$$

so the worst observer cost is at least $\delta^2/S_{\mathcal{K}}$. First-order optimality in the direction $\rho - \sigma_*$ gives

$$-\operatorname{Tr}\sigma_*^{-2}(\rho-\sigma_*) \geq 0, \qquad \operatorname{Tr}\rho\sigma_*^{-2} \leq S_{\mathcal{K}}.$$

Consequently $H_* := \delta\sigma_*^{-1}/S_{\mathcal{K}}$ has $\operatorname{Tr} H_* = \delta$, $\|H_*\| \leq |\delta| \leq \pi$, and worst cost at most $\delta^2/S_{\mathcal{K}}$.

It remains to realize H_*. Every $W \in SU(d)$ is a unitary commutator. Diagonalize $W = S\operatorname{diag}(w_j)S^*$, choose cyclic phases a_j with $a_j/a_{j-1} = w_j$ (possible because $\prod_j w_j = 1$), and let $T|j\rangle = |j+1\rangle$. For $A = \operatorname{diag}(a_j)$, $ATA^*T^* = \operatorname{diag}(w_j)$. Thus $U = ST^*S^*$ and $V = SA^*S^*$ satisfy $V^*U^*VU = W$. Apply this with $W = q^{-1}e^{iH_*}$, whose determinant is one. This attains the bound, including $d = 1$.

Finally $UV - qVU = UV(I - C)$, and on principal phases

$$\frac{4}{\pi^2}H^2 \leq (I-C)^*(I-C) \leq H^2.$$

Take weighted traces, maxima and infima and use $|\delta| = 2\pi\tau_d$ to obtain (A.1). □

If $I/d \in \mathcal{K}$, then $S_{\mathcal{K}} = d^2$, since $\operatorname{Tr}\sigma^{-1} \geq d^2$ for every density matrix. Thus testing all cyclic translates of a diagonal probe restores the uniform optimum. If $\mathcal{K} = \{\rho\}$, (A.1) gives the single-probe bound.

The dimension envelope needs arithmetic hypotheses. For badly approximable α, $\tau_d \geq a/d$; Dirichlet's pigeonhole bound supplies $d \leq Q$ with $\tau_d \leq 1/Q$, hence also $d \geq aQ$. Equation (A.1) then gives uniform-probe dimension $\Theta(e^{-1/2})$. Irrationality alone is insufficient. For a truncated thermal probe with weights proportional to $e^{-\beta j}$, $j = 0, \ldots, d-1$, and fixed $\beta > 0$, inverse trace is $\Theta_\beta(e^{\beta d})$. Under the uniform asymptotic condition $-\ln\tau_d = o(d)$, (A.1) gives $\ln(1/e_d) = \beta d/2 + o(d)$, hence the least dimension is $(2/\beta)\ln(1/e) + o(\ln(1/e))$. Neither envelope follows from irrationality without the stated qualification.

B Several phase controls sharing one shift

Consider unitary controls $U_1, \ldots, U_m, V$ with exact bipartite matching. Their tested squared defect in marginal state ρ is

$$\sum_{j=1}^{m} \|(U_jV - e^{2\pi i\alpha_j}VU_j)\sqrt{\rho}\|_2^2 \leq e^2.$$

No extra exact relation is imposed among the U_j.

Theorem B.1 (Proved: a multiphase price)**.** *The minimum defect on a uniform d-dimensional probe is*

$$\ell_d = 4\sum_{j=1}^{m} \sin^2\left(\frac{\pi\|d\alpha_j\|_{\mathbb{R}/\mathbb{Z}}}{d}\right). \tag{B.1}$$

The exact-matching entanglement and correlation prices are

$$E_m(e) = \inf_{\sum_r p_r\ell_{d_r} \leq e^2} \sum_r p_r \log_2 d_r, \tag{B.2}$$

$$C_m(e) = \inf_{t\ell_d + (1-t)\ell_{d'} \leq e^2} [h_2(t) + 2t\log_2 d + 2(1-t)\log_2 d'], \tag{B.3}$$

where allocations are finite probabilities and $0 \leq t \leq 1$. For $\theta = 2^{1/(m+1)}$ and $\alpha_j = \theta^j$,

$$D_m(e) = \Theta_m(e^{-m/(m+1)}), \qquad E_m(e) = \frac{m}{m+1}\log_2(1/e) + O_m(1), \tag{B.4}$$

$$C_m(e) = \frac{2m}{m+1}\log_2(1/e) + O_m(1). \tag{B.5}$$

Here D_m is the least uniform-probe dimension. The same minimum dimension is possible among exactly matched finite strategies.

Proof. For one relation, let $\theta_k \in (-\pi, \pi]$ be the principal eigenphases of its corrected commutator. Their sum is the forced determinant phase modulo 2π. Set $a = 2\pi\|d\alpha\|_{\mathbb{R}/\mathbb{Z}} \le \pi$. Taking absolute values and then trimming their sum to a cannot increase $\sum_k (1 - \cos\theta_k)$. For $u, v \ge 0$ with $u + v \le \pi$,

$$(1 - \cos u) + (1 - \cos v) \ge 2\left(1 - \cos\frac{u+v}{2}\right),$$

because $\cos u + \cos v = 2\cos((u+v)/2)\cos((u-v)/2)$. Repeated pair averaging therefore minimizes the sum at all phases equal to a/d. The uniform squared norm is at least $4\sin^2(a/(2d))$.

All these minima have a common construction. Take nearest integers p_j to $d\alpha_j$ and set $V|k\rangle = |k+1 \bmod d\rangle$ and $U_j|k\rangle = e^{2\pi i p_j k/d}|k\rangle$. Then $U_j V = e^{2\pi i p_j/d} V U_j$, giving (B.1). Every irreducible sector of dimension d has defect at least ℓ_d; the construction on Φ_d attains it, whether or not those controls are irreducible. Theorem 5.1 gives the lower bounds in (B.2)–(B.3), and these clock strategies and their flagged mixtures give the upper bounds. Two-sector sufficiency follows from the same theorem. A mixture with all sector dimensions at most D has average defect at least the smallest ℓ_d for $d \le D$, which proves the dimension assertion.

For the exponents put $n = m + 1$ and $M(x_0, \dots, x_m) = (x_1, \dots, x_m, 2x_0)$. Its irreducible characteristic polynomial is $x^n - 2$, by Eisenstein, and $w = (1, \theta, \dots, \theta^m)$ satisfies $Mw = \theta w$. Given $d \ge 1$, choose nearest integers p_j and set $v = (d, p_1, \dots, p_m) = dw + z$, where $\|z\| \le \sqrt{m}\,\Delta$ and $\Delta = \max_j \|d\theta^j\|_{\mathbb{R}/\mathbb{Z}}$. The integer determinant $\det[v, Mv, \dots, M^m v]$ is nonzero. Otherwise a rational polynomial of degree at most m annihilates v; its Bezout identity with $x^n - 2$, applied to M, forces $v = 0$.

Subtract $\theta^k v$ from column k for $1 \le k \le m$. The remaining column is $(M^k - \theta^k I)z$, of norm $O_m(\Delta)$, while the first column is $O_m(d)$. Hadamard's inequality yields

$$1 \le |\det[v, Mv, \dots, M^m v]| \le C_m d\Delta^m, \qquad \Delta \ge a_m d^{-1/m}. \tag{B.6}$$

This is the classical algebraic simultaneous-approximation mechanism associated with Perron; the determinant proof is included to specify the needed input.

Pigeonholing $Q^m + 1$ orbit points in Q^m cubes gives a denominator $d \le Q^m$ with $\Delta \le 1/Q$. Equation (B.6) forces $d \ge (a_m Q)^m$. The inequalities $2x/\pi \le \sin x \le x$ on $[0, \pi/2]$ give, respectively,

$$\sqrt{\ell_d} \ge 4a_m d^{-1-1/m}, \qquad \sqrt{\ell_d} \le 2\pi\sqrt{m}\,\Delta/d \le C'_m Q^{-m-1}$$

for that denominator. Choose Q proportional to $e^{-1/(m+1)}$ for the matching dimension bound. Finally $\ell_d \ge b_m d^{-r}$ with $r = 2 + 2/m$ implies

$$\sum_j p_j \log_2 d_j \ge \frac{1}{r}\left(\log_2 b_m - \sum_j p_j \log_2 \ell_{d_j}\right) \ge \frac{1}{r}\log_2 \frac{b_m}{\sum_j p_j \ell_{d_j}} \ge \frac{m}{m+1}\log_2(1/e) + O_m(1).$$

Jensen's inequality is used in the middle step. The single-sector construction matches it, and (5.5) gives the correlation price. □

For the cube-root pair $(\sqrt[3]{2}, \sqrt[3]{4})$ the three coefficients are $2/3$, $2/3$, and $4/3$. Reversible formation followed by separate reset, with no game transcript, costs exactly the prepared state's mutual information. Thus the last coefficient is also a cyclic local-reset correlation-work coefficient for this exactly matched calibration task. A coherent comparison of $U_j V$ and $e^{2\pi i\alpha_j} V U_j$ has minus-port probability one quarter of the tested squared norm. This predicts a gate-calibration curve; imposing exact matching experimentally and performing reversible reset remain part of the implementation problem. It is not an unrestricted nonlocal-game demand theorem.

C Approximate invariance and its limit

Proposition C.1 (PROVED: the entropy precision gap)**.** *Let U_s be free unitary generators, and let a finite word list $\mathcal{W}$ contain every relation to be tested. Suppose*

$$\|[U_s, \sqrt{\rho}]\|_2 \le e, \qquad \|(w(U) - \zeta_w I)\sqrt{\rho}\|_2 \le e, \quad |\zeta_w| = 1.$$

Define sector loss by the unnormalized sum of squared word defects, and let $s(e)$ be the least feasible $S(\rho)$. For $0 < e \le 1$,

$$f(C_{\mathcal{W}} e) \le s(e) \le f(e^2) + 1, \qquad C_{\mathcal{W}} = 2|\mathcal{W}| + 4 \sum_{w \in \mathcal{W}} |w|^2.$$

No additional exact relations beyond unitarity are assumed.

Proof. In an eigenbasis of ρ with eigenvalues λ_i, put $P_t = 1_{\{\rho \ge t\}}$ and $r(t) = \operatorname{rank} P_t$. Then $\rho = \int P_t dt$ and $\int r(t)dt = 1$. For a unitary U,

$$\begin{aligned} \int \|(I - P_t)UP_t\|_2^2 dt &= \sum_{ij} (\lambda_j - \lambda_i)_+ |U_{ij}|^2 \\ &= \tfrac{1}{2} \sum_{ij} |\lambda_j - \lambda_i| |U_{ij}|^2 \le \|[U, \sqrt{\rho}]\|_2. \end{aligned}$$

The second equality uses the vanishing signed sum, from unitarity. For the inequality factor the eigenvalue difference into the difference and sum of their square roots, and use Cauchy–Schwarz; the squared sum factor has total at most four.

Replace $P_t U_s P_t$ by a unitary polar completion $W_{s,t}$ on $P_t\mathcal{H}$. Its singular values $a_i \le 1$ obey $\sum_i (1 - a_i)^2 \le \sum_i (1 - a_i^2)$. Including leakage outside P_t gives

$$\|U_s P_t - W_{s,t} P_t\|_2^2 \le 2\|(I - P_t)U_s P_t\|_2^2.$$

The corresponding inverse-letter estimate follows from the same singular values and equality of forward and backward leakage traces. Telescoping a word of length L and using Cauchy–Schwarz over its L terms therefore gives integrated squared error at most $2L^2 e$. The original word residual has integrated square at most e^2. Squaring their sum bounds the rounded word defect by $2e^2 + 4L^2 e$ after integration.

Choose the uniform state $P_t/r(t)$ with probability $r(t)dt$. Entropy concavity gives $S(\rho) \ge \int r(t) \log_2 r(t) dt$. The rounded controls define finite representations of the free unitary algebra; decomposing them into irreducibles preserves tracial loss and lowers the average logarithmic dimension. The spectral projections change only finitely many times, so this is a finite allocation. Its aggregate loss is at most $C_{\mathcal{W}} e$ and its cost at most $S(\rho)$. This proves the lower bound. For the upper bound choose an allocation approaching $f(e^2)$ with at most two sectors and put $\rho = \bigoplus_j p_j I_{n_j}/n_j$. It commutes exactly with every control, each individual squared defect is at most e^2, and $S(\rho) = H(p) + \sum_j p_j \log_2 n_j \le 1 + f(e^2) + o(1)$. □

Proposition C.2 (PROVED: fixed-control rounding can cost unbounded entropy)**.** *Let U be diagonal with distinct phases, $|s\rangle = d^{-1/2} \sum_{k=0}^{d-1} |k\rangle$, $V = I + (e^{it} - 1)|s\rangle\langle s|$, and $\rho = |0\rangle\langle 0|$, where $t \notin 2\pi\mathbb{Z}$. The only state commuting with both U and V is I/d, but*

$$S(\rho) = 0, \quad [U, \sqrt{\rho}] = 0, \quad \|[V, \sqrt{\rho}]\|_2^2 = \frac{4(1 - \cos t)(d - 1)}{d^2} \longrightarrow 0.$$

Proof. A matrix commuting with U is diagonal. If it also commutes with $|s\rangle\langle s|$, all its diagonal entries agree, proving the first claim. The commutator with $|0\rangle\langle 0|$ has $2(d - 1)$ off-diagonal entries of modulus $|e^{it} - 1|/d$, giving the formula. The same controls and their transposes have vanishing matching error on $|00\rangle$, while exact matching forces the one-dimensional invariant-vector space spanned by Φ_d, by irreducibility and Schur's lemma. □

Thus no dimension-independent entropy guarantee can enforce exact invariance while keeping these controls. The example does not satisfy a game's obstructed relators, so it does not refute the published entanglement demand theorem. Proposition C.1 also concerns only representation-state entropy, not the one-way distillable entanglement of an arbitrary original mixed strategy. These are two separate reasons to leave the unrestricted stability extension open.

References

[1] Mir Alimuddin, Nathan Shettell, Raja Yehia, Antonio Acín, and Federico Centrone. Exclusive control of quantum memory erasure. *Physical Review Research*, 8:043078, 2026. URL: https://arxiv.org/abs/2512.05761, doi:10.1103/6jtz-g59p.

[2] Mario Berta, Fernando G. S. L. Brandão, Christian Majenz, and Mark M. Wilde. Conditional decoupling of quantum information. *Physical Review Letters*, 121:040504, 2018. URL: https://arxiv.org/abs/1808.00135, doi:10.1103/PhysRevLett.121.040504.

[3] Mario Berta, Fernando G. S. L. Brandão, Christian Majenz, and Mark M. Wilde. Deconstruction and conditional erasure of quantum correlations. *Physical Review A*, 98:042320, 2018. URL: https://arxiv.org/abs/1609.06994, doi:10.1103/PhysRevA.98.042320.

[4] Alexander B. Boyd, Dibyendu Mandal, and James P. Crutchfield. Thermodynamics of modularity: structural costs beyond the Landauer bound. *Physical Review X*, 8:031036, 2018. URL: https://arxiv.org/abs/1708.03030, doi:10.1103/PhysRevX.8.031036.

[5] Andrea Coladangelo and Jalex Stark. Robust self-testing for linear constraint system games. arXiv:1709.09267v2, 2019. Corrected version; Theorem 6.9 and Proposition 6.10. URL: https://arxiv.org/abs/1709.09267v2.

[6] Piotr Ćwikliński. Battery-explicit thermodynamic witnesses of Bell post-quantumness. *New Journal of Physics*, 28:084509, 2026. URL: https://arxiv.org/abs/2605.09149v4, doi:10.1088/1367-2630/ae98a4.

[7] Piotr Ćwikliński. Thermodynamic value of XOR-game-induced side information in a Szilard engine: extremal binary-channel bounds. arXiv:2605.12044v4, 2026. URL: https://arxiv.org/abs/2605.12044v4.

[8] Lídia del Rio, Johan Åberg, Renato Renner, Oscar Dahlsten, and Vlatko Vedral. The thermodynamic meaning of negative entropy. *Nature*, 474:61–63, 2011. Result labels refer to the arXiv text. URL: https://arxiv.org/abs/1009.1630, doi:10.1038/nature10123.

[9] Igor Devetak. Distillation of local purity from quantum states. *Physical Review A*, 71:062303, 2005. doi:10.1103/PhysRevA.71.062303.

[10] Igor Devetak and Andreas Winter. Distillation of secret key and entanglement from quantum states. *Proceedings of the Royal Society A*, 461:207–235, 2005. Theorem 10 in the arXiv text. URL: https://arxiv.org/abs/quant-ph/0306078, doi:10.1098/rspa.2004.1372.

[11] Seth Douglas and Nidhal Mghirbi. Only quantum channels save work by waiting, 2026. Unrefereed preprint. Result numbers refer to version 1.0.2, doi:10.5281/zenodo.23323044. doi:10.5281/zenodo.23310624.

[12] Seth Douglas and Nidhal Mghirbi. The price of keeping time, 2026. Unrefereed preprint. doi:10.5281/zenodo.23310632.

[13] Seth Douglas and Nidhal Mghirbi. The price of precision: usable entanglement in nonlocal games and the geometry of groups, 2026. Version 1.0.0, unrefereed preprint. Result numbers refer to this version. doi:10.5281/zenodo.23697346.

[14] Seth Douglas and Nidhal Mghirbi. The price of time order, 2026. Unrefereed preprint. Result numbers refer to version 1.0.2, doi:10.5281/zenodo.23323043. doi:10.5281/zenodo.23310622.

[15] Seth Douglas and Nidhal Mghirbi. The work cost of answering in real time, 2026. Unrefereed preprint. doi:10.5281/zenodo.23310628.

[16] Massimiliano Esposito, Katja Lindenberg, and Christian Van den Broeck. Entropy production as correlation between system and reservoir. *New Journal of Physics*, 12:013013, 2010. doi:10.1088/1367-2630/12/1/013013.

[17] Berry Groisman, Sandu Popescu, and Andreas Winter. Quantum, classical, and total amount of correlations in a quantum state. *Physical Review A*, 72:032317, 2005. URL: https://arxiv.org/abs/quant-ph/0410091, doi:10.1103/PhysRevA.72.032317.

[18] Michał Horodecki, Paweł Horodecki, Ryszard Horodecki, Jonathan Oppenheim, Aditi Sen(De), Ujjwal Sen, and Barbara Synak-Radtke. Local versus nonlocal information in quantum-information theory: formalism and phenomena. *Physical Review A*, 71:062307, 2005. Result labels refer to the arXiv text. URL: https://arxiv.org/abs/quant-ph/0410090, doi:10.1103/PhysRevA.71.062307.

[19] Ryszard Horodecki, Paweł Horodecki, and Michał Horodecki. Violating Bell inequality by mixed spin-$\frac{1}{2}$ states: necessary and sufficient condition. *Physics Letters A*, 200:340–344, 1995. doi:10.1016/0375-9601(95)00214-N.

[20] Hari Krovi and Igor Devetak. Local purity distillation with bounded classical communication. *Physical Review A*, 76:012321, 2007. doi:10.1103/PhysRevA.76.012321.

[21] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5:183–191, 1961. doi:10.1147/rd.53.0183.

[22] Noah Linden, Sandu Popescu, and William K. Wootters. Almost every pure state of three qubits is completely determined by its two-particle reduced density matrices. *Physical Review Letters*, 89:207901, 2002. doi:10.1103/PhysRevLett.89.207901.

[23] Koji Maruyama, Fumiaki Morikoshi, and Vlatko Vedral. Thermodynamical detection of entanglement by Maxwell's demons. *Physical Review A*, 71:012108, 2005. URL: https://arxiv.org/abs/quant-ph/0311083, doi:10.1103/PhysRevA.71.012108.

[24] Nidhal Mghirbi and Seth Douglas. An explicit nonlocal game of quantum value one that needs exponential entanglement, 2026. Version 1.0.0, unrefereed preprint. doi:10.5281/zenodo.23697255.

[25] Jonathan Oppenheim, Michał Horodecki, Paweł Horodecki, and Ryszard Horodecki. Thermodynamical approach to quantifying quantum correlations. *Physical Review Letters*, 89:180402, 2002. doi:10.1103/PhysRevLett.89.180402.

[26] Vern I. Paulsen, Simone Severini, Daniel Stahlke, Ivan G. Todorov, and Andreas Winter. Estimating quantum chromatic numbers. *Journal of Functional Analysis*, 270:2188–2222, 2016. doi:10.1016/j.jfa.2016.01.010.

[27] David Reeb and Michael M. Wolf. An improved Landauer principle with finite-size corrections. *New Journal of Physics*, 16:103011, 2014. doi:10.1088/1367-2630/16/10/103011.

[28] Benjamin Schumacher. Quantum coding. *Physical Review A*, 51:2738–2747, 1995. doi:10.1103/PhysRevA.51.2738.

[29] Ian Todd. Informational hardness: unifying Bell non-locality and Landauer cost, 2025. Preprint. doi:10.5281/zenodo.15251934.

[30] Thomas Vidick. Almost synchronous quantum correlations. arXiv:2103.02468, 2021. URL: https://arxiv.org/abs/2103.02468.

[31] Xingyao Wu, Jean-Daniel Bancal, Matthew McKague, and Valerio Scarani. Device-independent parallel self-testing of two singlets. *Physical Review A*, 93:062121, 2016. URL: https://arxiv.org/abs/1512.02074v2, doi:10.1103/PhysRevA.93.062121.