

Supplemental Material for “Two-Copy Onset of Strict Superadditivity for Measured Relative Entropy”

Maestro v0.1¹

¹*Pudim AI Research Ledger*

(Dated: Local candidate, August 22, 2026; Program Manager: Domingos S. P. Salazar)

I. THERMODYNAMIC SCOPE OF THE RELATIVE-ENTROPY IDENTITY

The structural theorem in the Letter concerns any faithful finite-dimensional pair (ρ, σ) . Its thermodynamic corollary uses the narrower isolated-driving construction. Let $\pi_0 = e^{-\beta H_0}/Z_0$, let $\rho_\tau = U\pi_0 U^\dagger$, and define $\pi_\tau = e^{-\beta H_\tau}/Z_\tau$. With $W = \text{Tr}(H_\tau \rho_\tau) - \text{Tr}(H_0 \pi_0)$ and $\Delta F = -\beta^{-1} \ln(Z_\tau/Z_0)$, unitary entropy conservation gives

$$\begin{aligned} D(\rho_\tau \| \pi_\tau) &= -S(\rho_\tau) + \beta \text{Tr}(H_\tau \rho_\tau) + \ln Z_\tau \\ &= \beta(W - \Delta F) \equiv \Sigma_{\text{iso}}. \end{aligned} \quad (\text{S1})$$

Thus $D(\rho_\tau \| \pi_\tau)$ is the dimensionless irreversible work, or nonequilibrium lag, for this specific Gibbs-prepared isolated protocol [S1, S2]. General entropy-production definitions require their own process and bookkeeping assumptions.

For comparison, the system–environment identity emphasized in Ref. [S2] assumes $\rho_{SE}(0) = \rho_S(0) \otimes \rho_E^{\text{th}}$, with $\rho_E^{\text{th}} = e^{-\beta H_E}/Z_E$, followed by a global unitary U_{SE} . In that framework,

$$\begin{aligned} \Sigma_{SE} &\equiv D(\rho_{SE}(\tau) \| \rho_S(\tau) \otimes \rho_E^{\text{th}}) \\ &= I(S:E)_\tau + D(\rho_E(\tau) \| \rho_E^{\text{th}}) = \Delta S_S + \beta Q_E. \end{aligned} \quad (\text{S2})$$

Here $Q_E = \text{Tr}[H_E(\rho_E(\tau) - \rho_E^{\text{th}})]$, so $Q_E > 0$ means energy gained by the environment. Equations (S1) and (S2) use distinct bookkeeping choices. A trivial environment makes $\Sigma_{SE} = 0$ under a system unitary, whereas Σ_{iso} may remain positive because it compares the driven state with the final Gibbs reference. All uses of entropy production in the Letter refer exclusively to Eq. (S1). The general two-copy theorem applies independently of Hamiltonians and thermodynamic interpretation.

II. PAIR-ROTATION DERIVATIVE

Let $R, S > 0$ be faithful states and let $(|x\rangle)$ be a rank-one projective measurement. For two selected basis vectors $|J\rangle, |K\rangle$, define

$$|J(\theta)\rangle = \cos \theta |J\rangle + e^{i\varphi} \sin \theta |K\rangle, \quad (\text{S3})$$

$$|K(\theta)\rangle = -e^{-i\varphi} \sin \theta |J\rangle + \cos \theta |K\rangle. \quad (\text{S4})$$

Here θ is the mixing angle and φ is the relative phase. Write $p_X = \langle X|R|X\rangle$, $q_X = \langle X|S|X\rangle$, and $t_X = p_X/q_X$. Let $D_{JK}(\theta, \varphi)$ denote the contribution of the two active outcomes to the classical output divergence. At $\theta = 0$ and fixed φ ,

$$\partial_\theta p_J|_0 = 2 \text{Re}(e^{i\varphi} R_{JK}), \quad \partial_\theta p_K|_0 = -\partial_\theta p_J|_0, \quad (\text{S5})$$

$$\partial_\theta q_J|_0 = 2 \text{Re}(e^{i\varphi} S_{JK}), \quad \partial_\theta q_K|_0 = -\partial_\theta q_J|_0. \quad (\text{S6})$$

For $g(p, q) = p \ln(p/q)$, $\partial_p g = 1 + \ln(p/q)$ and $\partial_q g = -p/q$. The directional derivative of the two active output terms is therefore

$$s_{JK}(\varphi) \equiv \left. \frac{\partial D_{JK}(\theta, \varphi)}{\partial \theta} \right|_{\theta=0} = 2 \text{Re} \left\{ e^{i\varphi} \left[\ln \left(\frac{t_J}{t_K} \right) R_{JK} + (t_K - t_J) S_{JK} \right] \right\}. \quad (\text{S7})$$

For the tensor power of a one-copy optimal measurement, choose $J = j^n$, $K = k^n$, $a = t_j$, and $b = t_k$. The stationarity relation in the Letter gives

$$\begin{aligned} R_{JK} &= L(a, b)^n \sigma_{jk}^n, & S_{JK} &= \sigma_{jk}^n, \\ t_J &= a^n, & t_K &= b^n. \end{aligned} \quad (\text{S8})$$

Using $a^n - b^n = \ln(a^n/b^n)L(a^n, b^n)$ and optimizing φ in Eq. (S7) yields

$$\max_{\varphi} s_{JK}(\varphi) = 2n \left| \ln \frac{a}{b} \right| [L(a^n, b^n) - L(a, b)^n] |\sigma_{jk}|^n. \quad (\text{S9})$$

Strict Jensen convexity gives

$$\begin{aligned} L(a^n, b^n) &= \int_0^1 (a^s b^{1-s})^n ds \\ &> \left(\int_0^1 a^s b^{1-s} ds \right)^n = L(a, b)^n \end{aligned} \quad (\text{S10})$$

for $a \neq b$ and $n \geq 2$. At $n = 2$, $L(a^2, b^2) = (a + b)L(a, b)/2$, so

$$\max_{\varphi} s_{JK}(\varphi) = 4|a - b| \left[\frac{a + b}{2} - L(a, b) \right] |\sigma_{jk}|^2 > 0. \quad (\text{S11})$$

III. FIXED-HORIZON SPECIALIZATION OF THE ADAPTIVE FRESH-COPY BOUND

For completeness, we specialize the conditional-divergence argument of Li, Tan, and Tomamichel [S3, Lemma 16(i)] to a deterministic horizon. Consider the hypotheses $\rho^{\otimes n}$ and $\sigma^{\otimes n}$. At round r , an arbitrary POVM $M_r(y^{r-1})$ acts on one fresh copy and may depend on all preceding classical outcomes. Only classical information passes between rounds. Let $P_{\rho}^{Y^n}$ and $P_{\sigma}^{Y^n}$ be the transcript laws. The classical chain rule gives

$$D(P_{\rho}^{Y^n} \| P_{\sigma}^{Y^n}) = \sum_{r=1}^n \mathbb{E}_{P_{\rho}} D(P_{\rho}^{Y_r | Y^{r-1}} \| P_{\sigma}^{Y_r | Y^{r-1}}). \quad (\text{S12})$$

Conditioned on a history, each term is the output divergence of one POVM applied to ρ and σ and is at most $D_M(\rho \| \sigma)$. Hence

$$D_{\text{seq}}^{(n)} \leq n D_M(\rho \| \sigma). \quad (\text{S13})$$

Repeating a one-copy optimum attains equality. This specialization covers sequential measurements on fresh copies with classical feed-forward. Separable and LOCC POVMs constitute broader restricted-measurement classes.

IV. QUBIT QUENCH AND EXACT FINITE-ANGLE PROBABILITIES

In the x -readout basis, take

$$\begin{aligned} \pi_f &= \frac{1}{2} \begin{pmatrix} 1 & c \\ c & 1 \end{pmatrix}, \quad \rho_{\tau} = \frac{1}{2} \begin{pmatrix} 1 + \delta & c L_{\delta} \\ c L_{\delta} & 1 - \delta \end{pmatrix}, \\ L_{\delta} &= \frac{2\delta}{\ln[(1 + \delta)/(1 - \delta)]}. \end{aligned} \quad (\text{S14})$$

The conditions $0 < \delta < 1$, $|c| < 1$, and $c^2 L_{\delta}^2 < 1 - \delta^2$ make both states faithful. The parametrization $c = \kappa \sqrt{1 - \delta^2} / L_{\delta}$, with $0 < \kappa < 1$, gives $c^2 L_{\delta}^2 = \kappa^2 (1 - \delta^2) < 1 - \delta^2$. To verify the remaining condition analytically, write $\delta = \tanh x$ with $x > 0$. Then $L_{\delta} = \tanh x / x > \text{sech } x = \sqrt{1 - \delta^2}$ because $\sinh x > x$. Hence $|c| < \kappa < 1$, so both states are faithful throughout the stated parameter domain.

The x -basis probabilities are $p_{\pm} = (1 \pm \delta)/2$ and $q_{\pm} = 1/2$, with likelihood ratios $t_{\pm} = 1 \pm \delta$. The only off-diagonal stationarity equation is

$$(\rho_{\tau})_{+-} = L(1 + \delta, 1 - \delta)(\pi_f)_{+-}, \quad (\text{S15})$$

which holds because $L(1 + \delta, 1 - \delta) = L_{\delta}$. Thus the positive operator $\Omega = \text{diag}(1 + \delta, 1 - \delta)$ is stationary in the concave variational problem for measured relative entropy [S4]. It is a global optimizer, so the x basis is one-copy optimal and

$$\Sigma_M^{(1)} = \frac{1 + \delta}{2} \ln(1 + \delta) + \frac{1 - \delta}{2} \ln(1 - \delta). \quad (\text{S16})$$

On two copies, use the ascent-phase rotation, corresponding to $\varphi_\star = \pi$ in Eq. (S4),

$$|J(\theta)\rangle = \cos\theta|+x, +x\rangle - \sin\theta|-x, -x\rangle, \quad (\text{S17})$$

$$|K(\theta)\rangle = \sin\theta|+x, +x\rangle + \cos\theta|-x, -x\rangle. \quad (\text{S18})$$

Let $u = \cos\theta$, $w = \sin\theta$, $r = (cL_\delta/2)^2$, and $s = (c/2)^2$. To isolate one controlled imperfection, introduce an active-subspace interference contrast $0 \leq v \leq 1$ multiplying only the selected off-diagonal contributions in both outcome laws. The four outcome probabilities are

$$p_J = u^2 p_+^2 + w^2 p_-^2 - 2vuwr, \quad q_J = \frac{1}{4} - 2vuws, \quad (\text{S19})$$

$$p_K = w^2 p_+^2 + u^2 p_-^2 + 2vuwr, \quad q_K = \frac{1}{4} + 2vuws, \quad (\text{S20})$$

$$p_{+-} = p_{-+} = p_+ p_-, \quad q_{+-} = q_{-+} = \frac{1}{4}. \quad (\text{S21})$$

The exact finite-angle output divergence is $D_{\text{out}}^{(2)}(\theta, \varphi_\star) = \sum_X p_X \ln(p_X/q_X)$. Differentiation at $v = 1$ gives

$$\left. \frac{\partial D_{\text{out}}^{(2)}(\theta, \varphi_\star)}{\partial \theta} \right|_{\theta=0} = 2\delta(1 - L_\delta)c^2 > 0. \quad (\text{S22})$$

The active-subspace contrast model in Eq. (S21) is not intended to represent arbitrary state-preparation, gate, or detector noise.

Every faithful ρ_τ is a Gibbs state of $H_i = -\beta^{-1} \ln \rho_\tau + \text{const.}$, while π_f fixes $H_f = -\beta^{-1} \ln \pi_f + \text{const.}$. Identifying $\pi_0 = \rho_\tau$ and $\pi_\tau = \pi_f$, a sudden isolated quench $H_i \rightarrow H_f$ realizes Eq. (S1); the state is unchanged during the instantaneous quench, while its equilibrium reference changes. The subsequent collective rotation and x -population detection interrogate the resulting state pair. Estimating the relative entropy additionally requires calibrated outcome probabilities for the reference preparation $\pi_f^{\otimes 2}$.

V. CALIBRATED SAMPLING ESTIMATE

For known reference probabilities q and multinomial frequencies $\hat{p}$ from N independent rounds, the plug-in estimator $D(\hat{p}||q)$ has the delta-method variance [S5]

$$\frac{V(p||q)}{N}, \quad V(p||q) = \sum_i p_i \left[\ln \frac{p_i}{q_i} - D(p||q) \right]^2. \quad (\text{S23})$$

For independent one-copy and two-copy measurements with the same number N of rounds, let $\hat{\Delta}_{\mathcal{M}} = \hat{D}_2/2 - \hat{D}_1$ denote the per-copy gain for a fixed two-copy readout $\mathcal{M}$. Its asymptotic variance is

$$\text{Var}(\hat{\Delta}_{\mathcal{M}}) \simeq \frac{V(p_2||q_2)/4 + V(p_1||q_1)}{N}. \quad (\text{S24})$$

The nominal rounds required for a z -standard-deviation separation are

$$N_z = \frac{z^2 [V(p_2||q_2)/4 + V(p_1||q_1)]}{\Delta_{\mathcal{M}}^2}. \quad (\text{S25})$$

For the recovered stationary-candidate PVM and the explicit pair rotation, Eq. (S25) is evaluated separately; the numerator and gain are never mixed across the two readouts. It gives $N_5 = 38349.7$ rounds per arm in the continuous asymptotic calculation for the recovered candidate at the best displayed grid point. An executable integer design rounds this upward to 38350 rounds per arm. This delta-method estimate omits finite-sample bias, uncertainty in q , multiple testing, drift, correlations, and systematic errors. It is a plausibility metric rather than a detection guarantee. Equal rounds N in the one- and two-copy arms consume $N + 2N = 3N$ signal-state copies. Copies used to prepare and estimate the reference laws q_1 and q_2 are additional and are excluded from this count.

VI. ARCHIVED BACKEND-SNAPSHOT PROVENANCE

The Research Object preserves a historical classical backend-snapshot noise-model study under `external/ibm_fake_v27/`. The upstream provider snapshot can be stale and need not represent current hardware [S6]. Its source pair and readout differ from the analytic qubit family in the Letter, so its outputs provide neither theorem evidence nor a hardware-performance claim. The package retains its protocol, circuit files, 10,000-repetition Monte Carlo output, environment lock, and derived figures solely for provenance and reconstruction. No IBM service connection or QPU measurement is used by that archived calculation.

VII. RECOVERED TWO-COPY NUMERICAL STATIONARY CANDIDATES

For $R = \rho_\tau^{\otimes n}$ and $S = \pi_f^{\otimes n}$, $n \in \{1, 2\}$, we evaluate the Berta–Fawzi–Tomamichel variational program directly. Positivity is enforced by the bijective log coordinate $\Omega = e^H$, giving

$$\mathcal{F}(H) = \text{Tr}(RH) + 1 - \text{Tr}(Se^H), \quad \nabla_H \mathcal{F} = R - D \exp_H(S). \quad (\text{S26})$$

The states are real, and averaging any Ω with its complex conjugate cannot lower the real, concave objective, so a real-symmetric optimizer exists. BFGS locates a stationary basin and a nonlinear least-squares root solve polishes $R - D \exp_H(S) = 0$. First-order consistency with a global maximizer is then tested in the original concave variable by

$$\|D \log_\Omega(R) - S\|_F < 2 \times 10^{-8}. \quad (\text{S27})$$

For an exact stationary point, concavity would imply global optimality; the finite residual in Eq. (S27) is a numerical stationarity check, not a rigorous objective-gap enclosure. If $\Omega = \sum_i \lambda_i |u_i\rangle\langle u_i|$, the recovered stationary-candidate PVM has $p_i = \langle u_i | R | u_i \rangle$, $q_i = \langle u_i | S | u_i \rangle$. We independently verify $\lambda_i = p_i/q_i$ and $\mathcal{F}(\Omega) = \sum_i p_i \ln(p_i/q_i)$ within the stated numerical tolerances.

The production grid contains 151 values of $\delta \in [0.20, 0.95]$ at each $\kappa \in \{0.70, 0.90, 0.99\}$. Pointwise, the optimized one-copy value agrees with the analytic expression above within 1.7×10^{-16} , and

$$\Sigma_M^{(1)} \leq \widehat{\Sigma}_{M,\text{num}}^{(2)} \leq D(\rho_\tau \| \pi_f), \quad \widehat{\Delta}_{\text{rot}} \leq \widehat{\Delta}_{\text{num}}. \quad (\text{S28})$$

The largest original-variable stationarity residual is 3.5×10^{-9} , the largest $|\lambda_i - p_i/q_i|$ residual is 5.9×10^{-9} , and the spectral-PVM value agrees with the variational value within 1.6×10^{-15} . A symmetry-reduced CVXPY 1.8.2/Clarabel audit at five corner and representative points uses a semidefinite quadrature approximation and agrees with the production values within 3.6×10^{-10} nats per copy.

On the displayed finite grid, the largest numerically recovered stationary-candidate gain is

$$\widehat{\Delta}_{\text{num}} = 0.0214758723 \text{ nats/copy} \quad \text{at} \quad (\delta, \kappa) = (0.810, 0.99). \quad (\text{S29})$$

At this point $\Sigma_M^{(1)} = 0.3791913303$, $\widehat{\Sigma}_{M,\text{num}}^{(2)} = 0.4006672026$, and $D = 0.5584510295$. The pair rotation attains 0.0073100901 nats per copy there. Across the full grid it captures 24.3%–49.4% of the recovered gap. Equation (S25) gives $N_5 = 38349.7$ rounds per arm in the continuous asymptotic calculation for the recovered stationary-candidate PVM at the best grid point; rounding upward gives 38350 rounds per arm and 115050 signal copies.

VIII. NUMERICAL METHODS AND REPRODUCIBILITY

The script `code/optimize_two_copy_sdp.py` generates `data/two_copy_sdp.csv` and `data/two_copy_sdp_validation.json`. Its analytic and primary-figure environment is Python 3.11.9 with NumPy 2.2.6, SciPy 1.17.0, Matplotlib 3.10.8, and CVXPY 1.8.2. The separately routed archived backend-snapshot environment is Python 3.11.9 with NumPy 2.4.6, SciPy 1.17.1, Matplotlib 3.11.1, Qiskit 2.5.2, Qiskit IBM Runtime 0.47.0, and Qiskit Aer 0.17.2. The project bundle preserves one disjoint lock for each route and records the actual runtime versions in `metadata/environment-provenance.yaml`. The filename `optimize_two_copy_sdp.py` records the initial semidefinite-program request; its production values use the exact log-coordinate objective in Eq. (S26), because CVXPY’s `quantum_rel ENTR` is itself a quadrature approximation.

The companion script `code/simulate_qubit_activation.py` optimizes the explicit rotation over $0 \leq \theta \leq \pi/2$ using a deterministic 4097-point grid, bounded refinement of every sampled local maximum, and endpoint comparison.

A seeded audit checks 80 faithful pairs; the largest slope residual is 4.21×10^{-9} . A separate matrix-level implementation constructs the tensor-product states and rotated projectors without calling the analytic probability routine. Across 50 additional seeded cases, its largest probability residual is 1.67×10^{-16} and its largest slope residual is 2.65×10^{-9} . The files are `data/validation.json`, `code/independent_qubit_audit.py`, and `data/independent_validation.json`. The theorem proof is analytic and independent of these numerical checks.

-
- [S1] S. Deffner and E. Lutz, *Physical Review Letters* **105**, 170402 (2010).
 - [S2] G. T. Landi and M. Paternostro, *Reviews of Modern Physics* **93**, 035008 (2021).
 - [S3] Y. Li, V. Y. F. Tan, and M. Tomamichel, *Communications in Mathematical Physics* **392**, 993 (2022), arXiv:2104.14706 [quant-ph].
 - [S4] M. Berta, O. Fawzi, and M. Tomamichel, *Letters in Mathematical Physics* **107**, 2239 (2017).
 - [S5] A. W. van der Vaart, *Asymptotic Statistics* (Cambridge University Press, 1998).
 - [S6] IBM Quantum, Fake provider: Qiskit ibm runtime documentation, <https://quantum.cloud.ibm.com/docs/en/api/qiskit-ibm-runtime/fake-provider> (2026), accessed August 19, 2026.