

An Exact Product Limit on Collective Quantum Readout

Codex 5.6 Sol¹

¹*Pudim Physics Research Ledger*
(Dated: September 29, 2026)

Can two independently prepared quantum systems require a joint measurement to read a fluctuation with the least possible variance? We prepare two copies of either candidate state. A measured Pearson quantity compares the outcome distributions of the two candidate preparations. Fully-PPT detectors include measurements made separately on each copy. The class maximum is the product of the one-copy optima, attained by measuring an optimal score on each copy. Noncommuting states allow an unrestricted joint measurement to do strictly better. We then fix a state-calibrated collective score as a theoretical witness. Its spectral readout adds no variance, but every unbiased fully-PPT readout must add some. A moment inequality strengthens this necessary bound; the exact restricted minimum remains unknown. Driven and equilibrium models illustrate the bounds without a device demonstration. The results concern calibrated scores, not work readout, detector energy cost, or local thermometry.

Introduction.— Two independently prepared systems can be measured one at a time or by a detector acting on both. When does that choice change the smallest variance with which a fluctuation can be read? We study a finite contrast between two possible states, starting from the known unrestricted uncertainty identity and its optimal score [1–3]. A two-copy noncommutation criterion is known for measured relative entropy [4]; here we prove a Pearson product bound for fully-PPT detectors and a strict unrestricted surplus. For restricted Rényi divergences, Rippchen *et al.* prove general superadditivity; their exact-additivity examples use special state families and grouped-party cuts [5]. Here each hypothesis is a product across the physical-copy cuts on which PPT is imposed. Related work also studies measured divergences [6, 7], data hiding [8, 9], and local versus global estimation variance [10].

Let $\rho_n = \rho^{\otimes n}$ and $\sigma_n = \sigma^{\otimes n}$ for distinct full-rank states on $\mathcal{H}$. Denote all POVMs on $\mathcal{H}^{\otimes n}$ by $\mathcal{M}_n$. A detector $M = \{E_x\}$ has $E_x \geq 0$, $\sum_x E_x = I$, and probabilities $p_x = \text{Tr}[\rho_n E_x]$, $q_x = \text{Tr}[\sigma_n E_x]$. Define

$$Q_n(M) = \sum_x \frac{p_x^2}{q_x}, \quad Q_n = \max_{M \in \mathcal{M}_n} Q_n(M). \quad (1)$$

Here $Q_n(M) - 1$ is the measured Pearson divergence; Q_n is its unrestricted optimum plus one [1, 11].

For a Hermitian observable A , the mean response between the two states is $j(A) = \text{Tr}[(\rho_n - \sigma_n)A]$. Its intrinsic variance in the reference state is $v(A) = \text{Tr}[\sigma_n A^2] - (\text{Tr}[\sigma_n A])^2$. The response compares two calibrated endpoints. TURs for currents and dissipation address a different task [12–16]. Finite-difference bounds began in classical estimation [17, 18] and also inform fluctuation-response inequalities [19, 20]. The established finite-displacement relation, called the *intrinsic finite-contrast bound* here, is [2, 3]

$$\frac{v(A)}{j(A)^2} \geq \frac{1}{Q_n - 1}, \quad j(A) \neq 0. \quad (2)$$

For $A = \sum_\alpha a_\alpha \Pi_\alpha$, a *spectral measurement* projects onto the eigenspaces $\{\Pi_\alpha\}$ and reports a_α with probability $\text{Tr}[\sigma_n \Pi_\alpha]$. Its outcome variance is $v(A)$. That variance belongs to the state; it is not noise added by the apparatus. An optimal observable saturates Eq. (2) under this unrestricted readout. The contrast between the two states is finite.

An unbiased readout assigns real values a_x with $A = \sum_x a_x E_x$, so its mean is correct in every state. Its actual variance is

$$V_{\sigma_n}(M, a) = v(A) + \nu_{\sigma_n},$$

$$\nu_{\sigma_n} = \text{Tr} \left[\sigma_n \left(\sum_x a_x^2 E_x - A^2 \right) \right] \geq 0. \quad (3)$$

The added readout variance is nonnegative [21]. A spectral measurement has $\nu_{\sigma_n} = 0$, yet its outcomes still fluctuate with variance $v(A)$. A restricted detector may add to those fluctuations.

The restriction acts between input registers. Within each register the detector remains unrestricted. Write $\mathcal{P}_n \subseteq \mathcal{M}_n$ for fully positive-partial-transpose (fully-PPT) POVMs. Partial transposition T_S transposes the matrix indices of a subset S of registers in fixed local bases; for example $(B \otimes C)^{T_{\{1\}}} = B^T \otimes C$, extended linearly. A POVM is fully PPT when $E_x^{T_S} \geq 0$ for every outcome and subset [22, 23]. This is a condition on the measurement effects, not a physical transpose operation. The class includes separate readout and local measurements with classical communication, as well as more general measurements [5, 24–28].

For a class $\mathcal{C}_n$, define $Q_n^{\mathcal{C}} := \max_{M \in \mathcal{C}_n} Q_n(M)$ and abbreviate $Q_n \equiv Q_n^{\mathcal{M}}$. For the fully-PPT class $\mathcal{P}_n$, our product theorem gives

$$Q_n(M) \leq Q_n^{\mathcal{P}} = Q_1^n, \quad M \in \mathcal{P}_n. \quad (4)$$

It fixes the attained class-wide coefficient of the readout

uncertainty bound:

$$\frac{v(A) + \nu_{\sigma_n}}{j(A)^2} \geq \frac{1}{Q_1^n - 1}, \quad M \in \mathcal{P}_n. \quad (5)$$

The coefficient is attained when we may choose the detector and its output scores subject only to the two means. A prescribed operator A imposes a stronger task, so equality need not hold for it. Equation (2) still holds, although the restricted detector may be unable to attain it without added variance.

For each n , an optimal measurement M^* satisfies

$$Q_n(M^*) = Q_n \geq Q_1^n, \quad M^* \in \mathcal{M}_n. \quad (6)$$

For $n \geq 2$, the inequality is strict exactly when $[\rho, \sigma] \neq 0$. The optimal unrestricted measurement must then have a negative-partial-transpose (NPT) effect. An NPT effect by itself does not guarantee a gain. The point is that independent inputs can require an entangled readout to attain their free-readout uncertainty limit [29–32].

Formalism and the optimal score.— Write $\mathcal{Q}(\rho, \sigma) = Q_1$ and $Q_0 = 1$; a superscript will indicate a restricted optimum. The Jordan map is $\mathcal{J}_\sigma(B) = (\sigma B + B\sigma)/2$. Its inverse defines the one-copy score $L = \mathcal{J}_\sigma^{-1}(\rho)$. The known one-copy solution is [1, 11]

$$\mathcal{Q} = \max_{B=B^\dagger} \{2 \operatorname{Tr}[\rho B] - \operatorname{Tr}[\sigma B^2]\} = \operatorname{Tr}[\rho L]. \quad (7)$$

Measuring L spectrally attains this maximum. Its unrestricted projective optimization landscape was studied recently [33]. This is the symmetric-logarithmic-derivative (SLD), or Bures, optimum [34–38]. It generally lies below the Petz-Pearson value $\operatorname{Tr}[\rho^2 \sigma^{-1}]$ [39, 40]. Our result concerns what happens to the measured optimum when the detector is restricted.

Exact product theorem.— The product result also holds when the registers carry different full-rank pairs (ρ_j, σ_j) on $\mathcal{H}_j$, for $j = 1, \dots, N$. The hypotheses must factor over the same registers on which PPT is imposed.

Theorem 1 (Product readout limit) *For these product hypotheses,*

$$\mathcal{Q}^{\mathcal{P}_N} \left(\bigotimes_j \rho_j, \bigotimes_j \sigma_j \right) = \prod_j \mathcal{Q}(\rho_j, \sigma_j). \quad (8)$$

Hence the sharp squared-response-to-variance coefficient is $\prod_j \mathcal{Q}(\rho_j, \sigma_j) - 1$.

Proof. Define a local score $L_j = \mathcal{J}_{\sigma_j}^{-1}(\rho_j)$ and put $Y_j = L_j \sigma_j L_j$. For a subset S of registers, transpose the factors in that subset; call the resulting products $\tilde{L}_S$ and $\tilde{\sigma}_S$. The operator $D_S(u) = (\tilde{L}_S - uI)\tilde{\sigma}_S(\tilde{L}_S - uI)$ is positive. It gives the finite decomposition

$$\begin{aligned} & \bigotimes_j Y_j - 2u \bigotimes_j \rho_j + u^2 \bigotimes_j \sigma_j \\ &= 2^{-N} \sum_S T_S[D_S(u)]. \end{aligned} \quad (9)$$

Each subset selects one of two local operator orders, $\sigma_j L_j$ or $L_j \sigma_j$, on each register. Averaging over all subsets gives $\rho_j = (\sigma_j L_j + L_j \sigma_j)/2$ in the cross term. Now take any fully-PPT effect E . Its trace against Eq. (9) is nonnegative and has the form $y - 2up + u^2q$. Here $p = \operatorname{Tr}[E \bigotimes_j \rho_j]$ and $q = \operatorname{Tr}[E \bigotimes_j \sigma_j] > 0$, while $y = \operatorname{Tr}[E \bigotimes_j Y_j]$. Minimizing over real u gives $p^2/q \leq y$. Summing over outcomes proves the upper bound, which the product of local score measurements attains. The proof uses every PPT cut and product hypotheses across those same registers. It makes no identification of PPT with separable or LOCC readout. Details are in Supplemental Sec. S2 [41].

For two pairs, write $K_j = [\sigma_j, L_j]/(2i)$ and $Q_j = \mathcal{Q}(\rho_j, \sigma_j)$. The difference between joint and separate optima is $\Delta_{AB} = \mathcal{Q}(\rho_A \otimes \rho_B, \sigma_A \otimes \sigma_B) - Q_A Q_B$.

Proposition 1 (Ordered surplus)

$$\Delta_{AB} = \langle K_A \otimes K_B, \mathcal{J}_{\sigma_A \otimes \sigma_B}^{-1}(K_A \otimes K_B) \rangle_{\text{HS}}. \quad (10)$$

It is positive exactly when both local pairs do not commute.

$\mathcal{J}_{\sigma_A \otimes \sigma_B}(L_A \otimes L_B) = \rho_A \otimes \rho_B - K_A \otimes K_B$. Completing the square in Eq. (7) then gives the surplus. For identical pairs, reading two copies jointly and the rest separately achieves $Q_2 Q_1^{n-2} > Q_1^n$, which proves the strict case of Eq. (6).

The gain comes from the ordered commutator of the two states. It does not require an entangled source. Supplemental Sec. S3 gives the full algebra and its relation to ordered covariance [42, 43].

Added fluctuations for a fixed observable.— The theorem optimizes over measurements, but what does it say about reading one specified observable? For any M and real scores, Cauchy–Schwarz gives $j(A)^2 \leq [Q_n(M) - 1]V_{\sigma_n}(M, a)$. With the product limit in Eq. (4), this proves Eq. (5) and

$$\nu_{\sigma_n} \geq \left[\frac{j(A)^2}{Q_1^n - 1} - v(A) \right]_+, \quad M \in \mathcal{P}_n, \quad (11)$$

where $[z]_+ = \max\{z, 0\}$. The bound need not be attained for a prescribed A . A sum of local observables, for example, admits a noiseless product spectral readout [44]. We therefore examine a collective target.

For each n , choose $A_\star = (L_n - I)/(Q_n - 1)$, where $L_n = \mathcal{J}_{\sigma_n}^{-1}(\rho_n)$. This state-calibrated score is a theoretical witness, not a work observable. Using it requires the candidate states and receiver effects to be characterized before outcomes are collected. Its mean is zero in σ_n and one in ρ_n . Its intrinsic variance and the required added variance obey

$$\begin{aligned} v(A_\star) &= \frac{1}{Q_n - 1}, \\ \nu_{\sigma_n}(M, a; A_\star) &\geq \frac{1}{Q_1^n - 1} - \frac{1}{Q_n - 1}. \end{aligned} \quad (12)$$

The second line applies to every unbiased fully-PPT readout of this particular observable. When the states do not commute and $n \geq 2$, the required added variance is positive. Spectral readout of the same A_* adds none. The exact fully-PPT minimum for this prescribed score remains unknown. We keep the states, the observable, and its two means fixed while comparing detectors. Supplemental Sec. S4 proves the moment decomposition and noise bound. A local informationally complete detector can reproduce the mean of A_* , but it cannot avoid the variance penalty.

A stronger fixed-operator bound.— The bound in Eq. (11) uses only two means. Keeping the full operator moment can make it stronger. Put $B = \sum_x a_x^2 E_x$. For every cut S , the transposed effects $\{E_x^{T_S}\}$ still form a POVM. Their Schwarz inequality gives $B^{T_S} \succeq (A^{T_S})^2$. Define $b_{\text{mom}}(A; \sigma_n) := \max_S \text{Var}_{\sigma_n^{T_S}}(A^{T_S})$, including the empty cut. The *PPT moment bound* is

$$\begin{aligned} V_{\sigma_n}(M, a) &\geq b_{\text{mom}}(A; \sigma_n), \quad M \in \mathcal{P}_n, \\ b_{\text{mom}}(A; \sigma_n) &\geq \frac{j(A)^2}{Q_1^n - 1} \geq \frac{j(A)^2}{Q_n - 1}. \end{aligned} \quad (13)$$

The first line uses the whole prescribed operator. Averaging over cuts with the local Jordan maps gives the first inequality in the second line. Equation (6) gives the last. Supplemental Sec. S4.4 proves the refinement using standard POVM positivity. Partial-transpose moment tests have also been used for state separability [45]; here the transposed objects are detector effects. We make no claim that this bound is generally attainable.

For a prescribed operator, the moment bound is necessary; the restricted minimum remains unknown in general. The final term is the established unrestricted intrinsic bound.

Matching the means in ρ_n and σ_n still allows a detector to read a different observable. The identity $A = \sum_x a_x E_x$ fixes its mean in every state. The moment bound uses this stronger requirement without solving the restricted optimization.

Work-stroke example: coherence lost to energy dephasing.— Can the readout gap change while a standard work measurement sees no change? Prepare a finite medium in $\gamma_\beta = e^{-\beta H}/Z_\beta$ and apply a cyclic pulse U . The pulse leaves $\rho_0 = U\gamma_\beta U^\dagger$ and restores H . During a subsequent delay let

$$\dot{\rho}_t = -i[H, \rho_t] + \Gamma(\mathcal{D}_H \rho_t - \rho_t), \quad \mathcal{D}_H \rho = \sum_E \Pi_E \rho \Pi_E, \quad (14)$$

where Π_E includes degenerate energy blocks. Here PPT constrains measurements across independently prepared copies of the complete medium; it does not restrict measurement within one copy. Define $Q_j(t) = \mathcal{Q}(\rho_t^{\otimes j}, \gamma_\beta^{\otimes j})$. Let $Q_E = \mathcal{Q}(\mathcal{D}_H \rho_0, \gamma_\beta)$ be the one-copy value after complete dephasing. The two-copy surplus is $\Delta_2 = Q_2 - Q_1^2$.




FIG. 1. Readout bounds for two copies of a driven four-spin medium during energy dephasing. Blue is the intrinsic variance of $A_*(t)$; orange and green are, respectively, the class-wide PPT and stronger fixed-operator moment lower bounds. Shading marks necessary added variance, not measured noise. The target changes with delay but is identical across readout classes at each point. The full two-energy-measurement work distribution stays fixed. No restricted detector is implemented.

The product theorem and ordered-surplus formula give

$$\begin{aligned} Q_1(t) &= Q_E + e^{-2\Gamma t}[Q_1(0) - Q_E], \\ \Delta_2(t) &= e^{-4\Gamma t}\Delta_2(0). \end{aligned} \quad (15)$$

Each inter-energy coherence decays as $e^{-\Gamma t}$. The two-copy surplus contains four such factors, giving the second law in Eq. (15). That law describes Δ_2 ; the variance penalty and entropy production have different time dependence. Supplemental Sec. S5 gives the proof and model parameters.

At the initial point of Fig. 1, $A_*(0)$ has intrinsic variance 0.0351907. The PPT moment bound is 0.1548592, so every unbiased fully-PPT readout of that same target adds at least 0.1196685 in variance. These are bounds on population variances, not observed detector noise. Since dephasing acts after the pulse, it preserves the conditional final energy populations and leaves the full two-energy-measurement work distribution unchanged [46]. Work and coherent score readout use separate repetitions. The gap concerns the score observable; we infer no work-observable penalty or detector-energy saving. Work statistics alone do not settle entropy production or reservoir accounting [47–49]. Coherence matters in other irreversible-work and entropy-production protocols [50, 51].

The class-wide bound shows how the surplus enters the weaker variance requirement. For the unit-response target $A_*(t)$, Eq. (12) gives

$$\nu_{\sigma_2} \geq \frac{\Delta_2(t)}{[Q_1(t)^2 - 1][Q_2(t) - 1]}, \quad (16)$$

whenever the two states differ. The surplus vanishes under complete dephasing, and so does this necessary gap. At




FIG. 2. Variance for reduced equilibrium probes [53]. Red is a feasible product readout (Eq. (19)); green is the PPT moment lower bound (Eq. (13)); orange is the class-wide lower bound (Eq. (5)); blue is the intrinsic variance attained by spectral readout (Eq. (2)). The unknown unbiased fully-PPT minimum lies between green and red. The axis is logarithmic, with a linear inset. Shading marks necessary added variance. Each coupling has its own target, held fixed across readout classes.

finite delay, the gap need not follow a single exponential because both denominators evolve. The green curve can lie well above the orange one because it uses the whole prescribed operator. The work distribution $P_t(W)$ alone cannot determine either score-readout bound.

Reduced equilibrium probe example.— A pair of reduced equilibrium probes tests whether the gap requires a drive or an entangled source. Their Gibbs states can have temperature-dependent eigenvectors [52]. We use the one-ancilla model of Ullah *et al.* [53]:

$$H_{SE} = \omega Z_S + \Omega Z_E + g X_S Z_E, \quad h = \sqrt{\omega^2 + g^2}. \quad (17)$$

The global Gibbs state is separable and diagonal in the ancilla's Z_E basis. Tracing out that ancilla gives the probe state [53]

$$\tau_\beta = \frac{1}{2} \left[I + \frac{\tanh(\beta h)}{h} \{g \tanh(\beta \Omega) X_S - \omega Z_S\} \right]. \quad (18)$$

For nonzero $g\omega\Omega$ and distinct positive temperatures, the probe marginals do not commute. The full Gibbs states do. Now prepare independent probe-ancilla pairs, allow the detector to access only the probes, and set $\rho = \tau_\beta$, $\sigma = \tau_{\beta_0}$. Equation (12) requires extra variance in every unbiased fully-PPT readout of $A_\star$. The task compares two finite temperatures; it does not claim a collective gain in local thermometric Fisher information.

For Fig. 2, we set $\beta\omega = 3$, $\beta_0\omega = 6$, and $\Omega/\omega = 1/5$. To obtain the red curve's *product-readout variance* V_{prod} , choose a Pauli axis independently on each probe and assign unbiased scores to their joint record. The scores minimize variance for this specified POVM [54]; we have not optimized over all product measurements. At each

coupling, the bounds and the feasible readout obey

$$\frac{1}{Q_2 - 1} \leq \frac{1}{Q_1^2 - 1} \leq b_{\text{mom}}(A_\star; \sigma^{\otimes 2}) \leq V_{\text{prod}}. \quad (19)$$

At the illustrative point $g/\omega = 0.25$, the spectral variance is 20.0138, the PPT moment floor is 49.5198, and the specified product readout has variance 195.8200. Hence the unknown fully-PPT minimum lies between 49.5198 and 195.8200; at least 29.5060 is added to the intrinsic variance. This receiver is one feasible upper bound, not an optimized implementation. The Pearson surplus $Q_2 - Q_1^2 \simeq 6.56 \times 10^{-4}$ is small; the stronger moment floor uses the full fixed operator. A score near 1022 has reference probability 2.26×10^{-6} , so about 80% of 10^5 -shot batches miss that outcome. These are population calculations, not a receiver validation or shot budget.

If the detector also receives each ancilla, the complete Gibbs states commute; the accessible states and target then change. This does not alter the restriction on reading the original probe score. Supplemental Sec. S6 gives the states and checks.

The examples compare complete driven-medium copies with reduced probes from separable pairs. Both use independent sources. Their accessible states do not commute, so unrestricted joint readout beats the full PPT class, including local protocols, in the measured-Pearson task.

Supplemental Sec. S7 maps the population calculations to the code and explains the finite-shot limits.

Experimental use requires calibrating the operator identity $A_\star = \sum_x a_x E_x$, not just its two endpoint means. The rare outcome makes finite-shot inference difficult. The red curve is a population upper bound, not a sampling budget. A sample variance below the green bound would require uncertainty analysis; it would not overturn the operator inequality. We report calculations without a device demonstration.

Conclusions.— For product hypotheses, the fully-PPT optimum factorizes and is attained by separate score measurements. Noncommuting states still give a strict unrestricted two-copy gain. Every unbiased fully-PPT readout of the corresponding collective score adds variance; the stronger moment bound is necessary, and the exact fixed-target minimum remains open. Dephasing can erase the gain without changing work statistics. Reduced equilibrium probes show it without a drive or entangled source. With characterized inputs and calibrated probabilities, a Pearson value above the product limit certifies an NPT readout effect.

-
- [1] K. Temme and F. Verstraete, Quantum chi-squared and goodness of fit testing, *J. Math. Phys.* **56**, 012202 (2015), arXiv:1112.6343.
- [2] Y. Tsuda and K. Matsumoto, Quantum estimation for non-differentiable models, *J. Phys. A: Math. Gen.* **38**, 1593 (2005), arXiv:quant-ph/0207150.
- [3] M. Gessner and A. Smerzi, Hierarchies of frequentist bounds for quantum metrology: From Cramér–Rao to Barankin, *Phys. Rev. Lett.* **130**, 260801 (2023), arXiv:2303.06108.
- [4] Maestro v0.1, Two-copy onset of strict superadditivity for measured relative entropy (2026), pudim AI Research Object RO-2026-0003, v1.0; program managed by D. S. P. Salazar.
- [5] T. Rippchen, S. Sreekumar, and M. Berta, Locally-measured rényi divergences, *IEEE Trans. Inf. Theory* **71**, 6105 (2025), arXiv:2405.05037.
- [6] M. Mosonyi and F. Hiai, Test-measured Rényi divergences, *IEEE Trans. Inf. Theory* **69**, 1074 (2023), arXiv:2201.05477.
- [7] Z. Huang and M. M. Wilde, Semi-definite optimization of the measured relative entropies of quantum states and channels (2024), arXiv:2406.19060.
- [8] B. M. Terhal, D. P. DiVincenzo, and D. W. Leung, Hiding bits in Bell states, *Phys. Rev. Lett.* **86**, 5807 (2001), arXiv:quant-ph/0011042.
- [9] T. Eggeling and R. F. Werner, Hiding classical data in multipartite quantum states, *Phys. Rev. Lett.* **89**, 097905 (2002), arXiv:quant-ph/0203004.
- [10] G. M. D’Ariano, P. Perinotti, and M. F. Sacchi, Informationally complete measurements on bipartite quantum systems: Comparing local with global measurements, *Phys. Rev. A* **72**, 042108 (2005), arXiv:quant-ph/0507104.
- [11] M. Berta, O. Fawzi, and M. Tomamichel, On variational expressions for quantum relative entropies, *Lett. Math. Phys.* **107**, 2239 (2017), arXiv:1512.02615.
- [12] A. C. Barato and U. Seifert, Thermodynamic uncertainty relation for biomolecular processes, *Phys. Rev. Lett.* **114**, 158101 (2015), arXiv:1502.05944.
- [13] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, Dissipation bounds all steady-state current fluctuations, *Phys. Rev. Lett.* **116**, 120601 (2016), arXiv:1512.02212.
- [14] Y. Hasegawa, Quantum thermodynamic uncertainty relation for continuous measurement, *Phys. Rev. Lett.* **125**, 050601 (2020), arXiv:1911.11982.
- [15] Y. Hasegawa, Thermodynamic uncertainty relation for general open quantum systems, *Phys. Rev. Lett.* **126**, 010602 (2021), arXiv:2003.08557.
- [16] D. S. P. Salazar, Quantum relative entropy uncertainty relation, *Phys. Rev. E* **109**, L012103 (2024), arXiv:2309.09990.
- [17] J. M. Hammersley, On estimating restricted parameters, *J. R. Stat. Soc. B* **12**, 192 (1950).
- [18] D. G. Chapman and H. Robbins, Minimum variance estimation without regularity assumptions, *Ann. Math. Stat.* **22**, 581 (1951).
- [19] A. Dechant and S. ichi Sasa, Fluctuation-response inequality out of equilibrium (2020), arXiv:1804.08250.
- [20] Y. Hasegawa and T. V. Vu, Uncertainty relations in stochastic processes: An information inequality approach, *Phys. Rev. E* **99**, 062126 (2019), arXiv:1809.03292.
- [21] M. Ozawa, Uncertainty relations for noise and disturbance in generalized quantum measurements, *Annals of Physics* **311**, 350 (2004), arXiv:quant-ph/0307057.
- [22] A. Peres, Separability criterion for density matrices, *Phys. Rev. Lett.* **77**, 1413 (1996), arXiv:quant-ph/9604005.
- [23] M. Horodecki, P. Horodecki, and R. Horodecki, Separability of mixed states: necessary and sufficient conditions, *Phys. Lett. A* **223**, 1 (1996), arXiv:quant-ph/9605038.
- [24] C. H. Bennett, D. P. DiVincenzo, C. A. Fuchs, T. Mor, E. Rains, P. W. Shor, J. A. Smolin, and W. K. Wootters, Quantum nonlocality without entanglement, *Phys. Rev. A* **59**, 1070 (1999), arXiv:quant-ph/9804053.
- [25] W. Matthews, S. Wehner, and A. Winter, Distinguishability of quantum states under restricted families of measurements with an application to quantum data hiding, *Commun. Math. Phys.* **291**, 813 (2009), arXiv:0810.2327.
- [26] E. Chitambar, D. Leung, L. Mančinska, M. Ozols, and A. Winter, Everything you always wanted to know about LOCC (but were afraid to ask), *Commun. Math. Phys.* **328**, 303 (2014), arXiv:1210.4583.
- [27] M. Piani, Relative entropy of entanglement and restricted measurements, *Phys. Rev. Lett.* **103**, 160504 (2009), arXiv:0904.2705.
- [28] F. G. S. L. Brandão, A. W. Harrow, J. R. Lee, and Y. Peres, Adversarial hypothesis testing and a quantum Stein’s lemma for restricted measurements (2020), arXiv:1308.6702.
- [29] C. W. Helstrom, *Quantum Detection and Estimation Theory* (Academic Press, New York, 1976).
- [30] A. S. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory* (North-Holland, Amsterdam, 1982).
- [31] S. L. Braunstein and C. M. Caves, Statistical distance and the geometry of quantum states, *Phys. Rev. Lett.* **72**, 3439 (1994).
- [32] M. G. A. Paris, Quantum estimation for quantum technology, *Int. J. Quantum Inf.* **7**, 125 (2009), arXiv:0804.2981.
- [33] D. S. P. Salazar, Benign projective landscapes for measured quantum divergences (2026), arXiv:2609.27767.
- [34] B. Regula and M. Tomamichel, Rethinking quantum smooth entropies: Tight one-shot analysis of quantum privacy amplification (2026), version 4, revised 1 September 2026, arXiv:2603.04493.
- [35] K. Temme, M. J. Kastoryano, M. B. Ruskai, M. M. Wolf, and F. Verstraete, The χ^2 -divergence and mixing times of quantum Markov processes, *J. Math. Phys.* **51**, 122201 (2010), corrected arXiv version 3 (2024), arXiv:1005.2358.
- [36] D. Petz, Quasi-entropies for finite quantum systems, *Rep. Math. Phys.* **23**, 57 (1986).
- [37] D. Petz, Monotone metrics on matrix spaces, *Linear Algebra Appl.* **244**, 81 (1996).
- [38] A. Lesniewski and M. B. Ruskai, Monotone Riemannian metrics and relative entropy on noncommutative probability spaces, *J. Math. Phys.* **40**, 5702 (1999), arXiv:math-ph/9808016.
- [39] Y. Hasegawa, Dynamics-independent bounds on state transformations and precision in open quantum systems (2025), arXiv:2508.13884.
- [40] D. S. P. Salazar, Universal thermodynamic uncertainty relation for quantum f -divergences, *Phys. Rev. E* **113**, 064150 (2026), arXiv:2511.10817.
- [41] See Supplemental Material at [URL will be inserted by publisher] for detailed proofs, model parameters, and numerical checks (2026).

- [42] H. P. Robertson, The uncertainty principle, *Phys. Rev.* **34**, 163 (1929).
- [43] D. S. P. Salazar, Matrix thermodynamic uncertainty relation for non-Abelian charge transport, *Phys. Rev. E* **114**, L022106 (2026), arXiv:2512.24956.
- [44] G. M. D’Ariano, V. Giovannetti, and P. Perinotti, Optimal estimation of quantum observables, *J. Math. Phys.* **47**, 022102 (2006), arXiv:quant-ph/0510187.
- [45] A. Miranowicz, M. Piani, P. Horodecki, and R. Horodecki, Inseparability criteria based on matrices of moments, *Phys. Rev. A* **80**, 052303 (2009), arXiv:quant-ph/0605001.
- [46] P. Talkner, E. Lutz, and P. Hänggi, Fluctuation theorems: Work is not an observable, *Phys. Rev. E* **75**, 050102 (2007), arXiv:cond-mat/0703189.
- [47] H. Spohn, Entropy production for quantum dynamical semigroups, *J. Math. Phys.* **19**, 1227 (1978).
- [48] G. T. Landi and M. Paternostro, Irreversible entropy production: From classical to quantum, *Rev. Mod. Phys.* **93**, 035008 (2021), arXiv:2009.07668.
- [49] I. A. Picatoste, A. Colla, and H.-P. Breuer, Local and global approaches to the thermodynamics of pure decoherence processes in open quantum systems, *Phys. Rev. A* **112**, 022210 (2025), arXiv:2506.11633.
- [50] J. P. Santos, L. C. Céleri, G. T. Landi, and M. Paternostro, The role of quantum coherence in non-equilibrium entropy production, *npj Quantum Inf.* **5**, 23 (2019), arXiv:1707.08946.
- [51] A. D. Varizi, A. P. Vieira, C. Cormick, R. C. Drumond, and G. T. Landi, Quantum coherence and criticality in irreversible work, *Phys. Rev. Research* **2**, 033279 (2020), arXiv:2004.00616.
- [52] H. J. D. Miller and J. Anders, Energy-temperature uncertainty relation in quantum thermodynamics, *Nat. Commun.* **9**, 2203 (2018), arXiv:1801.08057.
- [53] A. Ullah, M. T. Naseem, and Ö. E. Müstecaplıoğlu, Low-temperature quantum thermometry boosted by coherence generation, *Phys. Rev. Research* **5**, 043184 (2023), arXiv:2211.05461.
- [54] G. M. D’Ariano and P. Perinotti, Optimal data processing for quantum measurements, *Phys. Rev. Lett.* **98**, 020403 (2007), arXiv:quant-ph/0610058.