

Exclusive control of quantum memory erasure

Mir Alimuddin^{1,*}, Nathan Shettell^{2,3}, Raja Yehia,¹ Antonio Acín^{1,4} and Federico Centrone^{1,5,*}¹ICFO—*Institut de Ciències Fotoniques*, The Barcelona Institute of Science and Technology,
Av. Carl Friedrich Gauss 3, 08860 Castelldefels (Barcelona), Spain²MajuLab, CNRS-UCA-SU-NUS-NTU International Joint Research Laboratory, Singapore, Singapore³Centre for Quantum Technologies, National University of Singapore, 117543 Singapore, Singapore⁴ICREA—*Institució Catalana de Recerca i Estudis Avançats*, Lluís Companys 23, 08010 Barcelona, Spain⁵Universidad de Buenos Aires, *Instituto de Física de Buenos Aires (IFIBA)*, CONICET, Ciudad Universitaria,
1428 Buenos Aires, Argentina

(Received 16 March 2026; accepted 8 September 2026; published 20 October 2026)

Erasing memory is a fundamental operational task in quantum information processing, governed by Landauer's principle, which links information loss to thermodynamic work. We introduce and analyze *assisted quantum erasure*, where correlations with a remote system reduce the energetic cost of resetting a memory. We identify *exclusive control* of erasure as the central operational requirement: Only a designated party should be able to achieve the minimal cost, whereas any adversary must fail. In the device-dependent regime, we show that entanglement of formation exactly characterizes exclusivity, establishing entanglement as the decisive thermodynamic resource. Moving to a one-sided device-independent scenario, in which only the memory holder's device is trusted, we develop an operational erasure protocol based on random dephasing and conditional operations. Finally, in a fully device-independent setting, we show how Bell nonlocality and self-testing translate observed violations into lower bounds on any adversary's erasure capability, yielding a device-independent notion of exclusive thermodynamic control. Taken together, these results elevate quantum erasure from a thermodynamic constraint to an operational primitive: The erasure work cost quantifies secure, exclusive control over quantum memory, ensuring that an unauthorized agent cannot fully erase information under a bounded work budget.

DOI: [10.1103/6jtz-g59p](https://doi.org/10.1103/6jtz-g59p)

I. INTRODUCTION

The deep connection between information and thermodynamics has shaped foundational physics since Maxwell's demon challenged the second law [1]. Landauer's principle resolved this paradox by assigning an energetic cost to information erasure: at least $k_B T \ln 2$ of work per erased bit [2–5]. Erasure thus becomes a thermodynamically costly operation, making information a physical resource [6]. As quantum processors and networks scale, minimizing this cost becomes essential, especially under repeated operations constrained by finite energy budgets [7,8].

In the quantum regime, correlations between systems can be harnessed to reduce the work cost of erasure, allowing one party to lower the entropy of another's memory using only local operations and a classical message [9–13]. When this thermodynamic advantage exceeds what is possible with

classical correlations, it serves as a witness of nonclassicality. Erasure then becomes more than a resource sink: It is a tool for probing the structure and operational consequences of correlations in distributed quantum systems. Notably, the advantage need not be symmetric, since some parties may enable erasure at a lower cost than others. This raises a fundamental question: Can the ability to assist in low-cost erasure be made *exclusive* to a properly correlated party?

We formalize this idea as *exclusive control*: Only a designated party, by virtue of their quantum correlations with the system, should be able to assist erasure below a fixed thermodynamic threshold, while any other agent necessarily fails. In this view, erasure is not only a physical process but also a selective form of witnessing, enabled by privileged access to quantum information. Exclusivity is also practically important: In distributed scenarios, one often wants erasure to be enabled only when a designated party cooperates, even if other agents have partial access or side information. This mirrors quantum-cryptographic intuitions, in which only those sharing the appropriate quantum correlations gain secure access, whereas others are fundamentally constrained [14,15]. In our setting, untrusted agents may provide limited help, but the full thermodynamic advantage remains inaccessible without genuine quantum correlations.

We study three regimes where exclusivity can be guaranteed. In the *device-dependent* (DD) setting, where all operations are trusted, exclusivity is determined by entanglement:

*These authors contributed equally and share first authorship.

†Contact author: aliphy80@gmail.com‡Contact author: federico.centrone@icfo.eu

TABLE I. Summary of the three trust scenarios considered in this work, with the classical communication channel assumed to be untrusted in every case. In each scenario, Eve is treated as a competing helper, and the central question is whether she or any other party can match Alice's reduction of Bob's erasure cost under the same operational constraints.

Scenario	Trust assumptions	Quantity	Main statement
DD	Fully trusted state and devices	$W_{X \rightarrow B}$	Theorem 1: exclusivity iff an entanglement-of-formation gap holds.
SDI	Bob trusted, Alice untrusted	$\tilde{W}_{X \rightarrow B}$	Theorems 2 and 3: exclusivity follows from violating the LHS bound.
DI	Only Bell statistics trusted	$H(Z_B M_X)$	Bell statistics certify a lower bound on any adversary's erasure cost.

A trusted helper (Alice) outperforms an adversary (Eve) in assisted erasure precisely when Alice's entanglement with Bob exceeds Eve's. In the *semi-device-independent* regime, where only Bob's device is trusted, exclusivity is instead controlled by entropic uncertainty: An operational protocol based on random dephasing and conditional erasure leads to a state-independent lower bound on the assisted work cost that any local-hidden state (LHS) model must obey, and any observed violation of this bound guarantees exclusive control. Finally, in a fully *device-independent* (DI) scenario, where neither the source nor the measurement devices are characterized, Bell nonlocality and robust self-testing allow observed violations to be converted into lower bounds on the erasure work that any external agent must invest, providing a device-independent notion of exclusive thermodynamic control.

Taken together, these three regimes provide an operational and thermodynamic interpretation of the standard hierarchy of quantum correlations. Entanglement (DD), steering under uncertainty constraints (semi-DI), and Bell nonlocality (DI) each appear as the resource that determines who can make erasure cheap under a given level of trust. The rest of the article makes this hierarchy precise: We formalize assisted erasure and exclusivity, characterize exclusive control in the DD and semi-DI settings, show how to retain recoverability under adversarial attacks, and outline how Bell inequalities and self-testing yield device-independent bounds on exclusive thermodynamic control.

Table I summarizes the three trust scenarios considered in this work.

II. ASSISTED ERASURE

Throughout this work, we consider memories with a degenerate Hamiltonian. The equilibrium state is therefore maximally mixed and the reversible work cost is determined entirely by the entropy. Landauer's principle states that erasing information incurs a fundamental energy cost (and corresponding irreversible heat dissipation) [2–5]: Resetting a quantum memory B from ρ_B to a fixed pure state $|0\rangle$ requires, at least,

$$W_0 = k_B T S(B), \quad (1)$$

where $S(B) = -\text{Tr}[\rho_B \ln \rho_B]$ is the von Neumann entropy of B . Henceforth, we measure entropy in bits and we set $k_B T \ln 2 = 1$, so all work costs are measured in units of $k_B T \ln 2$. This baseline applies when the memory is treated in isolation and expresses a fundamental thermodynamic constraint on information processing [16].

In practice, erasure need not be performed in isolation. If Bob's memory B is correlated with a helper system A (by

Alice), suitable classical or quantum information processing can reduce Bob's *local* work cost [11–13, 17–20]. From Bob's perspective, the work required to reset B is lower, but the gain is fueled by the consumption of the initial AB correlations, so the global thermodynamic accounting remains consistent with Landauer's principle. In quantitative units, the optimal work bound is governed by the (classical or quantum) conditional entropy, $W_{\min} = S(B|A)$, highlighting that preestablished correlations carry a definite thermodynamic value [21–23]. Notably, in quantum communication settings, shared entanglement can make the conditional entropy negative, $S(B|A) < 0$, so that the corresponding erasure “cost” becomes negative and resetting B may even yield net work extraction [24]. In contrast, in the *classical-communication* regime of interest here, where Alice and Bob exchange only classical messages, the (classical) conditional entropy satisfies $S(B|A) \geq 0$, and the erasure cost is therefore necessarily non-negative.

Formally, Alice performs a measurement $A_M = \{\Pi_k\}$ on her subsystem and announces k . Bob then erases conditionally by bringing the system to the state $|0\rangle$. The minimal assisted work cost is

$$W_{A \rightarrow B} = \min_{A_M} \sum_k p_k S(\rho_{B|\Pi_k}) = S(B) - J_0(B|A), \quad (2)$$

with $p_k = \text{Tr}(\rho_A \Pi_k)$ and $\rho_{B|\Pi_k} = \text{Tr}_A[(\Pi_k \otimes I)\rho_{AB}]/p_k$. Here, $J_0(B|A) = S(B) - \min_{A_M} S(B|A_M)$ is the standard quantity used to quantify one-way classical correlations (from A to B), where the minimum is taken over measurements A_M on A [25]. Hence, any nonzero correlation accessible to Alice lowers Bob's *local* erasure cost relative to the unassisted baseline $W_0 = S(B)$.

Our main focus, however, is not merely cost reduction but exclusive control of the erasure: Bob's memory may also be correlated with an adversary Eve, who could in principle provide side information enabling a similarly low-cost reset. We therefore ask whether one can engineer a regime in which Alice retains exclusive control; i.e., she can assist Bob's erasure within his available work budget, while Eve cannot, despite being correlated with (B) . Quantum correlations, owing to their monogamy constraints, are natural candidates to enforce such asymmetric control.

Definition 1 (Exclusive erasure). Fix a per-run work budget $W_{\max}$ for Bob. We say that Alice holds *exclusive erasure* if

$$W_{A \rightarrow B} \leq W_{\max} < W_{E \rightarrow B} \leq W_0. \quad (3)$$

The cap must satisfy $W_{\max} < W_0$, otherwise exclusivity is vacuous, since Bob could always erase without any help by paying W_0 . Fig. 1 sketches this setting.




FIG. 1. Illustration of assisted erasure with competing parties. Bob's system begins in a high-entropy state, ρ_B (red). Alice and Eve each hold systems correlated with Bob's and perform local measurements (A_M and E_M) to assist in erasure. The resulting conditional states $\rho_{B|A_M}$ and $\rho_{B|E_M}$ show reduced entropy (blue and purple), reflecting Alice's and Eve's influence. Although both parties are correlated with Bob, only Alice achieves complete erasure, indicating a stronger correlation than Eve. The thermometer encodes the residual entropy of Bob's memory, visually distinguishing exclusive thermodynamic control. The figure depicts the direct-erasure realization in which the correlated register B itself is the memory; alternatively, B can serve as the auxiliary reset resource for a separate target memory M , as discussed in the text.

Equation (3) is the operational exclusivity we use throughout; the following results guarantee that there *exists* a nontrivial $W_{\max}$ between $W_{A \rightarrow B}$ and $W_{E \rightarrow B}$. For simplicity, we set Bob's capacity to $W_{\max} = W_{A \rightarrow B}$.

The above formulation also admits a useful operational realization in which the correlated register B is not itself the physical memory containing the information to be erased. Bob may instead hold an independent target memory M , while ρ_{AB} serves as an assisting thermodynamic resource. In this realization, Alice's measurement and announcement reduce the residual work $W_{A \rightarrow B}$ required to prepare B in a fixed pure state. Bob can then consume this pure auxiliary state in a trusted, input-independent reset (or irreversible) operation on M , for instance by swapping it with M and subsequently discarding or rethermalizing the auxiliary system. Eve can analogously attempt to prepare the same reset resource at cost $W_{E \rightarrow B}$. Hence, the exclusivity condition above determines which helper can enable the erasure of the target memory within Bob's available work budget. Importantly, this realization separates the system carrying the assisting correlations from the memory being erased and allows Bob to test the assisting resource before performing any irreversible operation on the stored information. The formal results below are expressed in terms of the assisted register B , and therefore apply both when B itself is the memory to be reset and when it is used as the thermodynamic resource for resetting a separate memory M .

In the device-dependent setting, all systems and measurement devices are fully characterized and trusted. We model any potential adversary, Eve, by assuming that she holds a system E , so that $\rho_{AB} = \text{Tr}_E[\rho_{ABE}]$. Classical correlations

between Alice and Bob may still lower Bob's local erasure cost, but they can be freely copied and distributed so that any such advantage can be reproduced by Eve. The following no-go lemma shows that separable states can therefore never generate exclusive erasure.

Lemma 1. If ρ_{AB} is separable, then any reduction of Bob's erasure cost achieved with Alice's assistance can also be reproduced by Eve. In particular, Alice cannot hold an exclusive erasure, i.e., $W_{A \rightarrow B} \geq W_{E \rightarrow B}$.

Together with Definition 1, this yields a nontrivial choice of $W_{\max}$, showing that entanglement is a necessary resource for exclusive erasure. For example, if Alice and Bob share a pure entangled state, then $W_{A \rightarrow B} = 0$ and, by monogamy, Bob is uncorrelated with any third party. Consequently, $W_{E \rightarrow B} = W_0$ and Alice enjoys strict exclusivity. For mixed states, the situation is subtler: Eve may share correlations with Bob via a purification $|\psi\rangle_{ABE}$. In that case, exclusivity requires that Alice's correlations with Bob dominate Eve's across all possible extensions. These correlations can be quantified using the entanglement of formation, an average measure of the entanglement (in the so-called entangled bits, or ebits) required to prepare the state. For a shared bipartite state $\sigma_{AB} = \sum_i p_i |\psi_i\rangle\langle\psi_i|$, the entanglement of formation [26] is defined as

$$E_f(A : B)_{\sigma_{AB}} := \min_{\{p_i, |\psi_i\rangle\}} \sum_i p_i S(\text{Tr}_A |\psi_i\rangle\langle\psi_i|).$$

The following theorem makes Alice's exclusive control on Bob's memory explicit in terms of entanglement of formation.

Theorem 1. Let Alice and Bob share ρ_{AB} and let $|\psi\rangle_{ABE}$ be a purification, with Eve holding the purifying system. Alice achieves a strictly lower erasure cost than Eve if and only if

$$E_f(A : B)_{\rho_{AB}} > E_f(E : B)_{\rho_{EB}}.$$

The proof of Lemma 1 and Theorem 1 is provided in the Appendix. Theorem 1 completes the DD characterization: When devices are trusted, exclusive thermodynamic control is governed entirely by an entanglement-of-formation *gap*. Alice can exert exclusive influence if and only if they are more entangled with Bob than any adversary; conversely, separability never suffices (Lemma 1). The proof follows from the Koashi-Winter relation [27], which connects competing classical correlations to competing entanglements of formation and ultimately leads to exclusive erasure.

III. STEERABILITY AS THE EXCLUSIVE RESOURCE

We now relax trust in Alice's device and their correlations, while keeping Bob's operations fully characterized. In this one-sided, semi-device-independent (SDI) regime, Alice's action is a possibly adversarial black box, so certification cannot rely on its internal form or on entanglement measures. It must come solely from the observed influence of Alice's announcements on Bob's system. Throughout, Bob's local state and operations are characterized and trusted, while the source and Alice's device are untrusted.

We adopt a four-step routine (see Fig. 2) that mirrors assisted erasure while enforcing a verifiable benchmark:

(1) Bob chooses at random one of two orthonormal bases $R = \{|r_i\rangle\}$ or $S = \{|s_j\rangle\}$ announcing the result to Alice. Bob




FIG. 2. One-sided semi-DI protocol: Bob's random dephasing, Alice's announcement, conditional erasure, and verification.

applies a dephasing channel Λ_R or Λ_S to B , corresponding to the chosen basis. This induces classical distributions $\vec{R} = \{r_i\}$ and $\vec{S} = \{s_j\}$ with $r_i = \langle r_i | \rho_B | r_i \rangle$, $s_j = \langle s_j | \rho_B | s_j \rangle$.

(2) Alice performs an (untrusted) measurement on ρ_A and announces a classical outcome. We denote the resulting classical registers by A_R and A_S , respectively.

(3) Bob applies a trusted, outcome-conditioned operation $\mathcal{E}$ to complete the erasure. Bob's budget in this step is limited to $W_{\max}$.

(4) Bob performs a cost-free verification test to certify whether the previous step has indeed erased the state.

The last step is crucial because Bob does not trust Alice's devices, the shared correlations with memory, or even the classical communication channel; therefore, he must verify that the state has indeed been erased (see the Appendix for details).

In general, the random dephasing of the first step of the routine injects entropy into Bob's system, raising the unassisted average erasure cost to

$$\tilde{W}_0 = \frac{1}{2}(H(R) + H(S)) \geq S(B) = W_0, \quad (4)$$

with equality only in special cases (e.g., $\rho_B = \mathbb{I}/d$ or when both the dephasing coincide with the spectral basis of ρ_B). Here, $H(X)$ is the Shannon entropy of the probability distribution of the outcomes measured in basis X . Crucially, if R and S are incompatible, the Maassen-Uffink relation [28] gives a nontrivial *state-independent* bound

$$\tilde{W}_0 \geq \frac{1}{2} C(R, S), \quad C(R, S) := \min_{i,j} \log_2 \frac{1}{|\langle r_i | s_j \rangle|^2}, \quad (5)$$

which depends only on the overlap of the two bases.

Dephasing in base R yields a classical-quantum state $\sigma_{AR}^{(R)} = \sum_i r_i \rho_{A|r_i} \otimes |i\rangle \langle i|_R$. While in Step 2, Alice's black-box outcome A_R carries at most $I_{\text{acc}}(A_R : R) \leq \chi_R = S(\sum_i r_i \rho_{A|i}) - \sum_i r_i S(\rho_{A|i})$ bits of *accessible* information about R [29,30], implying the classical conditional entropy

$$H(R|A_R) = H(R) - I_{\text{acc}}(A_R : R),$$

and analogously for base S . Optimizing over her black-box choices on each branch gives the operational assisted cost

$$\tilde{W}_{A \rightarrow B} = \frac{1}{2}(H(R|A_R) + H(S|A_S)) \leq \tilde{W}_0, \quad (6)$$

which depends only on observed announcements and is directly comparable to the state-independent benchmark $\frac{1}{2} C(R, S)$.

The random dephasing serves two purposes in the one-sided DI setting: (1) it turns Bob's quantum state into a classical register (R or S) whose uncertainty is quantified by Shannon entropies and constrained by the state-independent complementarity $C(R, S)$; (2) it yields an operational, device-agnostic metric of help $\tilde{W}_{X \rightarrow B}$ that depends only on observed announcements, not on trusting Alice's internal measurements or their initial shared correlations. The dephased classical register in R or S is then reset to $|0\rangle$ in Step 3 by applying an erasure operation chosen according to Alice's communication (Step 2) and the Bob's dephasing operation (Step 1). The physical implementation of this erasure map incurs the associated work cost.

To certify success without reading out the memory, in Step 4 Bob dephases the register in a fixed computational basis and couples it to a Szilard-type work-extraction device, i.e., a thermodynamic engine whose maximal extractable work is achieved only when the register is in a pure state and decreases with any residual entropy. Perfect erasure leaves a pure state and permits extraction of $\log_2 d$ work per run, while any residual mixing reduces the extractable work (see the Appendix).

In the following subsection, we leverage this setup: Whenever the observed work $\tilde{W}_{A \rightarrow B}$ drops below the complementarity benchmark $\frac{1}{2} C(R, S)$, Alice's advantage cannot be reproduced by any LHS model, which in turn underpins exclusive control in the one-sided DI setting. However, note that, even if the protocol passes verification (Step 4) and the observed assisted cost is strictly below the *unassisted* value $\tilde{W}_0$, this alone does not certify exclusive control. In particular, if the correlations between Alice and Bob admit an LHS model, the whole procedure can be simulated classically and reproduce the same apparent reduction. In that case, an adversary can match Alice's help. To capture this limitation, we first bound the assisted cost achievable under LHS models.

Theorem 2. If the shared state ρ_{AB} admits a local hidden state model (from A to B), then

$$\tilde{W}_{A \rightarrow B} \geq \frac{1}{2} C(R, S). \quad (7)$$

Intuitively, with an LHS model, Alice's announcements cannot jointly predict both dephased registers better than complementarity allows [31]; the residual uncertainty imposes a thermodynamic bound on assisted erasure. Using more than two incompatible dephasings on B raises this bound and detects a larger set of steerable states [32].

In the one-sided DI setting, classical models can again undermine exclusivity. Indeed:

Lemma 2. If ρ_{AB} admits an LHS model, then one can construct (for any Eve) a strategy with

$$\tilde{W}_{A \rightarrow B} = \tilde{W}_{E \rightarrow B}. \quad (8)$$

Thus, steerability is *necessary* for Alice's exclusive advantage. It is also *sufficient* once Alice's observed cost dips below the complementarity benchmark:

Theorem 3. If

$$\tilde{W}_{A \rightarrow B} < \frac{1}{2} C(R, S), \quad (9)$$

then Alice holds an exclusive advantage,

$$\tilde{W}_{A \rightarrow B} < \tilde{W}_{E \rightarrow B}. \quad (10)$$

Proof idea. The strict violation of the LHS bound in Eq. (7) certifies steering. Combining this with the entropic uncertainty relations with quantum memory (applied to Eve's side) yields $H(R|E) + H(S|E) \geq C(R, S)$, so any Eve must pay at least $\frac{1}{2} C(R, S)$ on average, whereas Alice pays less. Full proofs are provided in the Appendix. These additional noncommuting operations provide a sufficient strategy when neither the shared correlations nor Alice's measurement device is trusted. Equivalently, any correlations admitting an LHS model necessarily destroy exclusivity, establishing steering as a necessary resource. We show that the additional noncommuting operations are sufficient to guarantee both steering and exclusive erasure. Whether alternative methods of certifying steering can yield a stronger exclusive-erasure advantage in the semi-device-independent setting remains an open question.

In the next section, we show how, when combined with an authentication step that sacrifices some copies of the state, DD and SDI exclusive advantage implies that Eve cannot cause additional irrecoverable loss of AB correlations beyond what Bob explicitly chooses to erase.

IV. MEMORY RETRIEVAL AFTER ADVERSARIAL ATTACKS

In both DD and SDI settings, exclusivity implies that a malicious or untrusted helper cannot reproduce the honest assisted cost under the same work cap. In the external-memory realization discussed above, Bob performs the authorization test before acting on the target memory M , so a rejected request leaves the stored information untouched. The remaining question is therefore how much of the assisting AB resource must be consumed in order to authenticate the request. In the external-memory realization introduced above, a rejected request leaves the target memory M untouched; the results below instead quantify how much of the assisting AB resource can be preserved and reduce to memory recoverability when B itself is the memory. The following results quantify this preservation in the asymptotic regime.

In what follows, we work in an asymptotic regime consisting of identically and independently distributed (i.i.d.) rounds: Alice and Bob share n copies of a known state ρ_{AB} , Eve holds the purifying system E and may perform arbitrary operations and send classical messages, and Bob's laboratory and operations are treated as trusted and noiseless at the level of our idealized analysis.

In the asymptotic i.i.d. regime, the relevant resource is the total entropy of the memory state. For a known state ρ , each copy carries $S(\rho)$ units of entropy. A Schumacher-type compression [33] allows one to apply a global unitary that concentrates all of this entropy essentially into a smaller "noisy" subsystem, leaving the remaining registers in a standard pure state. A work budget $W < S(\rho)$ per copy then limits how much of that noisy subsystem can be irreversibly erased; the untouched part can be reversibly reencoded into fewer intact copies of ρ . In the Appendix, we formalize this as a compression-based erasure routine.

In our setting, this generic routine is combined with the protocol's *authentication* step, applied only to a random test subset of vanishing relative size. Bob never applies irreversible erasure to the remaining data block unless the compressed test block behaves exactly as predicted for Alice-Bob's honest strategy. Any deviation attributable to Eve is therefore detected during the authentication stage, before the data block is erased. As a result, unauthorized attempts cannot cause additional, undetected loss of Alice-Bob correlations beyond the vanishing test fraction and the part that Bob himself chooses to erase.

It is useful to distinguish the present authentication test from the verification step discussed in the semi-DI protocol. Here, Bob does not verify that erasure has already succeeded on the data block. Rather, he uses a sacrificial test subset to check whether the classical control information he receives is consistent with Alice's honest assistance before committing irreversible work on the remaining copies. In this sense, the subroutine plays the role of an authentication mechanism for the control message. This is relevant because, unlike standard quantum-cryptographic protocols such as quantum key distribution, we do not assume from the outset an authenticated classical channel. If such authentication were available in the fully device-dependent case, then Bob would simply ignore any message not certified to come from Alice, and exclusivity would already be enforced at the level of access to the control channel.

Theorem 4 (DD case). Let ρ_{AB} be shared by Alice and Bob with a purification $|\psi\rangle_{ABE}$, and assume the device-dependent exclusivity condition

$$W_{A \rightarrow B} < W_{E \rightarrow B}$$

holds for some per-run work cap $W_{\max}$ with $W_{A \rightarrow B} \leq W_{\max} < W_{E \rightarrow B}$. Suppose Bob implements the authentication routine on a random test subset of size $\varepsilon_n n$, where $\varepsilon_n \rightarrow 0$ and $n\varepsilon_n \rightarrow \infty$, and reserves the remaining copies as a data block. Then, under the i.i.d. and trust assumptions stated above, for any adversarial strategy by Eve that respects the same work cap, the following holds in the asymptotic limit ($n \rightarrow \infty$):

- (1) either the authentication step fails, in which case Bob aborts and keeps the untouched data block of size $(1 - \varepsilon_n)n + o(n)$,
- (2) or the authentication step succeeds, in which case the effective evolution on the data block is asymptotically indistinguishable from Alice's honest assistance, and Bob can safely erase the memory.

In particular, the recoverable fraction satisfies

$$R_n \geq 1 - \varepsilon_n - o(1) \xrightarrow{n \rightarrow \infty} 1,$$

and no adversarial strategy can induce additional irrecoverable loss of Alice-Bob correlations beyond Bob's intentional erasure and the vanishing test fraction.

Proof idea. Using the compression routine described above, Bob can reversibly compress the random test block conditioned on Alice's (or Eve's) announcement so that, in the honest case, the compressed block has an asymptotically well-defined amount of extractable work, equal to $\varepsilon_n n(1 - W_{A \rightarrow B}) + o(n\varepsilon_n)$. In the dishonest case, even though Bob applies the same honest compression maps, unitarity preserves

the total entropy of the compressed test block, so its total extractable work is upper bounded by $\varepsilon_n n(1 - W_{E \rightarrow B}) + o(n\varepsilon_n)$. Since $W_{A \rightarrow B} < W_{E \rightarrow B}$, these two asymptotic work rates are separated by a nonzero gap, and Bob can choose an acceptance threshold strictly between them. Hence, for sufficiently large n , the honest protocol is accepted while any dishonest strategy is rejected. If Bob rejects, only the test subset is consumed and the untouched data block is preserved; if he accepts, the data block can be processed exactly as in the honest protocol under the same work cap.

The same reasoning applies to the semi-device-independent protocol, where the resource is the *postdephasing* (basis-conditioned) state used in the assisted erasure routine.

Theorem 5 (Semi-DI case). Consider the one-sided device-independent protocol based on random dephasing in bases R, S , and let $\tilde{W}_{X \rightarrow B}$ denote the observed assisted per-run work cost averaged over the two bases for $X \in \{A, E\}$. Suppose the semi-DI exclusivity condition

$$\tilde{W}_{A \rightarrow B} < \tilde{W}_{E \rightarrow B}$$

holds for some admissible work cap. If, under the same i.i.d. and trust assumptions, Bob applies the authentication strategy on random test subsets of the R and S blocks, of total size $\varepsilon_n n$ with $\varepsilon_n \rightarrow 0$ and $n\varepsilon_n \rightarrow \infty$, then in the asymptotic limit any adversarial attempt either

(1) fails authentication, in which case Bob aborts and keeps the untouched data block of size $(1 - \varepsilon_n)n + o(n)$, or

(2) is asymptotically indistinguishable, on the data block, from Alice's honest assistance under the same work cap.

Consequently, the recoverable fraction again satisfies

$$R_n \geq 1 - \varepsilon_n - o(1) \xrightarrow{n \rightarrow \infty} 1,$$

and all *basis-conditioned* AB correlations that fuel the semi-DI protocol remain recoverable up to Bob's deliberate erasure and the vanishing test fraction.

In both regimes, exclusivity therefore comes with a strong robustness guarantee: As long as Bob adheres to the prescribed work cap and verifies only on a vanishing test fraction, no untrusted helper can inflict additional irretrievable damage on his memory beyond that sacrificial fraction. On the remaining data block, either the protocol aborts and the memory is kept untouched, or the posterasure state is asymptotically the one produced by the honest, Alice-assisted protocol under the same work cap. Any further loss of information arises solely from Bob's own deliberate erasure.

Example: two-qubit Werner state. To illustrate the above notions of exclusive control in both the DD and one-sided device-independent settings, we consider the family of two-qubit Werner states

$$\rho_{AB}(p) = p|\Psi^-\rangle\langle\Psi^-| + \frac{1-p}{4}\mathbb{1}_4, \quad (11)$$

with singlet fraction $p \in [0, 1]$. Bob's reduced state is maximally mixed, $\rho_B = \mathbb{1}/2$, so the unassisted erasure cost is $W_0 = S(B) = 1$.

In the DD regime, Alice's optimal assisted cost is

$$W_{A \rightarrow B} = h_2\left(\frac{1-p}{2}\right), \quad (12)$$




FIG. 3. Per-run work costs for assisted erasure on the two-qubit Werner state $\rho_W(p)$. The solid curve shows Alice's assisted cost in both DD and SDI scenarios $W_{A \rightarrow B} = \tilde{W}_{A \rightarrow B}$, the dashed curve shows Eve's assisted cost $W_{E \rightarrow B}^{DD}$ in the device-dependent scenario (given by the entanglement of formation), and the dotted horizontal line marks the one-sided device-independent LHS floor $\tilde{W}_{E \rightarrow B}^{LHS} = \frac{1}{2}$ arising from two mutually unbiased dephasings on Bob. Vertical lines indicate the known thresholds for entanglement ($p > 1/3$), steering ($p > 1/2$), and Bell nonlocality ($p > 1/\sqrt{2}$) [34]. The intersections $W_{A \rightarrow B} = W_{E \rightarrow B}^{DD}$ and $W_{A \rightarrow B} = \frac{1}{2}$ (not explicitly labeled here) mark the onset of exclusive thermodynamic control in the device-dependent and semi-device-independent regimes, witnessing entanglement and steering, respectively.

where h_2 is the binary entropy. Eve's assisted cost is given by the entanglement of formation $E_f(A : B)$, which for two-qubit Werner states takes the standard Wootters form

$$W_{E \rightarrow B}^{DD} = E_f(A : B) = h_2\left(\frac{1 + \sqrt{1 - C^2}}{2}\right), \quad (13)$$

with concurrence $C = \max(0, \frac{3p-1}{2})$. Exclusive DD control holds whenever $W_{A \rightarrow B} < W_{E \rightarrow B}^{DD}$.

In the one-sided semi-DI regime with two mutually unbiased bases on Bob's side (e.g., σ_x and σ_z), Alice's assisted cost coincides with her DD cost,

$$\tilde{W}_{A \rightarrow B} = W_{A \rightarrow B}, \quad (14)$$

while any LHS model must satisfy the state-independent bound

$$\tilde{W}_{E \rightarrow B} \geq \frac{1}{2}C(\sigma_x, \sigma_z) = \frac{1}{2}, \quad (15)$$

so semi-DI exclusivity holds whenever $W_{A \rightarrow B} < \frac{1}{2}$.

Figure 3 shows $W_{A \rightarrow B}$, $W_{E \rightarrow B}^{DD}$, and the semi-DI LHS floor $\tilde{W}_{E \rightarrow B}^{LHS} = \frac{1}{2}$ as functions of the Werner parameter p , together with the known thresholds for entanglement, steering, and Bell nonlocality [34]. The crossings $W_{A \rightarrow B} = W_{E \rightarrow B}^{DD}$ and $W_{A \rightarrow B} = \frac{1}{2}$ mark, respectively, the onset of exclusive thermodynamic control in the DD and semi-DI regimes. Interestingly, for this highly symmetric family the DD and semi-DI thresholds lie relatively close, indicating that relaxing device assumptions does not dramatically enlarge the parameter range of exclusive control for Werner states. The full derivation of these expressions is given in the Appendix.

Device-independent thermodynamic control. We now ask whether exclusive thermodynamic control can be certified

without characterizing either the source or the measurement devices. In the fully DI setting, Alice and Bob interact only with black boxes and base their conclusions exclusively on the observed input-output statistics. As usual in a DI protocol, the quantum devices are untrusted, while Bob's classical processing and the thermodynamic apparatus used to store and erase the output register are assumed to operate as prescribed.

Consider a standard Bell scenario in which Alice and Bob choose inputs x, y and obtain outputs a, b , defining an observed behavior $P(a, b|x, y)$. For designated settings, Bob stores his output in a physical classical memory Z_B , while Z_A denotes Alice's corresponding classical register. An adversary Eve is allowed to hold an arbitrary quantum system E purifying the underlying realization. Although the state, Hilbert-space dimension, and measurements inside the black boxes are unknown, device-independent methods can infer from the observed behavior a lower bound

$$H(Z_B|E) \geq h_{\text{DI}}[P], \quad (16)$$

where $h_{\text{DI}}[P]$ is obtained solely from the observed Bell statistics [35,36].

This bound has a direct thermodynamic interpretation. Any party assisting Bob is required to communicate with his laboratory through a classical message M_X , with $X = A, E$. In our units, the residual work required to erase Bob's physical output memory after receiving such a message is

$$W_{X \rightarrow B}^{\text{DI}} = H(Z_B|M_X). \quad (17)$$

For Eve, the message M_E can be generated by an arbitrary measurement or classical processing of her quantum side information E . By data processing,

$$H(Z_B|M_E) \geq H(Z_B|E), \quad (18)$$

and therefore,

$$W_{E \rightarrow B}^{\text{DI}} \geq H(Z_B|E) \geq h_{\text{DI}}[P]. \quad (19)$$

Hence, for every quantum realization compatible with the observed Bell behavior, no external party holding the purification can reduce the residual entropy, and therefore the erasure work of Bob's physical memory, below the device-independent value $h_{\text{DI}}[P]$. This is the thermodynamic content of the DI certification: The internal physical realization of the black boxes need not be known in order to bound how much an adversary can lower the cost of an irreversible operation in Bob's laboratory.

For the legitimate helper, Alice may simply communicate her observed output, $M_A := Z_A$. This gives an explicit achievable cost

$$W_{A \rightarrow B}^{\text{DI}} \leq H(Z_B|Z_A), \quad (20)$$

where $H(Z_B|Z_A)$ is directly determined from the observed Alice-Bob statistics. Consequently, whenever

$$H(Z_B|Z_A) < h_{\text{DI}}[P], \quad (21)$$

Alice necessarily achieves a strictly lower erasure cost than any adversary,

$$W_{A \rightarrow B}^{\text{DI}} < W_{E \rightarrow B}^{\text{DI}}. \quad (22)$$

Bob can therefore choose a work threshold satisfying

$$W_{A \rightarrow B}^{\text{DI}} \leq W_{\text{max}} < W_{E \rightarrow B}^{\text{DI}}, \quad (23)$$

which certifies exclusive thermodynamic control directly from the observed statistics. Notice that the relevant condition involves both the device-independent lower bound on Eve's residual uncertainty and Alice's observed correlation with Bob; a nonzero Bell violation need not by itself guarantee a positive erasure-cost gap.

For near-maximal Bell violations, self-testing provides a particularly transparent realization of this separation. Robust self-testing results imply that correlations close to the maximal CHSH violation constrain the underlying state and measurements, up to local isometries, to be close to the ideal maximally entangled realization [37–40]. The corresponding Alice-Bob outputs can therefore be made strongly correlated, making $H(Z_B|Z_A)$ small, while the same nonlocal statistics impose a device-independent lower bound on Eve's uncertainty $H(Z_B|E)$. Near the ideal point the two effects combine to yield a finite thermodynamic gap between the work accessible with Alice's assistance and that achievable by any external party.

Importantly, the role of Bell nonlocality here is not merely to certify private information. Through Landauer's principle, the device-independent entropy bound becomes a lower bound on a physical work cost: Irrespective of which quantum state and measurements realize the black-box behavior, Eve cannot assist Bob in erasing the stored output below the certified thermodynamic threshold. Optimizing this bound over the classical messages that may enter Bob's laboratory, or deriving tighter work-Bell trade-offs through semidefinite-program relaxations, could strengthen the quantitative certification. Such refinements are not required, however, for the basic DI conclusion established here: Observed nonlocal correlations can certify an exclusive separation in the thermodynamic ability of different parties to assist an irreversible memory-erasure process.

V. DISCUSSION

We introduced *exclusive control of erasure* as an operational criterion for identifying which parties can lower the erasure cost below a fixed budget. In the device-dependent setting, we established a necessary and sufficient condition: Exclusivity arises precisely when the designated helper possesses a strictly larger *entanglement of formation* with Bob than any adversary, implying that classical or separable states cannot yield exclusivity. In the one-sided semi-device-independent regime, exclusivity is instead governed by an *entropic uncertainty threshold*. Random dephasing in two incompatible bases induces a state-independent lower bound $\frac{1}{2}C(R, S)$ on the assisted work cost that any LHS model must obey; only helpers whose announcements beat this bound can achieve exclusive control. At the structural level, the correlations that allow one to surpass the uncertainty threshold are necessarily the steerable ones, but operationally, the relevant parameter is the uncertainty gap set by $C(R, S)$.

More generally, the correlated subsystem need not store the information being erased; it may instead serve as an assisting thermodynamic resource that authorizes the reset of a separate

memory, thereby separating the resource that provides exclusive control from the physical register on which the irreversible operation is ultimately performed.

Beyond characterizing when exclusivity holds, we showed that when combined with an authentication step, exclusivity guarantees that malicious or untrusted intervention cannot cause irrecoverable damage beyond a vanishing sacrificial test fraction and Bob's declared work budget. In both DD and semi-DI settings, the compression routine is applied only to a random test subset, whose total extractable work is compared with the honest expectation. This allows Bob to treat any external interaction as either (1) a detected failure, in which case he aborts and preserves the untouched data block, or (2) an effective honest run, in which the data block is asymptotically processed exactly as it would be under Alice's assistance alone and the same cap $W_{\max}$. Thus, on accepted runs, Eve cannot force additional loss of AB correlations on the data block: Bob never loses more than he *chooses* to erase, apart from the vanishing test fraction used for authentication. This elevates erasure from a passive thermodynamic cost to an active primitive for access control, in which "who can erase" is itself a protected resource.

Our framework naturally admits broader operational connections. Suppose the designated helper is allowed to send quantum, rather than purely classical, messages. In that case, exclusive erasure intersects with the notion of negative conditional entropy [24], in which quantum side information enables reductions in erasure cost that are unattainable with classical correlations alone. Analyzing exclusivity in this regime could enlarge the class of states that support exclusive thermodynamic control in the device-dependent setting. Moreover, our entropic characterization of the one-sided semi-device-independent protocol resonates with the use of entropic uncertainty relations in quantum cryptography [32], suggesting that exclusive erasure can serve as a new primitive for resource-constrained cryptographic architectures, in which thermodynamic work bounds act as security witnesses even under limited trust in devices and operations. More broadly, this perspective is consistent with other recent efforts to give steering a direct operational meaning in physical tasks, including metrological and dynamical settings [41–43]. Taken together, these connections show that the hierarchy of quantum correlations is also a hierarchy of exclusive thermodynamic control over distant memories.

See the Appendix for detailed proofs and additional derivations.

ACKNOWLEDGMENTS

This work was supported by the Government of Spain (Severo Ochoa CEX2019-000910-S, FUNQIP, and NextGeneration EU PRTR-C17.I1), European Union (QSNP, 101114043, Quanteria Veriqtas), Fundació Cellex, Fundació Mir-Puig, and Generalitat de Catalunya (CERCA program). M.A. acknowledges funding by the European Union (QURES, 101153001). N.S. was supported by the National Research Foundation, Singapore, through the National Quantum Office, hosted in A*STAR, under its Centre for Quantum Technologies Funding Initiative (S24Q2d0009). In addition, N.S. acknowledges the Plan France 2030 through the project

BACQ as part of the MetriQs-France program (Grant No. ANR-22-QMET-0002). R.Y. acknowledges funding from the Horizon Europe program (QUCATS) as well as LUBOS. A.A. acknowledges the ERC AdG CERQUITE, the AXA Chair in Quantum Information Science. F.C. acknowledges funding by the European Union (EQC, 101149233). Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

DATA AVAILABILITY

No datasets were generated or analyzed in this study. All numerical results and figures were obtained directly from the analytical expressions presented in the manuscript.

APPENDIX A

We use the notation and definitions introduced in the main text throughout.

1. Proof of Lemma 1

Before proving the theorem, we can prove Lemma 1. It is a specific case of the above theorem, but sets the stage for the formalism we will employ.

Proof. A classical ensemble admits a decomposition

$$\rho_{AB} = \sum_i p_i |\psi_i\rangle\langle\psi_i|_A \otimes |\phi_i\rangle\langle\phi_i|_B. \quad (\text{A1})$$

Since the state is separable, we may straightforwardly make the following tripartite extension:

$$\rho_{ABE} = \sum_i p_i |\psi_i\rangle\langle\psi_i|_A \otimes |\phi_i\rangle\langle\phi_i|_B \otimes |\psi_i\rangle\langle\psi_i|_E, \quad (\text{A2})$$

for which $\rho_{AB} = \rho_{AE}$. This shows that Eve can replicate Alice's correlations with Bob and achieve the same erasure cost. Moreover, Eve can assign a perfectly distinguishable register that encodes the classical index i , e.g.,

$$\sigma_{ABE} = \sum_i p_i |\psi_i\rangle\langle\psi_i|_A \otimes |\phi_i\rangle\langle\phi_i|_B \otimes |i\rangle\langle i|_E, \quad (\text{A3})$$

where $\langle i|j\rangle = \delta_{i,j}$. In this case, $J_o(B|E) = S(B) \geq J_o(B|A)$, which implies

$$W_{A \rightarrow B} \geq W_{E \rightarrow B}. \quad (\text{A4})$$

Thus, classical correlations alone never grant Alice exclusive control over Bob's memory. ■

2. Proof of Theorem 1

Now, we show the proof of Theorem 1.

Proof. Consider the cost gap

$$\begin{aligned} \Delta &= W_{E \rightarrow B} - W_{A \rightarrow B} \\ &= J_o(B|A) - J_o(B|E). \end{aligned} \quad (\text{A5})$$

The Koashi-Winter relation [27] states that

$$J_o(B|A) + E_f(E : B) = J_o(B|E) + E_f(A : B) = S(B). \quad (\text{A6})$$

Substituting, we find

$$\Delta = E_f(A : B) - E_f(E : B). \quad (\text{A7})$$

Hence, Alice achieves exclusive advantage, i.e., $\Delta > 0$, if and only if $E_f(A : B) > E_f(E : B)$, independent of the purification. ■

3. Verification step in the SDI setting

In the one-sided semi-device-independent setting, Bob's operations are trusted, but neither the shared state nor Alice's measurement device are assumed to be reliable. As a result, even if Alice's announcement lowers Bob's *average* assisted erasure cost, it does not by itself guarantee that Bob's memory has actually been reset to the target pure state in each run. A verification step is therefore necessary: Without it, Bob could apply the prescribed correction and still end up with a mixed state, while incorrectly concluding that erasure has succeeded.

A simple example illustrates the point. Suppose Alice and Bob are meant to share the Bell state

$$|\phi^+\rangle = \frac{|00\rangle + |11\rangle}{\sqrt{2}},$$

and Bob follows the honest protocol based on Alice's announced basis and outcome. If the shared state and Alice's measurements are indeed the intended ones, Bob's conditional correction resets his memory to $|0\rangle$ at zero work cost. However, if Alice's device measures in the wrong basis, or if the shared state differs from the expected one, then the same announced message may leave Bob's conditional state mixed. In that case, the prescribed unitary correction no longer produces the erased state, even though the observed assisted cost may still satisfy $\tilde{W}_{A \rightarrow B} < \tilde{W}_0$.

For this reason, in the semi-DI protocol Bob must verify that the reset has actually succeeded. He does so by testing the final memory after the erasure step. In the honest case, the memory is in the pure state $|0\rangle$ and Bob can extract the maximal work $\log d$ from it using an ideal work extractor. Any shortfall certifies that the reset was incomplete, either because the shared state was not the intended one, or because Alice's device was faulty. In such cases, the protocol must be aborted (or Bob must pay additional work to complete the erasure).

4. Proof of Theorem 2

Proof. Assume that ρ_{AB} admits an LHS model from A to B . Then, there exist a classical variable λ with distribution $\{p_\lambda\}$, a family of states $\{\rho_B^\lambda\}$ on Bob's system, and classical response probabilities $p(a|x, \lambda)$ such that, for every measurement setting x on A and outcome a , Bob's (unnormalized) conditional state can be written as

$$\sigma_{a|x} := \text{Tr}_A[(M_{a|x} \otimes \mathbb{1}) \rho_{AB}] = \sum_\lambda p_\lambda p(a|x, \lambda) \rho_B^\lambda. \quad (\text{A8})$$

In words: Before any measurement choice is made, Bob's system is drawn from a preexisting ensemble $\{\rho_B^\lambda\}$, while Alice's announced outcome is generated by classical postprocessing of λ .

We now consider the semi-DI protocol in the main text, in which Bob chooses uniformly at random one of two dephasing settings R or S (with orthonormal bases $R = \{|r_i\rangle\}$ and $S = \{|s_j\rangle\}$), applies the corresponding dephasing to B , and Alice announces a classical message (denoted A_R or A_S depending on the setting). The semi-DI assisted work cost is

defined operationally as

$$\tilde{W}_{A \rightarrow B} = \frac{1}{2}(H(R|A_R) + H(S|A_S)), \quad (\text{A9})$$

where $H(\cdot|\cdot)$ are classical conditional Shannon entropies of Bob's dephasing outcome given Alice's announcement.

Fix λ . If Bob dephases in basis R , then the outcome distribution is

$$p_i^{(R, \lambda)} := \langle r_i | \rho_B^\lambda | r_i \rangle, \quad \text{and similarly} \quad p_j^{(S, \lambda)} := \langle s_j | \rho_B^\lambda | s_j \rangle.$$

Since Alice's announcement is obtained from λ by classical postprocessing (and does not give access to anything beyond λ), it cannot reduce Bob's uncertainty below what would be possible if λ itself were known. Hence,

$$\begin{aligned} H(R|A_R) &\geq \sum_\lambda p_\lambda H(\bar{p}^{(R, \lambda)}), \\ H(S|A_S) &\geq \sum_\lambda p_\lambda H(\bar{p}^{(S, \lambda)}), \end{aligned} \quad (\text{A10})$$

where $\bar{p}^{(R, \lambda)} = \{p_i^{(R, \lambda)}\}_i$ and $\bar{p}^{(S, \lambda)} = \{p_j^{(S, \lambda)}\}_j$. Substituting Eq. (A10) into Eq. (A9) gives

$$\tilde{W}_{A \rightarrow B} \geq \frac{1}{2} \sum_\lambda p_\lambda (H(\bar{p}^{(R, \lambda)}) + H(\bar{p}^{(S, \lambda)})). \quad (\text{A11})$$

Now, for each fixed λ , the two distributions $\bar{p}^{(R, \lambda)}$ and $\bar{p}^{(S, \lambda)}$ arise from measuring the *same* state ρ_B^λ in the two bases R and S . Therefore, the Maassen-Uffink entropic uncertainty relation yields

$$\begin{aligned} H(\bar{p}^{(R, \lambda)}) + H(\bar{p}^{(S, \lambda)}) &\geq C(R, S), \\ C(R, S) &:= \min_{i,j} \log \frac{1}{|\langle r_i | s_j \rangle|^2}. \end{aligned} \quad (\text{A12})$$

Averaging Eq. (A12) over λ and using Eq. (A11) concludes that

$$\tilde{W}_{A \rightarrow B} \geq \frac{1}{2} C(R, S),$$

which is exactly Eq. (7). This proves the theorem. ■

5. Proof of Lemma 2

Proof. Assume that ρ_{AB} admits an LHS model from A to B . Then there exist a random variable λ with distribution $\{p_\lambda\}$, a collection of states $\{\sigma_B^\lambda\}$, and response functions $q_A(k|X, \lambda)$ (for $X \in \{R, S\}$) such that the (unnormalized) conditional states on Bob satisfy

$$\sigma_{B|k, X} = \sum_\lambda p_\lambda q_A(k|X, \lambda) \sigma_B^\lambda. \quad (\text{A13})$$

In particular, after Bob applies the dephasing Λ_X and regards B as the classical register X_B , the induced joint distribution of (X_B, A_X) is

$$P(i, k|X) = \sum_\lambda p_\lambda q_A(k|X, \lambda) \langle x_i | \sigma_B^\lambda | x_i \rangle,$$

where $\{|x_i\rangle\}$ is the basis defining Λ_X .

Now, consider the tripartite extension in which Eve holds a classical register storing λ . Upon learning X , Eve outputs a message E_X by sampling $q_E(k|X, \lambda) := q_A(k|X, \lambda)$. Then, the induced distribution of (X_B, E_X) is *identical* to that of (X_B, A_X) for each $X \in \{R, S\}$, hence $H(X_B|E_X) = H(X_B|A_X)$ for both

settings. Averaging over R, S gives

$$\begin{aligned}\tilde{W}_{E \rightarrow B} &= \frac{1}{2}(H(R|E_R) + H(S|E_S)) \\ &= \frac{1}{2}(H(R|A_R) + H(S|A_S)) = \tilde{W}_{A \rightarrow B}.\end{aligned}$$

This constructs an Eve strategy matching Alice's semi-DI assisted cost, proving the lemma. ■

6. Proof of Theorem 3

Proof. Let

$$\Sigma_A := H(R|R') + H(S|S'), \quad \Sigma_E := H(R|R'') + H(S|S''),$$

so that

$$\tilde{W}_{A \rightarrow B} = \frac{\Sigma_A}{2}, \quad \tilde{W}_{E \rightarrow B} = \frac{\Sigma_E}{2}.$$

The hypothesis of the theorem,

$$\tilde{W}_{A \rightarrow B} < \frac{C(R, S)}{2},$$

is equivalent to

$$\Sigma_A < C(R, S). \quad (\text{A14})$$

Now, let E denote Eve's quantum side information before she measures it to generate the classical outcomes R'' and S'' . Since R'' and S'' are obtained from E by measurement and classical postprocessing, data processing implies

$$H(R|R'') \geq H(R|E), \quad H(S|S'') \geq H(S|E). \quad (\text{A15})$$

Similarly, since R' and S' are classical outcomes obtained by measuring Alice's quantum system A ,

$$H(R|R') \geq H(R|A), \quad H(S|S') \geq H(S|A). \quad (\text{A16})$$

For any tripartite extension ρ_{ABE} , Berta *et al.* [44] proved the entropic uncertainty relations with quantum memory

$$H(R|E) + H(S|A) \geq C(R, S) \quad (\text{A17})$$

and

$$H(S|E) + H(R|A) \geq C(R, S). \quad (\text{A18})$$

Adding Eqs. (A17) and (A18) gives

$$H(R|E) + H(S|E) + H(R|A) + H(S|A) \geq 2C(R, S).$$

Using Eq. (A16), we obtain

$$H(R|E) + H(S|E) \geq 2C(R, S) - \Sigma_A.$$

Finally, by Eq. (A15),

$$\begin{aligned}\Sigma_E &= H(R|R'') + H(S|S'') \\ &\geq H(R|E) + H(S|E) \geq 2C(R, S) - \Sigma_A.\end{aligned}$$

Since $\Sigma_A < C(R, S)$ by Eq. (A14), it follows that

$$\Sigma_E > \Sigma_A.$$

Multiplying by $1/2$ yields

$$\tilde{W}_{E \rightarrow B} > \tilde{W}_{A \rightarrow B},$$

which proves the theorem. ■

7. Derivation details for the Werner state example

We now illustrate our framework using the two-qubit Werner state [45]

$$\rho_{AB}(p) = p|\Psi^-\rangle\langle\Psi^-| + \frac{1-p}{4}\mathbb{1}_4 \quad (\text{A19})$$

$$= \frac{1}{4} \left(\mathbb{1}_A \otimes \mathbb{1}_B - p \sum_{i=x,y,z} \sigma_{iA} \otimes \sigma_{iB} \right), \quad (\text{A20})$$

where $|\Psi^-\rangle$ is the singlet state and σ_i are the Pauli matrices. The Pauli parametrization is particularly convenient for writing the conditional states after local measurements [46,47].

The unconditional local state of Bob is maximally mixed,

$$\rho_B = \text{Tr}_A[\rho_{AB}] = \frac{\mathbb{1}_2}{2},$$

so its stand-alone erasure cost is $W_0 = S(B) = 1$ (in $k_B T$ units).

The conditional local state after Alice's projective measurement and announcement is

$$\rho_{B|A}^\pm = \text{Tr}_A[(\Pi_A^\pm \otimes \mathbb{1}_B) \rho_{AB}], \quad (\text{A21})$$

where $\Pi_A^\pm = \frac{1}{2}(\mathbb{1} \pm \sum_i r_i \sigma_i)$ defines Alice's local measurement along Bloch direction $\vec{r} = (r_x, r_y, r_z)$, and the sign corresponds to the $\pm$ outcome. For measurements along a Pauli axis, one finds explicitly

$$\rho_{B|A}^\pm = \frac{1}{2} \begin{pmatrix} 1 \mp p & 0 \\ 0 & 1 \pm p \end{pmatrix}. \quad (\text{A22})$$

The assisted work cost is

$$W_{A \rightarrow B} = S(B) - J(B|A) \quad (\text{A23})$$

$$= 1 - [2 - S(\rho_{AB}) - D(B|A)]$$

$$= S(\rho_{AB}) + D(B|A) - 1, \quad (\text{A24})$$

where we used $S(B) = 1$ for two-qubit Werner states and the identity $J(B|A) = 2 - S(\rho_{AB}) - D(B|A)$ between one-way classical correlation and discord.

For the Werner family,

$$S(\rho_{AB}) = -\frac{1+3p}{4} \log_2 \frac{1+3p}{4} - 3 \frac{1-p}{4} \log_2 \frac{1-p}{4},$$

$$\begin{aligned}D(B|A) &= \frac{1-p}{4} \log_2(1-p) - \frac{1+p}{2} \log_2(1+p) \\ &\quad + \frac{1+3p}{4} \log_2(1+3p).\end{aligned}$$

Inserting these into Eq. (A23) and simplifying yields

$$W_{A \rightarrow B} = 1 - \frac{1-p}{2} \log_2(1-p) - \frac{1+p}{2} \log_2(1+p) \quad (\text{A25})$$

$$= h_2\left(\frac{1-p}{2}\right), \quad (\text{A26})$$

where $h_2(x) = -x \log_2 x - (1-x) \log_2(1-x)$ is the binary entropy.

To compute the work $W_{E \rightarrow B}$ associated with an adversarial helper in the DD scenario, we use that in our setting it is given

by the entanglement of formation of the Werner state [48],

$$W_{E \rightarrow B} = E_f(A : B), \quad (\text{A27})$$

with

$$E_f(A : B) = h_2\left(\frac{1 + \sqrt{1 - C^2}}{2}\right), \quad (\text{A28})$$

where the concurrence is $C = \max(0, \frac{3p-1}{2})$.

In the semi-DI scenario, we dephase Bob in two mutually unbiased bases, say, the eigenbases of σ_z and σ_x , and define

$$\tilde{W}_{A \rightarrow B} = \frac{1}{2}(H(\sigma_{z_B}|\sigma_{z_A}) + H(\sigma_{x_B}|\sigma_{x_A})), \quad (\text{A29})$$

where $H(\cdot|\cdot)$ are classical conditional entropies of the joint outcome distributions. For the Werner state, the error probability in any Pauli basis is $(1-p)/2$, so each contribution equals $h_2(\frac{1-p}{2})$, and hence

$$\tilde{W}_{A \rightarrow B} = W_{A \rightarrow B}. \quad (\text{A30})$$

On the other hand, Eve's assisted erasure cost in the one-sided DI setting is bounded below by the LHS floor from Theorem 2,

$$\tilde{W}_{E \rightarrow B} \geq \frac{1}{2}C(\sigma_z, \sigma_x) = \frac{1}{2}\log_2 2 = \frac{1}{2}. \quad (\text{A31})$$

Putting these expressions together, the per-run work costs $W_{A \rightarrow B}$, $W_{E \rightarrow B}^{\text{DD}} = E_f$, and the semi-DI LHS bound $\tilde{W}_{E \rightarrow B}^{\text{LHS}} = \frac{1}{2}$ as functions of the Werner parameter p are shown in Fig. 3.

APPENDIX B: MEMORY ERASURE WITH BOUNDED ENERGY

In this Appendix, we recall a standard state-transformation routine that we use as a building block in our recoverability arguments. It captures the idea that, for a *known* memory state and a bounded work budget, one can reversibly compress the entropy into a smaller subsystem, erase part of it, and then recover a reduced number of intact copies of the original state.

Lemma B1. Consider n copies of a quantum memory state $\rho \in D(\mathcal{H}_d)$ that we wish to erase, where the available work per copy is bounded by $W \leq S(\rho)$. Then, in the asymptotic limit $n \rightarrow \infty$, there exists a protocol consisting of

(1) unitaries U_1 and U_2 (thermodynamically reversible), and

(2) an irreversible erasure channel Λ_{irr} applied to a subset of registers,

such that the overall map

$$\mathcal{E} = U_2 \circ \Lambda_{\text{irr}} \circ U_1$$

implements

$$\rho^{\otimes n} \xrightarrow{\mathcal{E}} \rho^{\otimes n(1-W/S(\rho))},$$

i.e., the protocol irreversibly removes exactly a fraction $W/S(\rho)$ of the total entropy and leaves a fraction

$$R = 1 - \frac{W}{S(\rho)}$$

of the memory intact, in the form of $(1 - W/S(\rho))n$ effective copies of ρ .

Proof. Suppose Bob holds n copies of $\rho \in D(\mathcal{H}_d)$. The total entropy is $nS(\rho)$, and the total work budget is nW .

Step 1. By the standard Shannon-Schumacher compression picture, there exists a unitary U_1 (possibly acting jointly on the n systems and an ancilla) such that

$$\rho^{\otimes n} \xrightarrow{U_1} |0\rangle\langle 0|^{\otimes n(1-\frac{S(\rho)}{\log d})} \otimes \left(\frac{I}{d}\right)^{\otimes \frac{nS(\rho)}{\log d}}.$$

The first factor is a large pure register with no entropy, and the second factor is a tensor product of maximally mixed d -dimensional registers containing essentially all the entropy. Since U_1 is unitary, this step is reversible and costs no work.

Step 2. Erasing one d -dimensional maximally mixed register costs $\log d$ units of work. With total budget nW , Bob can therefore erase

$$\frac{nW}{\log d}$$

such registers. Define the irreversible erasure channel

$$\Lambda_{\text{irr}}(\sigma) = |0\rangle\langle 0| \quad \forall \sigma \in D(\mathcal{H}_d)$$

and apply it to exactly $\frac{nW}{\log d}$ of the maximally mixed registers:

$$\left(\frac{I}{d}\right)^{\otimes \frac{nS(\rho)}{\log d}} \xrightarrow{\Lambda_{\text{irr}}^{\otimes \frac{nW}{\log d}}} |0\rangle\langle 0|^{\otimes \frac{nW}{\log d}} \otimes \left(\frac{I}{d}\right)^{\otimes \frac{nS(\rho)}{\log d}(1-\frac{W}{S(\rho)})}.$$

This consumes exactly nW units of work.

Step 3 (Reversible recovery of fewer copies). The remaining maximally mixed registers carry a total entropy of

$$nS(\rho)\left(1 - \frac{W}{S(\rho)}\right) = n(S(\rho) - W).$$

By Schumacher dilation, there exists a unitary U_2 such that

$$|0\rangle\langle 0|^{\otimes *} \otimes \left(\frac{I}{d}\right)^{\otimes \frac{nS(\rho)}{\log d}(1-\frac{W}{S(\rho)})} \xrightarrow{U_2} \rho^{\otimes n(1-W/S(\rho))},$$

where the precise number of $|0\rangle$ registers is irrelevant. Again, U_2 is reversible and cost-free.

Composing the three steps, the overall transformation is

$$\rho^{\otimes n} \xrightarrow{\mathcal{E} = U_2 \circ \Lambda_{\text{irr}} \circ U_1} \rho^{\otimes n(1-W/S(\rho))},$$

so a fraction $W/S(\rho)$ of the total entropy is erased and a fraction $R = 1 - W/S(\rho)$ of the memory remains in intact copies of ρ . ■

1. Proof of Theorem 4

Proof. Suppose Bob receives n copies of a memory state. In the honest case, Alice's classical communication induces an ensemble

$$\{p_i^A, \rho_i^A\}_{i=1}^N, \quad p_i^A = \frac{n_i^A}{n},$$

and the average work required to erase Bob's memory is

$$W_{A \rightarrow B} = \sum_i p_i^A S(\rho_i^A).$$

In the dishonest case, Eve may instead induce an unknown ensemble

$$\{p_i^E, \rho_i^E\}_{i=1}^N, \quad p_i^E = \frac{n_i^E}{n},$$

with corresponding cost

$$W_{E \rightarrow B} = \sum_i p_i^E S(\rho_i^E).$$

We assume the exclusivity condition

$$W_{A \rightarrow B} < W_{E \rightarrow B},$$

and define the gap

$$\Delta := W_{E \rightarrow B} - W_{A \rightarrow B} > 0.$$

Let $\varepsilon_n \in (0, 1)$ be such that

$$\varepsilon_n \rightarrow 0, \quad n\varepsilon_n \rightarrow \infty \quad (n \rightarrow \infty).$$

For each label i , Bob sets aside a test block of size

$$m_i^X := \lfloor \varepsilon_n n_i^X \rfloor, \quad X \in \{A, E\},$$

and leaves the remaining $(n_i^X - m_i^X)$ copies untouched. The untouched part therefore has total size

$$n_{\text{data}} = n - \sum_i m_i^X = (1 - \varepsilon_n)n + o(n).$$

For each i , Bob applies to the test copies the Schumacher compression unitary U_i designed for the honest state ρ_i^A . By Lemma B 1, in the honest i.i.d. case,

$$(\rho_i^A)^{\otimes m_i^A} \xrightarrow{U_i} (|0\rangle\langle 0|)^{\otimes m_i^A(1-S(\rho_i^A))} \otimes (\text{noisy block}), \quad (\text{B1})$$

up to vanishing typicality errors. Summing over all i , the compressed test block contains

$$K_A = \sum_i m_i^A (1 - S(\rho_i^A)) = \varepsilon_n n (1 - W_{A \rightarrow B}) + o(n\varepsilon_n)$$

pure qudits and a noisy block carrying the remaining entropy.

Bob now couples the *entire* compressed test block to an ideal work extractor. Since only the pure part contributes work, the total extractable work in the honest case is

$$W_{\text{ext}}^{(A)} = K_A + o(n\varepsilon_n) = \varepsilon_n n (1 - W_{A \rightarrow B}) + o(n\varepsilon_n). \quad (\text{B2})$$

In the dishonest case, Bob applies the same honest compression unitaries U_i to the test copies $(\rho_i^E)^{\otimes m_i^E}$. The output need not have the form (B1), but the total entropy of the compressed test block is preserved because the map is unitary:

$$S(U_i(\rho_i^E)^{\otimes m_i^E} U_i^\dagger) = m_i^E S(\rho_i^E).$$

Hence, the total entropy of the full compressed dishonest test block is

$$S(\Omega_E^{\text{test}}) = \sum_i m_i^E S(\rho_i^E) = \varepsilon_n n W_{E \rightarrow B} + o(n\varepsilon_n).$$

For a block of m qudits in state ω , the maximal extractable work is $m - S(\omega)$ (in units of $k_B T \ln 2$). Therefore, the total extractable work from the dishonest test block is bounded by

$$W_{\text{ext}}^{(E)} \leq \varepsilon_n n (1 - W_{E \rightarrow B}) + o(n\varepsilon_n). \quad (\text{B3})$$

By Eqs. (B2) and (B3), the honest and dishonest asymptotic work yields are

$$W_{\text{ext}}^{(A)} = \varepsilon_n n (1 - W_{A \rightarrow B}) + o(n\varepsilon_n),$$

$$W_{\text{ext}}^{(E)} \leq \varepsilon_n n (1 - W_{E \rightarrow B}) + o(n\varepsilon_n).$$

Since $W_{A \rightarrow B} < W_{E \rightarrow B}$, the two leading-order work rates are separated by the gap

$$\Delta := W_{E \rightarrow B} - W_{A \rightarrow B} > 0.$$

Set

$$W_{\text{th}} := \varepsilon_n n \left(1 - \frac{W_{A \rightarrow B} + W_{E \rightarrow B}}{2} \right),$$

namely, the midpoint between the honest and dishonest asymptotic work values. Then,

$$W_{\text{ext}}^{(A)} - W_{\text{th}} = \frac{\Delta}{2} \varepsilon_n n + o(n\varepsilon_n),$$

whereas

$$W_{\text{ext}}^{(E)} - W_{\text{th}} \leq -\frac{\Delta}{2} \varepsilon_n n + o(n\varepsilon_n).$$

Because $n\varepsilon_n \rightarrow \infty$, the $o(n\varepsilon_n)$ corrections are negligible compared with $\varepsilon_n n \Delta$. Therefore, for all sufficiently large n ,

$$W_{\text{ext}}^{(A)} > W_{\text{th}}, \quad W_{\text{ext}}^{(E)} < W_{\text{th}}.$$

Bob can thus use the deterministic block-work test: He accepts if the extracted work from the compressed test block satisfies

$$W_{\text{ext}} \geq W_{\text{th}}. \quad (\text{B4})$$

In the honest i.i.d. case, the Schumacher asymptotics imply that Bob accepts with probability $1 - o(1)$. In the dishonest case, Eq. (B3) implies that Eve cannot reach the threshold (B4) for sufficiently large n , so the acceptance probability tends to zero.

If Bob rejects, he discards only the test block and keeps the untouched data block of size

$$n_{\text{data}} = (1 - \varepsilon_n)n + o(n).$$

Thus, the recoverable fraction obeys

$$R_n \geq 1 - \varepsilon_n - o(1) \xrightarrow{n \rightarrow \infty} 1.$$

If Bob accepts, he proceeds to the intended erasure protocol on the data block under the declared work cap. In that case, the evolution on the untested copies is asymptotically the same as in the honest Alice-assisted protocol.

Therefore, under the exclusivity condition and the above test-subset protocol, any dishonest attempt either fails authentication and leaves an asymptotically unit fraction of the memory untouched, or is operationally indistinguishable from the honest strategy on the remaining data block. ■

2. Proof of Theorem 5

Proof. In the one-sided semi-device-independent protocol, Bob chooses uniformly at random one of two dephasing settings $X \in \{R, S\}$ in each run. Over n runs, this produces two

blocks: an R -block of size n_R and an S -block of size n_S , with

$$n_R + n_S = n, \quad n_R = \frac{n}{2} + o(n), \quad n_S = \frac{n}{2} + o(n)$$

in the asymptotic limit. The relevant memory in each block is classical, namely, Bob's dephasing outcome in the chosen basis together with Alice's classical announcement.

Let

$$\begin{aligned} \tilde{W}_{A \rightarrow B} &= \frac{1}{2}(H(R|A_R) + H(S|A_S)), \\ \tilde{W}_{E \rightarrow B} &= \frac{1}{2}(H(R|E_R) + H(S|E_S)) \end{aligned}$$

denote the honest and adversarial assisted erasure costs, respectively, and assume

$$\tilde{W}_{A \rightarrow B} < \tilde{W}_{E \rightarrow B}.$$

Set

$$\Delta := \tilde{W}_{E \rightarrow B} - \tilde{W}_{A \rightarrow B} > 0.$$

Choose a sequence $\varepsilon_n \in (0, 1)$ such that

$$\varepsilon_n \rightarrow 0, \quad n\varepsilon_n \rightarrow \infty \quad (n \rightarrow \infty).$$

Bob now chooses at random a test subset of size

$$\begin{aligned} m_R &= \lfloor \varepsilon_n n_R \rfloor \quad \text{from the } R \text{ block,} \\ m_S &= \lfloor \varepsilon_n n_S \rfloor \quad \text{from the } S \text{ block,} \end{aligned}$$

and leaves the remaining runs untouched.

Since the postdephasing memory is classical, Bob can apply Shannon compression separately to the two test blocks, conditioned on Alice's honest announcements.

For the R -test block, the total entropy is asymptotically

$$m_R H(R|A_R),$$

so classical compression maps it to

$$m_R(\log d - H(R|A_R))$$

pure d -dimensional registers together with a noisy block carrying the entropy. Likewise, for the S -test block, the number of pure registers is asymptotically

$$m_S(\log d - H(S|A_S)).$$

Therefore, the total extractable work from the *combined* compressed test blocks in the honest case is

$$\begin{aligned} W_{\text{ext}}^{(A)} &= m_R(\log d - H(R|A_R)) \\ &\quad + m_S(\log d - H(S|A_S)) + o(n\varepsilon_n) \\ &= \varepsilon_n n(\log d - \tilde{W}_{A \rightarrow B}) + o(n\varepsilon_n). \end{aligned} \quad (\text{B5})$$

In the dishonest case, Bob applies the same honest compression maps to the R - and S -test blocks. These maps need

not produce the honest pure/noisy decomposition, but they are still reversible and therefore preserve the total entropy of each test block.

Hence, the total entropy of the compressed test blocks is

$$m_R H(R|E_R) + m_S H(S|E_S) = \varepsilon_n n \tilde{W}_{E \rightarrow B} + o(n\varepsilon_n).$$

The total extractable work from the full compressed test block is therefore bounded above by

$$W_{\text{ext}}^{(E)} \leq \varepsilon_n n(\log d - \tilde{W}_{E \rightarrow B}) + o(n\varepsilon_n). \quad (\text{B6})$$

By Eqs. (B5) and (B6), the honest and dishonest extractable works differ by

$$W_{\text{ext}}^{(A)} - W_{\text{ext}}^{(E)} \geq \varepsilon_n n \Delta + o(n\varepsilon_n).$$

Since $\Delta > 0$ and $n\varepsilon_n \rightarrow \infty$, for sufficiently large n the subleading corrections are smaller than $\frac{\Delta}{4} \varepsilon_n n$. Bob may therefore choose the acceptance threshold

$$W_{\text{th}} = \varepsilon_n n \left(\log d - \frac{\tilde{W}_{A \rightarrow B} + \tilde{W}_{E \rightarrow B}}{2} \right), \quad (\text{B7})$$

so that, for all sufficiently large n ,

$$W_{\text{ext}}^{(A)} > W_{\text{th}}, \quad W_{\text{ext}}^{(E)} < W_{\text{th}}.$$

Thus,

(1) in the honest case, the total extracted work from the compressed test blocks exceeds the threshold and Bob accepts;

(2) in the dishonest case, no adversarial strategy satisfying the same thermodynamic constraints can reach the threshold, so Bob rejects.

If Bob rejects, only the tested fraction

$$m_R + m_S = \varepsilon_n n + o(n)$$

has been consumed by the authentication step. The remaining

$$n - (m_R + m_S) = (1 - \varepsilon_n)n + o(n)$$

runs are untouched and therefore remain fully available for future use. Hence, the recoverable fraction obeys

$$R_n \geq 1 - \varepsilon_n - o(1) \xrightarrow{n \rightarrow \infty} 1.$$

If Bob accepts, he proceeds to the intended erasure protocol on the untested runs under the same work cap. Therefore, under the exclusivity condition, any dishonest attempt either fails authentication and leaves an asymptotically unit fraction of the memory untouched, or is operationally indistinguishable from the honest strategy on the remaining data block. This proves the theorem. ■

-
- [1] K. Maruyama, F. Nori, and V. Vedral, *Colloquium: The physics of Maxwell's demon and information*, *Rev. Mod. Phys.* **81**, 1 (2009).
 [2] R. Landauer, Irreversibility and heat generation in the computing process, *IBM J. Res. Dev.* **5**, 183 (1961).

- [3] C. H. Bennett, The thermodynamics of computation—A review, *Int. J. Theor. Phys.* **21**, 905 (1982).
 [4] M. B. Plenio and V. Vitelli, The physics of forgetting: Landauer's erasure principle and information theory, *Contemp. Phys.* **42**, 25 (2001).

- [5] P. Chattopadhyay, A. Misra, T. Pandit, and G. Paul, Landauer principle and thermodynamics of computation, *Rep. Prog. Phys.* **88**, 086001 (2025).
- [6] A. Auffèves, Quantum technologies need a quantum energy initiative, *PRX Quantum* **3**, 020101 (2022).
- [7] S. Samsi, D. Zhao, J. McDonald, B. Li, A. Michaleas, M. Jones, W. Bergeron, J. Kepner, D. Tiwari, and V. Gadepally, From words to watts: Benchmarking the energy costs of large language model inference, in *2023 IEEE High Performance Extreme Computing Conference (HPEC)* (IEEE, Piscataway, NJ, 2023), pp. 1–9.
- [8] Y. Li, M. Mughees, Y. Chen, and Y. R. Li, The unseen AI disruptions for power grids: LLM-induced transients, [arXiv:2409.11416](https://arxiv.org/abs/2409.11416).
- [9] M. Hotta, Quantum energy teleportation: An introductory review, [arXiv:1101.3954](https://arxiv.org/abs/1101.3954).
- [10] N. A. Rodríguez-Briones, H. Katiyar, E. Martín-Martínez, and R. Laflamme, Experimental activation of strong local passive states with quantum information, *Phys. Rev. Lett.* **130**, 110801 (2023).
- [11] K. Beyer, K. Luoma, and W. T. Strunz, Steering heat engines: A truly quantum Maxwell demon, *Phys. Rev. Lett.* **123**, 250606 (2019).
- [12] T. Biswas, C. Datta, and L. P. García-Pintos, Quantum thermodynamic advantage in work extraction from steerable quantum correlations, *Phys. Rev. Lett.* **135**, 110402 (2025).
- [13] W. Ji, Z. Chai, M. Wang, Y. Guo, X. Rong, F. Shi, C. Ren, Y. Wang, and J. Du, Spin quantum heat engine quantified by quantum steering, *Phys. Rev. Lett.* **128**, 090602 (2022).
- [14] C. Branciard, E. G. Cavalcanti, S. P. Walborn, V. Scarani, and H. M. Wiseman, One-sided device-independent quantum key distribution: Security, feasibility, and the connection with steering, *Phys. Rev. A* **85**, 010301(R) (2012).
- [15] D. Cavalcanti, P. Skrzypczyk, G. Aguilar, R. V. Nery, P. S. Ribeiro, and S. Walborn, Detection of entanglement in asymmetric quantum networks and multipartite quantum steering, *Nat. Commun.* **6**, 7941 (2015).
- [16] G. T. Landi and M. Paternostro, Irreversible entropy production: From classical to quantum, *Rev. Mod. Phys.* **93**, 035008 (2021).
- [17] D. Basilewitsch, R. Schmidt, D. Sugny, S. Maniscalco, and C. P. Koch, Beating the limits with initial correlations, *New J. Phys.* **19**, 113042 (2017).
- [18] M. Esposito and C. Van den Broeck, Second law and Landauer principle far from equilibrium, *Europhys. Lett.* **95**, 40004 (2011).
- [19] J. P. Peterson, R. S. Sarthour, A. M. Souza, I. S. Oliveira, J. Goold, K. Modi, D. O. Soares-Pinto, and L. C. Céleri, Experimental demonstration of information to energy conversion in a quantum system at the Landauer limit, *Proc. R. Soc. A* **472**, 20150813 (2016).
- [20] C. O. A. R. Neto and B. de L. Bernardo, Thermodynamics of ancilla-assisted erasure of quantum information: Coa Ribeiro Neto et al., *Quantum Inf. Proc.* **24**, 145 (2025).
- [21] D. Reeb and M. M. Wolf, An improved Landauer principle with finite-size corrections, *New J. Phys.* **16**, 103011 (2014).
- [22] J. M. Parrondo, J. M. Horowitz, and T. Sagawa, Thermodynamics of information, *Nat. Phys.* **11**, 131 (2015).
- [23] M. Huber, M. Perarnau-Llobet, K. V. Hovhannisyan, P. Skrzypczyk, C. Klöckl, N. Brunner, and A. Acín, Thermodynamic cost of creating correlations, *New J. Phys.* **17**, 065008 (2015).
- [24] L. d. Rio, J. Åberg, R. Renner, O. Dahlsten, and V. Vedral, The thermodynamic meaning of negative entropy, *Nature (London)* **474**, 61 (2011).
- [25] L. Henderson and V. Vedral, Classical, quantum and total correlations, *J. Phys. A: Math. Gen.* **34**, 6899 (2001).
- [26] W. K. Wootters, Entanglement of formation of an arbitrary state of two qubits, *Phys. Rev. Lett.* **80**, 2245 (1998).
- [27] M. Koashi and A. Winter, Monogamy of quantum entanglement and other correlations, *Phys. Rev. A* **69**, 022309 (2004).
- [28] H. Maassen and J. B. M. Uffink, Generalized entropic uncertainty relations, *Phys. Rev. Lett.* **60**, 1103 (1988).
- [29] M. M. Wilde, *Quantum Information Theory* (Cambridge University Press, Cambridge, 2013).
- [30] A. S. Holevo, Bounds for the quantity of information transmitted by a quantum communication channel, *Prob. Peredachi Inf.* **9**, 3 (1973).
- [31] J. Schneeloch, C. J. Broadbent, S. P. Walborn, E. G. Cavalcanti, and J. C. Howell, Einstein-Podolsky-Rosen steering inequalities from entropic uncertainty relations, *Phys. Rev. A* **87**, 062103 (2013).
- [32] P. J. Coles, M. Berta, M. Tomamichel, and S. Wehner, Entropic uncertainty relations and their applications, *Rev. Mod. Phys.* **89**, 015002 (2017).
- [33] B. Schumacher, Quantum coding, *Phys. Rev. A* **51**, 2738 (1995).
- [34] M.-C. Yang, J.-L. Li, and C.-F. Qiao, The decompositions of Werner and isotropic states, *Quantum Inf. Proc.* **20**, 255 (2021).
- [35] A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani, Device-independent security of quantum cryptography against collective attacks, *Phys. Rev. Lett.* **98**, 230501 (2007).
- [36] S. Pironio, A. Acín, S. Massar, A. Boyer de la Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, and C. Monroe, Random numbers certified by Bell's theorem, *Nature (London)* **464**, 1021 (2010).
- [37] D. Mayers and A. Yao, Self testing quantum apparatus, in *Proceedings 44th Annual IEEE Symposium on Foundations of Computer Science* (IEEE Computer Society, Los Alamitos, CA, 2003), pp. 503–513.
- [38] B. W. Reichardt, F. Unger, and U. Vazirani, Classical command of quantum systems, *Nature (London)* **496**, 456 (2013).
- [39] J. Kaniewski, Analytic and nearly optimal self-testing bounds for the Clauser-Horne-Shimony-Holt and Mermin inequalities, *Phys. Rev. Lett.* **117**, 070402 (2016).
- [40] I. Šupić and J. Bowles, Self-testing quantum states and measurements: A review, *Quantum* **4**, 337 (2020).
- [41] B. Yadin, M. Fadel, and M. Gessner, Metrological complementarity reveals the Einstein-Podolsky-Rosen paradox, *Nat. Commun.* **12**, 2410 (2021).
- [42] C.-Y. Hsieh and M. Gessner, General quantum resources provide advantages in work extraction tasks, *Phys. Rev. A* **114**, 022456 (2026).
- [43] F. Centrone and M. Gessner, Breaking local quantum speed limits with steering, *Phys. Rev. Res.* **6**, L042067 (2024).
- [44] M. Berta, M. Christandl, R. Colbeck, J. M. Renes, and R. Renner, The uncertainty principle in the presence of quantum memory, *Nat. Phys.* **6**, 659 (2010).

- [45] R. F. Werner, Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model, [Phys. Rev. A **40**, 4277 \(1989\)](#).
- [46] D. W. Lyons and S. N. Walck, Entanglement classes of symmetric Werner states, [Phys. Rev. A **84**, 042316 \(2011\)](#).
- [47] N. Li, S. Luo, and Y. Sun, Information-theoretic aspects of Werner states, [Ann. Phys. **424**, 168371 \(2021\)](#).
- [48] S. Díaz-Solórzano and E. Castro, Exact analytical solution of entanglement of formation and quantum discord for Werner state and generalized Werner-like states, [arXiv:1802.05877](#).