The work cost of answering in real time

Seth Douglas Nidhal Mghirbi

October 2026

Abstract

Many devices must answer each request before the next one arrives, while remembering what came before. If such a device is built from energy-conserving dynamics and heat baths, how much work must it spend? We answer this for every finite interactive quantum process whose outputs carry the same Hamiltonians as its inputs, at every positive finite temperature, as the optimal work per episode, one complete run of the process, over many episodes at vanishing error. Processes that do not respect time translation obey the same law when the device's whole controller need only come back approximately, with its actual mean energy and entropy refunded and its startup capital declared. The least work per episode is $k_{\mathrm{B}}T \ln 2$ times a single quantity: an entropy change plus an energy change, both measured against the same adaptive observer, the most demanding one, and minimized over all ways to build the process in the order its outputs are due. For processes that respect time translation, a device whose controller and reservoirs start thermal and whose work is deterministic and below this rate has worst-case error tending to one.

A process costs nothing exactly when it is a limit of devices that run energy-conserving steps, in the required order, on one shared heat bath. Order alone can cost work, as is known for classical information engines; we find the exact price: a two-qubit process that is free when its outputs may wait costs $k_{\mathrm{B}}T \ln 2\,[1 - h(p)]$ per episode when they may not, where p is the bias of a coin the process draws, even when everything is classical. For memoryless channels the law reduces to the real-time work law of our companion paper. In a classical example, the same accounting restricted to agents whose memory is a function of the past reproduces the online cost that Thompson et al. derive for such agents, while an agent that commits its randomness in advance pays less.

The authors developed this work with AI models from more than one developer and take responsibility for every claim. How the work was done, and what the machine checks cover, is described at the end of the paper.

Contents

1 Introduction **2**

2 Setting **5**

3 The work law **9**
3.1 Bookkeeping 9
3.2 What each episode must leave behind 9
3.3 The lower bound 10
3.4 The construction 10

4 The strong converse **11**

5 What is free **12**
5.1 Driving with Gibbs inputs . . . 12
5.2 Exact zero gain . . . 13
5.3 Small gain . . . 13
5.4 Covariance and the work rate . . . 14

6 Without time covariance **14**

7 Examples **15**
7.1 Channels . . . 15
7.2 Order costs work . . . 16
7.3 The delay process . . . 16
7.4 Delayed channels . . . 17
7.5 A correlated process . . . 17
7.6 Classical agents . . . 18

8 Discussion **19**

How this work was done **20**

A Testers, dilations and port bounds **20**

B The construction **21**

C The strong converse **24**

D Repairing small gain **26**

E The compiler without time covariance **28**
E.1 The contract and the theorem . . . 28
E.2 Finite dilations, minimax and the converse . . . 29
E.3 Recording the observer's strategy . . . 30
E.4 A joint tangent and a positive heat effect . . . 31
E.5 The hot-qubit contact . . . 32
E.6 Forming the source from Gibbs supplies . . . 32
E.7 A cyclic cache and its joint return . . . 34
E.8 Whole-history phase lift and finite completion . . . 35
E.9 Order of choices and the ledger . . . 36
E.10 Scope . . . 37

F The correlated process **38**

G Results from the companion papers **39**

1 Introduction

Landauer's principle says that erasing a bit costs at least $k_{\mathrm{B}}T\ln 2$ of work [1]. For a quantum channel applied n times, Faist, Berta and Brandão found the least work per use, the thermodynamic capacity, when the device may collect all inputs before it releases any output [2, 3]: for every channel when Gibbs-preserving maps are free, and for time-covariant channels with thermal operations. In a companion paper we showed that releasing each output before the next input arrives can cost extra work, but only for quantum channels [4].

Most devices that interact with their surroundings also remember. A controller acts on what it measured earlier, an agent responds to a history, a protocol correlates its rounds. Such a device

implements a process with memory, also called a quantum comb or a quantum strategy [5, 6]. For strategies with classical inputs and outputs, Thompson, Riechers, Garner, Elliott and Gu derived the extra work that an agent whose memory is a function of its past must dissipate when it answers online rather than in long blocks, and showed that quantum memory can reduce it [7]. Here we ask the question for every finite interactive quantum process. If a device built from energy-conserving dynamics and heat baths must release each output before the next input, how much work must it spend?

The answer is one formula. Its single ingredient, K_β, is defined in Section 2: an entropy change plus an energy change, both measured against the same adaptive observer, the most demanding one, and minimized over all ways to build the process in the order its outputs are due.

Main results. Throughout, a process has n slots. In slot j an input with Hamiltonian H_j arrives, and an output with the same Hamiltonian must be released before the next input. The process may keep memory between slots. A device runs the process again and again; one run of all n slots is an *episode*. We assume that the process commutes with one common time translation of all its inputs and outputs; without this assumption the same law holds when the device's clock need only return approximately (Theorem 9). Separate covariance at each slot is not needed: the delay process of Section 7 has only the common one.

Theorem 1 (Real-time work law for processes)**.** *Let $\mathcal{P}$ be such a process. In real time, the optimal work per episode at vanishing error is $k_{\mathrm{B}}T \ln 2\, K_\beta(\mathcal{P})$. Moreover, for every fixed $\delta > 0$, any implementation whose controller and reservoirs start in Gibbs states and that draws at most $Nk_{\mathrm{B}}T \ln 2\,(K_\beta(\mathcal{P}) - \delta)$ of deterministic work over N episodes, meaning that its battery ends in a definite energy level with probability tending to one, has error tending to one.*

As for channels, bounding only the mean work gives no such converse: a device that runs the optimal implementation with probability λ and does nothing otherwise spends λ times the rate and errs with probability about $1 - \lambda$. Nor can the Gibbs initial state of the controller simply be dropped (Section 4).

Theorem 2 (What is free)**.** *For a process $\mathcal{P}$ as above the following are equivalent:*

(i) $K_\beta(\mathcal{P}) = 0$*;*

(ii) driven by fresh Gibbs inputs, implementations of $\mathcal{P}$ can release outputs nearly in their product Gibbs state while leaving their bath nearly uncorrelated with all of them: for every $\delta > 0$ some implementation within δ of $\mathcal{P}$ has total Gibbs defect at most δ (Section 5);

(iii) $\mathcal{P}$ is a limit of causally ordered thermal implementations: one heat bath in its Gibbs state, retained throughout, and a sequence of unitaries that each conserve energy exactly, with every output released on time;

(iv) the real-time work rate of $\mathcal{P}$ is zero.

Each condition forces covariance under the common time translation, but not at each slot separately.

Milz et al. take Gibbs-preserving channels as the free channels and ask which higher-order transformations preserve them [8]. Requiring this completely, they find exactly the causally ordered combs that are thermal in their sense: inserting the Gibbs state into the first slot returns the Gibbs state, uncorrelated with the remaining comb, which is again thermal [8, Def. 8, Thm. 6]. Causal order is thus derived from equilibrium preservation. Theorem 2 starts from causally ordered processes

and uses thermal operations instead: its condition (ii) is a dilation-level, approximate form of their Gibbs-insertion condition, and in condition (iii) "thermal" means energy-conserving unitaries on one retained Gibbs bath, not Gibbs preservation. (An n-slot process here is an $(n-1)$-slot supermap in their counting.) The processes that are free in our sense are strictly fewer than theirs; this already follows from the single-channel case [9, 10, 11]. Milz et al. propose analogous constructions for more restrictive frameworks such as thermal operations, and ask whether causal order can again be derived [8, Sec. VI]. Theorem 2 gives the thermal-operation counterpart of their thermality condition for processes whose causal order is given; it does not derive causal order.

Proposition 3 (Order costs work)**.** *Let $\mathcal{P}_{p,V}$ read a bit a from a qubit A and release it, and then apply $V^{a\oplus r}$ to a second qubit Q, where r is a fresh coin equal to 1 with probability p and $V \in \{X, Z\}$; the Hamiltonians are diagonal, and for $V = X$ the qubit Q has a degenerate one. As one channel on AQ, $\mathcal{P}_{p,V}$ is a thermal operation. If a must be released before Q arrives, its real-time work rate is $k_{\mathrm{B}}T \ln 2\,[1 - h(p)]$ per episode; for $V = X$ the bound is already forced by a classical observer.*

The cheapest device keeps a masked record. Before the episode it prepares a bit M equal to 1 with probability p; at the first slot it adds a into M and releases a; at the second it applies V^M to Q. The record $M = a \oplus r$ shares $1 - h(p)$ bits with the released a, and that correlation must be erased after a has left. A device that held a until Q arrived could uncompute the record for free.

Beyond time-covariant processes, the law holds for every finite process when the controller need only return approximately (Theorem 9): on a rare branch the device records the observer's whole strategy, which tells it how much surplus energy to release as heat without touching any released output. Delayed channels and a correlated process with hidden memory are worked examples (Section 7).

Where the price comes from. A device that answers in real time may have to keep, in its memory and its baths, correlations with outputs it has already released. It cannot erase those correlations against outputs that are gone. For classical agents the same tension appears in the finding that agents that extract work efficiently must balance prediction against forgetting [12]. A classical channel needs no such correlation and pays nothing extra [4]. A quantum channel may need some, and that is the surcharge of [4]. A process with memory may have it forced on it by the task itself, classical or not, as is known for classical information ratchets, generators and agents [13, 14, 7, 12]: in Proposition 3 the device must remember a after releasing it.

Prior work. Faist, Berta and Brandão implement channels that do not respect time translation with thermal operations and a small source of coherence when the input is a fixed i.i.d. state [3, Sec. 8.2], and ask whether thermal operations can implement every i.i.d. channel universally, which they expect to require an arbitrary amount of coherence [3, Sec. 9]. Theorem 9 settles the analogous question for real-time implementation of every finite process, with, as they anticipated, an unbounded phase reference; their batch question, for a clock that must return, is answered in [15]. Lie, Son, Boes, Ng and Wilming show that a channel is a thermal operation exactly when it has a dilation, by an arbitrary unitary and environment state, that leaves the environment unchanged when the input is the Gibbs state; energy conservation, a Gibbs bath and time covariance then follow [16, Def. 1, Lemma 3, Prop. 4]. Theorem 2 is a robust form of their criterion for processes with several steps: a small defect at Gibbs inputs implies closeness to a thermal implementation, with constants that do not depend on the bath. Theorems 1 and 2 give the real-time cost of every finite process whose inputs and outputs carry the same Hamiltonians, and characterize the processes whose cost is zero. For Markovian dynamics, Luo et al. show that thermal Lindbladians are exactly those

realized by energy-conserving collision models with fresh thermal units [17]; the free processes of Theorem 2 use one bath retained across all slots. For repeated independent uses of one channel at zero Hamiltonian, the rate of Theorem 1 is the purity cost of [18], the zero set of Theorem 2 is the closure of the finitely factorizable channels [19], and a strong converse for purity is in [20]; the case of one channel at nonzero temperature is in [4].

Work extraction. Zambon and Adesso study work extraction from non-Markovian multitime processes [21]. An experimenter who stores all outputs until the end can extract more work than one who extracts step by step, by an amount bounded by the non-Markovianity of the process [21, Thm. 2]; this is the extraction-side counterpart of the gap between batch and real-time operation. Since, unlike for states and channels, the work cost of implementing a comb is not known to equal the work it can yield, they take as the cost of a comb the least work value of an ancilla state plus the asymptotic costs [2, 3] of the channels in a dilation. Theorem 1 concerns a different quantity, the cost of implementing a process in real time, episode after episode, against every observer; it does not settle their question. Fiderer, Barth, Smith and Briegel define the work capacity of a classical environment with memory, the largest rate at which an agent can extract work from it, and ask whether, for channels with memory, it can match the least work the environment must spend to implement its channel [12]. Theorem 1 gives the least cost of implementing a finite process in real time when it must be accurate against every observer; against a given agent the cost can be lower, and their question remains open. Huang et al. define the time-ordered free energy, the most work an agent without persistent quantum memory can extract online from a temporally correlated sequence of quantum states [22]; that is extraction under a causal constraint, whereas we price implementation with arbitrary quantum memory.

Classical modularity. For classical processes, the price of operating step by step rather than on the whole process at once is the modularity dissipation of Boyd, Mandal and Crutchfield, the loss of global correlations under local control [13], which Loomis and Crutchfield trace to temporal locality [14]; its simplest instance is the Landauer–Bennett price of a record whose original has left [23, 24]. Proposition 3 gives its exact rate in a process that is free when its outputs may wait.

Online agents. Thompson et al. prove their classical online cost optimal among causal agents, whose memory is a function of the past [7]. A classical agent that commits its randomness in advance lies outside that class and can pay less (Section 7). When the inputs stay on the tape, this is the analogue, for a transducer whose outputs depend on informative inputs, of the known efficiency of retrodictive pattern generators [25, 13, 14]; when each output overwrites its input, even this agent pays $1 - h(p)$ (Section 7).

Outline. Section 2 sets up the model and defines K_β. Section 3 proves the work law, Section 4 the strong converse and Section 5 Theorem 2. Section 6 drops time covariance, and Section 7 treats channels, the ordered and delay processes, and classical agents. The longer constructions are in the appendices. A companion paper [26] places this law among general schedules of arrivals, releases and loans.

2 Setting

Throughout, $\beta = 1/k_{\mathrm{B}}T$ with $0 < T < \infty$, and entropies are in bits unless stated otherwise.

Processes. A process with n slots takes an input in slot 1 and releases an output, then takes an input in slot 2, and so on; it may keep memory between slots. Slot j has an input I_j and an output O_j, both of dimension d_j and with Hamiltonian H_j. Write $d = \prod_j d_j$. Such a process is a quantum comb [5, 27, 6]: a sequence of channels, each from I_j and a memory to O_j and a new memory. Its Choi operator $R_{\mathcal{P}}$ on $\bigotimes_j I_j O_j$ satisfies the causal constraints

$$\mathrm{Tr}_{O_j} R^{(j)} = I_{I_j} \otimes R^{(j-1)}, \qquad R^{(0)} = 1, \tag{1}$$

where $R^{(j)}$ is the Choi operator of the first j slots and $R^{(n)} = R_{\mathcal{P}}$. We compare processes with the strategy distance: the largest trace distance $\frac{1}{2}\|\cdot\|_1$ that an observer can reach between them, choosing each input after seeing the earlier outputs and keeping quantum references [5, 6]. Faist, Berta and Brandão suggested such measures for implementations with memory effects [2].

Time covariance. Let $D_j(u) = e^{-\mathrm{i}uH_j}$. A process commutes with the common time translation if shifting every input by the same time u shifts every output by u. In terms of its Choi operator,

$$\Big[R_{\mathcal{P}},\ \bigotimes_j D_j(u)_{O_j} \otimes \overline{D_j(u)}_{I_j}\Big] = 0 \qquad \text{for all real } u. \tag{2}$$

This does not require each slot to be covariant on its own. The delay process releases the ground state in slot 1, releases the first input in slot 2 and discards the second input. It carries coherence from slot 1 to slot 2, and it satisfies (2) only for the common translation.

Lemma 4 (Covariant realization)**.** *A process satisfies (2) if and only if it can be realized by isometries V_j from I_j and a memory M_{j-1} to O_j and M_j that conserve energy exactly, for suitable Hamiltonians on the memories. No assumption on the rank of $R_{\mathcal{P}}$ is needed.*

For a compact symmetry group this is the covariant minimal dilation of Chiribella, D'Ariano and Perinotti [28, Eq. (15)]; we include the short proof because the time translations form a non-compact group and the memory Hamiltonians are needed explicitly.

Proof. A sequence of energy-conserving isometries gives a process that satisfies (2), since the time translations pass through each V_j. Conversely, every prefix $R^{(j)}$ inherits (2) by partial trace. Take the minimal realization [27], in which M_j is the support A_j of $\overline{R^{(j)}}$ and

$$V_j = \Big(I_{O_j} \otimes \overline{R^{(j)}}^{1/2}\, \overline{R^{(j-1)}}^{-1/2}\Big)\big(|I\rangle\!\rangle_{O_j O'_j} \otimes \mathrm{id}_{I_j \to I'_j}\big),$$

with inverses taken on supports. The prefix constraint (1) gives $V_j^* V_j = I$. By (2) the square roots commute with the time translations, the maximally entangled vector is invariant under $D_j(u) \otimes \overline{D_j(u)}$, and the wire $I_j \to I'_j$ intertwines the input translation. So V_j conserves energy when A_j carries the restriction of $\sum_{l \le j}(H_l - H_l^T)$, with H_l on the copy I'_l and $-H_l^T$ on the copy O'_l. The supports are invariant, so rank deficiency causes no difficulty. □

The device. The device model is that of [4, Sec. 2]. A device consists of a controller C, a battery B and a reservoir R. The controller holds all hardware: memory, clocks, catalysts and any work registers not charged to the battery. The reservoir is a finite collection of cells, each with its own Hamiltonian, all declared in advance and all starting in their Gibbs states at inverse temperature β, independent of each other and of everything else. Every operation is a unitary that conserves the total energy of the systems it acts on. Measurements and records are part of the dilation and are counted. There is no external coherence and no Gibbs-preserving map supplied as a black box.

Real-time operation. The device runs N episodes of $\mathcal{P}$, one after another. Within an episode, input $j+1$ is supplied only after output j has left, and the device never acts again on a delivered output. Between slots, and after the last one, it may act on its internal systems as it likes. The inputs come from an observer who may choose them adaptively, across slots and across episodes, and keep quantum references. The *error* ε_N is the strategy distance between the whole interaction and N ideal episodes of $\mathcal{P}$.

Work. The controller must return. For the constructions of Theorems 1 and 2 it returns exactly to its initial state, independent of all other systems (Theorem 9 relaxes this); for the lower bounds it is enough that its free energy returns up to $o(N)$. The battery is mechanical: work is its mean loss of energy, and its entropy changes by $o(N)$. It is *stationary* if it starts in a state that commutes with its Hamiltonian, as an energy eigenstate does; a stationary battery holds no phase reference. Work is *deterministic* when the battery starts in an energy eigenstate, ends, with probability tending to one for every observer, in a fixed energy eigenstate, and its mean energy differs from that eigenvalue by $o(N)$. The last clause matters only for a controller that returns approximately, as in Sections 6 and 7: it prevents a broad reference from lending the device energy that the battery repays in one rare, large event. In the strong converse (Theorem 8) the controller also starts in a thermal state. For a family of implementations with $\varepsilon_N \to 0$, the *work rate* is $\limsup_N W_N/(Nk_{\mathrm{B}}T\ln 2)$, in bits per episode, and the optimal real-time rate is the infimum over such families. These rates allow reservoirs, controllers and running times of any finite size. Table 1 lists the conditions that each main result uses.

Result	Controller starts	Controller returns	Battery	Error and conclusion
Rate law and free set (Theorems 1 and 2)	in any declared state	in the constructions exactly, independent of everything else; for the lower bound, its free energy up to $o(N)$	mechanical, entropy change $o(N)$	error tending to zero; optimal work K_β per episode, zero exactly for the free processes
Strong converse (Theorem 8)	in a Gibbs state, independent of the reservoirs	not required	starts in an energy eigenstate; deterministic work below $N(K_\beta - \delta)$	error tending to one, exponentially fast if the battery never misses its target
Without covariance (Theorem 9)	in any declared state; its free energy is invested capital	approximately: jointly within $\eta_N \to 0$, with actual mean energy and entropy changes $o(N)$	ends in a fixed eigenstate except with vanishing probability; entropy and mean absolute deviation $o(N)$	error $O(N^{-2})$; optimal work K_β per episode; no converse is claimed

Table 1: The contracts of the main results. Every error is the strategy distance of the whole interaction, over all episodes, against observers that adapt and keep quantum references. When the controller returns only approximately, its energy and entropy refunds are separate requirements, since closeness in trace distance alone does not bound them. Startup capital is declared and is not part of the work. Throughout, the number N of episodes is known in advance, and each output carries the same Hamiltonian as its input.

Remark 1 (Why covariance). With the controller returned exactly and a stationary battery, a device can only implement processes that satisfy (2); for channels this is the familiar fact that thermal

operations are time-covariant [3, Sec. 3]. The argument behind the returning-clock proposition of [4] applies unchanged: the device as a whole commutes with time evolution, so when the inputs, the references and the battery are invariant, so is the final state of everything that is kept. Feeding energy-diagonal inputs in all other episodes isolates one episode, whose implemented process then satisfies (2); its strategy distance to $\mathcal{P}$ is at least half the largest distance between $\mathcal{P}$ and its time translates. Common covariance is therefore necessary for vanishing error, and by Lemma 4 it is all that is needed. The battery must be stationary: a battery in a broad superposition of energy levels is a phase reference of its own, which the device can spend while the controller returns exactly [4]. With a controller that returns only approximately, covariance is not needed (Theorem 9).

Ports. We take each input and its output to have the same Hamiltonian; isospectral ports, identified by a fixed unitary, are the same. If $H_{I_j} \neq H_{O_j}$, a device must change the Hamiltonian of the system it hands back, and the answer depends on how that change is modelled. For instance, a qubit with energies 0 and E that is replaced by a maximally mixed qubit with zero Hamiltonian yields work if the device may switch the Hamiltonian itself, but not if it must emit a fresh qubit from its own supplies. We do not treat this case.

The cost. An *ε-extension* of $\mathcal{P}$ is a process $\widetilde{\mathcal{P}}$ with one extra output F, released before the first input and never acted on afterwards, such that discarding F leaves a process within strategy distance ε of $\mathcal{P}$, and such that $\widetilde{\mathcal{P}}$ itself satisfies all the causal constraints (1). A *tester* is an observer for one episode: it starts in a pure state, prepares the inputs adaptively, keeps quantum references, and ends with a register Z that holds everything it kept, including all outputs.

Definition 1 (Real-time cost of a process)**.** For a tester t run against an extension $\widetilde{\mathcal{P}}$, let $S(Z|F)_t$ be the conditional entropy of its final register given F, and let

$$e_\beta(t) = \frac{\beta}{\ln 2} \sum_j \operatorname{Tr} H_j (\rho_{O_j} - \rho_{I_j})$$

be the energy the process adds, in bits, where ρ_{I_j} and ρ_{O_j} are the states of the j-th input and output in this run. Set

$$K_{\beta,\varepsilon}(\mathcal{P}) = \inf_{\widetilde{\mathcal{P}}} \sup_t \left[S(Z|F)_t + e_\beta(t) \right], \qquad K_\beta(\mathcal{P}) = \lim_{\varepsilon \downarrow 0} K_{\beta,\varepsilon}(\mathcal{P}),$$

with the infimum over ε-extensions and the supremum over testers.

F is an inert mathematical purifier of the device's initial state, not a message available to the tester. Formally the extension emits it first, so that the extension is itself a process, and never touches it again; none of it carries information about the coming inputs. The tester's operations act only on its own memory and the released outputs of $\mathcal{P}$, never on F; conditional entropy is used only to compute the device's entropy increase, and it does not authorize either party to operate on F. If B is the device's retained body, the purified experiment ends in a pure state on BZF, so $S(Z|F) = S(B_{\text{final}}) - S(B_{\text{initial}})$: after the episode, $S(Z|F)$ equals the entropy the episode has added to the device (Lemma 7). If the tester could read F, the value would change; in the example of Section 7 with hidden memory, knowing b would let the observer choose late inputs branch by branch. So K_β is the least entropy, plus energy, that one episode must add to a device that does not know its inputs in advance.

Both terms are evaluated on the same tester, and this matters. For the delay process with $H = \operatorname{diag}(0, E)$, the entropy term alone is largest for a maximally mixed second input and the

energy term alone is largest for a second input in its ground state; taking the two maxima separately would give 1, while $K_\beta = \log_2(1 + e^{-\beta E})$ (Section 7).

Proposition 5 (Basic properties). *(a) For one slot, $\mathcal{P}$ is a channel Φ and $K_\beta(\mathcal{P})$ is its real-time cost $\kappa_\beta(\Phi)$ from [4].*

(b) For n uses of a time-covariant channel Φ, viewed as one process, $K_\beta = n\,\kappa_\beta(\Phi)$.

(c) $K_\beta(\mathcal{P}) \geq 0$, and $|K_\beta(\mathcal{P})| \leq 2\log_2 d + \frac{\beta}{\ln 2}\sum_j 2\|H_j\|$.

(d) At zero Hamiltonian, K_β involves entropy alone: it is the purity cost of answering in real time.

Proof. (a) A one-slot extension is a state σ_{QOF} in which F is independent of the reference Q of the input, so Definition 1 becomes the definition of [4] once two smoothings are compared. Smoothing input by input and smoothing in strategy distance have the same limit as $\varepsilon \to 0$, and the limit commutes with the supremum over inputs by the minimax argument of Lemma 6. (b) For the lower bound, apply the chain rule to $S(Z|F)$ across the n uses: each use adds at least the one-use cost on its own input. For the upper bound, take a product of one-use extensions; each use's entropy and energy then depend only on that use's average input. (c) Feed every slot a Gibbs input. Every extension is itself a finite process, so it has an exact dilation with F as the purifier of the bath. Against this tester, $S(Z|F) + e_\beta$ is the sum over slots of the relative entropies $D(\rho_{O_j B_j} \| \gamma_j \otimes \rho_{B_j}) \geq 0$, where B_j is the bath after slot j (Section 5); here the equal input and output Hamiltonians are used. So every extension has value at least 0. The bound on $|K_\beta|$ holds because a pure tester's final register has effective dimension at most d^2 (Appendix A). (d) is Definition 1 with $e_\beta = 0$. □

3 The work law

The lower bound comes from the bookkeeping identity of [4] and a bound on the entropy each episode must add to the device. The upper bound is a construction. Section 4 upgrades the lower bound to every fixed error.

3.1 Bookkeeping

Write $D = CBR$ for everything except the inputs and outputs, and let $A = S(D_{\text{fin}}) - S(D_{\text{init}})$ be the entropy it gains. Let ΔE_{sys} be the total energy delivered in the outputs minus the total energy taken from the inputs, and let w be the free energy drawn from the battery and the controller, in units of $k_{\text{B}}T \ln 2$. The accounting identity of [4] holds verbatim:

$$w = A + \tfrac{\beta}{\ln 2}\,\Delta E_{\text{sys}} + D(R_{\text{fin}}\|\gamma_R) + I(C{:}B{:}R)_{\text{fin}} \;\geq\; A + \tfrac{\beta}{\ln 2}\,\Delta E_{\text{sys}}. \tag{3}$$

It concerns only the device, so correlations with the observer do not enter; a reservoir cell that is set aside stays in R, so discarding never helps.

3.2 What each episode must leave behind

Lemma 6 (Minimax). $K_{\beta,\varepsilon}(\mathcal{P}) = \sup_t \inf_{\widetilde{\mathcal{P}}} [S(Z|F)_t + e_\beta(t)]$.

Proof. For a fixed extension, $S(Z|F)_t + e_\beta(t)$ is concave in the tester: the bath state is affine in t, entropy is concave, and the payload energy is affine. Mixtures of extensions are realized by adding a classical flag to F, whose entropy cancels in $S(Z|F)$, and stay ε-close to $\mathcal{P}$. Given $\eta > 0$, finitely many extensions $\widetilde{\mathcal{P}}_1, \ldots, \widetilde{\mathcal{P}}_L$ have $\min_l [\cdot]_l < b + \eta$ at every tester, where b is the right-hand side. The

set of points $x \in \mathbb{R}^L$ dominated by $([\cdot]_1(t), \ldots, [\cdot]_L(t))$ for some t is closed, convex and downward closed, and excludes $(b+\eta, \ldots, b+\eta)$. A separating hyperplane with non-negative normal gives flag probabilities whose mixture has value below $b+\eta$ at every tester. □

Lemma 7 (Causal accumulation)**.** *Suppose the error is at most ε, and the observer runs a fixed tester t in every episode, with a fresh copy each time and all earlier registers left alone. Then each episode raises $S(D) + \frac{\beta}{\ln 2} E_{\text{payload}}$ by at least $\inf_{\widetilde{\mathcal{P}}}[S(Z|F)_t + e_\beta(t)]$, the infimum over ε-extensions. Hence*

$$A + \tfrac{\beta}{\ln 2}\Delta E_{\text{sys}} \;\geq\; N \inf_{\widetilde{\mathcal{P}}}\big[S(Z|F)_t + e_\beta(t)\big].$$

Proof. Let D_k be the device before episode k and let F_k purify it; F_k contains the earlier references and outputs. The fresh tester starts pure and independent of $D_k F_k$. The episode acts on the tester's inputs and on D_k, and F_k is never touched, so the episode, seen from the observer with F_k released first, is an extension of the process the device actually implements in episode k. That process is within ε of $\mathcal{P}$: an observer could run this tester in episode k and the ideal comparison elsewhere. After the episode the global state is pure, so $S(D_{k+1}) - S(D_k) = S(Z_k F_k) - S(F_k) = S(Z_k|F_k)$, and the payload energy changes by exactly $e_\beta(t)$ in bits. Operations between episodes are unitary on D and do not change $S(D)$, so the increments add up. □

The device's memory may be correlated with every earlier output; the lemma allows that. What it rules out is a memory that already knows the coming inputs.

3.3 The lower bound

Choose, by Lemma 6, a tester t whose value $\inf_{\widetilde{\mathcal{P}}}[S(Z|F)_t + e_\beta(t)]$ is within η of $K_{\beta,\varepsilon_N}(\mathcal{P})$. Lemma 7 and (3) give

$$w \;\geq\; N\big[K_{\beta,\varepsilon_N}(\mathcal{P}) - \eta\big] \tag{4}$$

for every $\eta > 0$. For a mechanical battery the actual work is $k_{\mathrm{B}} T \ln 2\,(w - \Delta S_B)$ when the controller returns, and $\Delta S_B = o(N)$. Letting $\varepsilon_N \to 0$ gives the lower half of Theorem 1. Covariance is not used: the lower bound holds for every process and every controller.

3.4 The construction

The matching construction has four layers (Appendix B). At each N it compiles one finite implementation of an approximate extension with error N^{-3}, so that the error accumulated over N episodes still vanishes; every overhead depends only on d, the Hamiltonians, β and N, never on the size or spectrum of the extension's bath.

A covariant dilation. A finite extension is realized by one bath, in a state τ, and one unitary contact per slot. Averaging over a finite grid of phases of the torus generated by the energy gaps, controlled by a flag in the bath, makes the realization covariant under the common time translation without increasing its cost.

What an observer can do to the bath. Every tester is dominated by the flat tester, which feeds maximally mixed inputs: the bath state it leaves satisfies $\omega_t \leq d^2 \omega_0$. A pure tester's final register has effective dimension at most d^2, however large its references. These bounds involve only the ports.

Paying for entropy and energy together. A single sharp energy bank, shared by all episodes, records the energy that flows between the payload and the bath, without measuring the payload's energy. An analytic interpolation over a strip in the complex plane, applied to the whole episode at

once, turns the port bounds into an operator inequality that holds even when the bank, the observer and the earlier baths are entangled. It charges entropy and energy on the same tester. Tracing out the bank averages the whole history over one common phase, never over separate phases at separate times.

Paying the work. Converters between subspaces of given Gibbs weight, built from reservoirs with exactly conserving swaps, form the baths at the start and reset them after the last output, moving the battery by the logarithm of the Gibbs-volume ratios. The bank's large starting energy is borrowed and returned, and the battery's entropy and energy variance tend to zero. A final swap with a thermal cell returns the controller exactly. The net debit is $Nk_{\mathrm{B}}T\ln 2\,(K_\beta(\mathcal{P})+o(1))$.

This proves the upper half of Theorem 1. The same bank also handles processes whose Choi operator is not of full rank: it pays for changes of relative likelihood, so the target never has to be implemented exactly.

4 The strong converse

The lower bound (4) is a rate statement. The following theorem covers every error below one, provided the work is deterministic. For batch implementations of channels, the converse of Faist, Berta and Brandão already holds at fixed small error [3, Prop. 4.2]; Theorem 8 is a real-time counterpart for processes that covers every error below one. For a battery projector P_B that commutes with H_B, write

$$b_\beta(P_B)=\frac{\beta E_{\mathrm{init}}}{\ln 2}+\log_2\mathrm{Tr}\big(P_B e^{-\beta H_B}\big).$$

Theorem 8 (Strong converse)**.** *Let $\mathcal{P}$ satisfy (2) and let $\delta>0$. There are constants $a,b,c>0$, depending only on $\mathcal{P}$, the Hamiltonians, β and δ, with the following property. Consider any real-time implementation whose controller and reservoirs start in independent Gibbs states, of any size, and whose battery starts in a pure energy eigenstate of energy E_{init}. Let P_B satisfy $b_\beta(P_B)\le N(K_\beta(\mathcal{P})-\delta)$, and let q_N be the probability that the battery ends outside P_B when the observer of the proof is used. Then*

$$1-\varepsilon_N\ \le\ q_N+c\,e^{-bN}+2^{-aN}.$$

For a battery that ends in a single eigenstate of energy E_{fin}, the condition on P_B says exactly that the work $E_{\mathrm{init}}-E_{\mathrm{fin}}$ is at most $Nk_{\mathrm{B}}T\ln 2\,(K_\beta(\mathcal{P})-\delta)$. So with deterministic work below the rate, the error tends to one, exponentially fast if the battery never misses its target. This is the second statement of Theorem 1. A projector of rank R_B whose energies are all at least $E_{\mathrm{init}}-W$ has $b_\beta(P_B)\le\beta W/\ln 2+\log_2 R_B$: the accepted degeneracy counts as work. Without this, a large degenerate battery in a pure state could supply purity at no energy cost. The hypothesis that the controller starts in a Gibbs state cannot simply be dropped. A controller that only returns approximately, with its actual energy and entropy changing by $o(N)$, can carry an incoherent energy bank whose loans depend on the branch and cancel in the mean; with it, a covariant reset runs at half the vanishing-error rate with error exactly $1/2$ [15], while the battery's debit is the same on every branch. What decides this is how spread out the controller's capital is, not its size: spread $o(N/\sqrt{\eta})$, where η is the controller's return error, still forces a strong converse for every channel and without time covariance, but only at the thermodynamic capacity, which can lie below the real-time rate, and a bank of spread of order $N/\sqrt{\eta}$ can break even that [15]; whether a converse at the real-time rate survives such a controller is open.

Outline of the proof. The proof follows the fixed-error converse for channels in [4], with one new ingredient: an observer that runs a whole episode adaptively and still conserves energy exactly. Details are in Appendix C.

Step 1: an observer that conserves energy. Before N or the device is fixed, choose a tester whose value $\inf_{\widetilde{\mathcal{P}}}[S(Z|F)_t + e_\beta(t)]$ is within a small margin of $K_\beta(\mathcal{P})$, that has full rank, and that commutes with the common time translation; averaging over the translations and mixing in a little of the flat tester achieves this, since the value is concave and invariant. Apply the construction of Lemma 4 to the tester's own causal constraints. This realizes the tester by isometries that conserve energy when each stage of its memory carries the matching charge. Its final register Z, of dimension at most d^2, then carries the charge $H_Z = \sum_j (H_{O_j} - H_{I_j}^T)$, and $\operatorname{Tr} H_Z \nu_Z$ is exactly the energy delivered minus the energy injected in the episode.

Step 2: a tradeoff valid for every device. For this tester, the least $S(Z|F)$ over all extensions is a convex function of the observer's final state. A separating hyperplane gives an affine lower bound, which may be averaged over the charge so that it depends only on the energy-block pinching of Z. The observer measures a random fraction γ of the episodes with an informationally complete measurement inside each energy block, built from rank-one effects and dilated so that it conserves charge. This turns the affine bound into a min-tradeoff function for entropy accumulation [29]. It holds whatever the device's state, and its constants do not involve the device.

Step 3: accumulation and an energy event. Run this tester afresh in every episode, while the device may keep arbitrary correlations between episodes. Entropy accumulation bounds the smooth min-entropy of the observer's registers, given a purification E of the initial hardware, from below by $Nh - C\sqrt{N}$ on the event G that the empirical score is high. In the ideal experiment the observer's total charge is at least $N(\Delta E - u)$, with $\Delta E = \operatorname{Tr} H_Z \nu_Z$ and any fixed $u > 0$, except with exponentially small probability.

Step 4: a Gibbs-weighted cap. Every operation conserves the total charge, and the hardware starts thermal. So the same smooth min-entropy, after G, the energy event and P_B have all passed with probability p, is at most $b_\beta(P_B) - \beta N(\Delta E - u)/\ln 2 + N\theta_\gamma + \log_2(1/p) + 4$, where $\theta_\gamma \to 0$ as $\gamma \to 0$. The partition function of the hardware cancels exactly, whatever its size; with equal input and output Hamiltonians, swapping a carrier in or out does not change it. Comparing the two bounds with $b_\beta(P_B) \le N(K_\beta - \delta)$ forces $p < 2^{-aN}$, and the union bound $p \ge 1 - \varepsilon_N - q_N - ce^{-bN}$ gives the theorem. □

Covariance is used only in Step 1, to choose the tester and to give it an energy-conserving realization; per-slot covariance is never needed. The constants are not uniform as $\delta \to 0$, as for channels.

5 What is free

A *dilation* of a process is a finite realization of it: a bath B in a state τ, and for each slot a unitary U_j on I_jB (equivalently O_jB, since the ports are square). Its *gain* is $g = \sup_t[S(Z|F)_t + e_\beta(t)]$, with F the purifier of τ. By Definition 1, $K_\beta(\mathcal{P})$ is the limit, as $\varepsilon \to 0$, of the least gain of dilations within ε of $\mathcal{P}$.

5.1 Driving with Gibbs inputs

Feed every slot a fresh Gibbs state $\gamma_j = e^{-\beta H_j}/Z_j$, and let τ_j be the bath after slot j, with $\tau_0 = \tau$. Entropy conservation and $\log_2 \gamma_j = -\beta H_j/\ln 2 - \log_2 Z_j$ give, for each slot,

$$a_j := S(\tau_j) - S(\tau_{j-1}) + \tfrac{\beta}{\ln 2} \operatorname{Tr} H_j(\rho_{O_j} - \gamma_j) = D(U_j(\gamma_j \otimes \tau_{j-1})U_j^* \,\|\, \gamma_j \otimes \tau_j) \;\ge\; 0. \tag{5}$$

The Gibbs tester is a legal tester, and its score is $\sum_j a_j \le g$. This proves $K_\beta \ge 0$ (Proposition 5).

The slot-by-slot defects add up to a single relative entropy. Let registers $A_1, \dots, A_n$ start in $\gamma_1, \dots, \gamma_n$, and let slot j act on A_j and the bath, so that A_j plays the role of the j-th input and then holds the j-th output. Let Ω_j be the joint state of the bath and all A_l after slot j, and $\xi_j = \tau_j \otimes \gamma_{A_1} \otimes \cdots \otimes \gamma_{A_n}$. Then

$$D(\Omega_j \| \xi_j) = \sum_{l \le j} a_l \qquad \text{for every } j. \tag{6}$$

Indeed $S(\Omega_j)$ equals the initial entropy, $\log \xi_j$ is a sum of the bath's $\log \tau_j$ and Gibbs terms, unused registers cancel, and each used one contributes its energy change. Equation (6) includes every correlation with earlier outputs. It says that the gain against Gibbs inputs is small exactly when the released outputs stay close to Gibbs and the bath stays nearly uncorrelated with all of them. This is condition (ii) of Theorem 2, and it shows that a single tester, the one that feeds Gibbs states, decides whether a process is free. The exact criterion of Lie et al. for one channel also uses the Gibbs input alone [16], and so does the thermality condition of Milz et al. at the Gibbs-preserving level, which inserts only Gibbs states [8, Def. 8].

5.2 Exact zero gain

Write G_j for the inverse of the first j contacts, acting on the bath and $A_1, \dots, A_j$, and let $\sigma = \tau \otimes \gamma_{A_1} \otimes \cdots \otimes \gamma_{A_n}$. Regard the registers A_l as part of an enlarged bath with source σ, and gauge the contacts by

$$W_j = (I \otimes G_j)(U_j \otimes I_A)(I \otimes G_{j-1}^*). \tag{7}$$

The internal gauges cancel between consecutive contacts and the last one acts on the bath alone, so the whole adaptive process, with every deadline, is unchanged. The point of the gauge is that one source now serves every contact. By (6) and the inequality $\|\sqrt{\rho} - \sqrt{\sigma}\|_2^2 \le D_{\mathrm{nat}}(\rho\|\sigma)$ [4],

$$\big\| [W_j, \sqrt{\gamma_j \otimes \sigma}] \big\|_2 \;\le\; 3\sqrt{\ln 2\, g} \qquad \text{for every } j. \tag{8}$$

If $g = 0$, every W_j preserves $\gamma_j \otimes \sigma$. On the support of σ, which every W_j leaves invariant, all W_j then commute with $H_j + H_B$ for the bath Hamiltonian $H_B = -\beta^{-1} \ln \sigma$: this is an exact causally ordered thermal realization of the process. Conversely, a causally ordered thermal realization has $S(Z|F)_t + e_\beta(t) = -D(\omega_t \| \sigma) \le 0$ for every tester, where ω_t is the final bath. So zero gain is attained exactly by thermal realizations. For one channel this is the argument of Lie et al. [16, Prop. 4]; the gauge (7) lets one source serve every contact.

5.3 Small gain

The difficulty is that a small gain may be achieved only with baths of unbounded size and arbitrarily small eigenvalues. Repairing each contact separately, against its own bath state, gives a different bath Hamiltonian at each slot, and the repairs do not compose. The common source (8) removes this obstacle. The repair has four steps (Appendix D).

(i) Mix a small multiple of the identity into the source, and approximate each port Hamiltonian by one whose gaps are multiples of a mesh h_E.

(ii) Round the logarithms of the source's eigenvalues to the mesh βh_E, with *one* random offset for all contacts. Energy blocks of the rounded bath Hamiltonian then match bins of the source's logarithms, and with high probability over the offset each W_j has little weight outside the blocks. Separately good offsets would not give one compatible bath Hamiltonian.

(iii) Pinch each W_j into the blocks and complete it to a unitary by its polar decomposition. Each completed contact conserves energy exactly for the rounded bath.

(iv) Compare the original and repaired processes by replacing one contact at a time, along purified histories that keep all earlier outputs and references. Each replacement costs little, because the bath before contact j is close to the source in the square-root distance, and the errors add up over the n contacts.

Irrational gaps are handled as in [4]: they are approximated by commensurate ones with the same integer labels, and one Gibbs lattice register, shared by all contacts and retained throughout, absorbs the difference so that the original Hamiltonians are conserved exactly. Tracing out the lattice averages the whole history over one common phase. The result is a causally ordered thermal process within

$$\varepsilon + C\, g^{1/(6\max\{1,r\})} \tag{9}$$

of $\mathcal{P}$, where ε is the distance of the dilation to $\mathcal{P}$, r is the number of rationally independent gaps of all the port Hamiltonians together, and C depends only on n, the dimensions, the Hamiltonians and β, possibly exponentially in n. It does not depend on the bath.

5.4 Covariance and the work rate

Each repaired process commutes with the common time translation of its approximating Hamiltonians, and these converge to the original ones. So every process with $K_\beta = 0$ commutes with the common time translation. For one channel and an exact equilibrating dilation, this is part 2 of [16, Prop. 4]; what is new is that only the common translation is forced, and that the bound does not depend on the bath. It need not commute with the translations of each slot separately: the thermal delay process, which swaps the first input into a bath qubit and swaps it out at the second slot, is a causally ordered thermal process with this property.

The operational statement (iv) follows from the work law in one direction: zero rate implies $K_\beta = 0$ by the lower bound (4), which does not use covariance. In the other direction, take a thermal realization within $o(1/N)$ of $\mathcal{P}$, prepare N independent Gibbs copies of its bath in advance, and retire one after each episode. No bath is ever reset, so no work is drawn; a controller clock that indexes the episodes is formed and returned at a cost of $O(\log N)$, which is sublinear. This completes the proof of Theorem 2.

Examples. Werner–Holevo, used n times, preserves equilibrium yet has $K_\beta = n \log_2 3$; it is not free. The ordered process of Proposition 3 and the delay process have positive cost, so they are not free either, although the first is a thermal operation as one channel. The thermal delay process, and any sequence of thermal operations sharing one bath, are free.

6 Without time covariance

By Remark 1, a controller that returns exactly, with a stationary battery, can only implement processes that commute with the common time translation. As for single channels [4], the restriction disappears if the controller need only return approximately: its final joint state with everything else close to its initial state tensored with the rest, and its mean energy and entropy changing by $o(N)$, uniformly over observers.

Theorem 9 (Real-time work law without covariance)**.** *Let $\mathcal{P}$ be a fixed finite process whose slots return the same physical carrier with the same Hamiltonian, not necessarily commuting with any time translation. Consider implementations of N episodes, with N known in advance, by a controller in any declared initial state, whose free energy is counted as invested capital, independent Gibbs stocks and energy-conserving unitaries. Require, uniformly over all adaptive observers, including adaptation across episodes and entanglement with references, that the complete error tend to zero, that the controller return jointly with everything else within $\eta_N \to 0$, that its actual mean energy and entropy change by $o(N)$, and that the battery end in a fixed energy eigenstate except with vanishing probability, with entropy and mean absolute deviation $\mathbb{E}|H_{W,\mathrm{f}} - E_{\mathrm{nom}}|$ both $o(N)$. Then the optimal work per episode is $k_{\mathrm{B}}T \ln 2\, K_\beta(\mathcal{P})$; at each N it is attained by a single implementation that serves every observer, with error $O(N^{-2})$.*

The lower bound is that of Section 3, which uses no covariance; it needs only the weaker signed condition that the battery's mean energy end within $o(N)$ of its nominal endpoint, while the construction meets the stronger absolute one. The construction, given in full in Appendix E, adds two ingredients to Appendix B. First, on a rare branch, of probability δ, the device keeps each input before the last and emits half of a maximally entangled pair in its place. What it holds before the last output is then exactly the observer's tester operator, up to normalization. Second, a joint tangent, built as in Appendix B but for this enlarged dilation, bounds the energy $J(t)$ that one episode adds to the payload and the bath by a constant w for every tester, with w plus the cost of forming the source at most the gain. The slack $w - J(t)$ is linear in the tester and nonnegative on the tester cone. By conic duality it equals $\mathrm{Tr}(D_{\mathrm{hist}}t)$ for a positive operator D_{hist} whose norm depends only on the ports. The retained record holds the tester only with probability δ, so the correction it must apply has norm $O(1/\delta)$; a heat contact of that norm moves exactly this slack, on average, into a thermal qubit of large gap, so that the phase reference's mean energy is refunded up to a remainder that vanishes as the gap grows. No released output is touched. A broad phase reference that approximates the whole history replaces the sharp energy bank of Appendix B; its width, the hot gap and the cache size depend on the chosen extension. The sources are prepared before the first input from a cache in the style of the correlated catalysts of Duan, Feng, Li and Ying and of Shiraishi and Sagawa [30, 31].

As for channels, the controller's energy comes back only on average: on typical branches the reference can end higher or lower by order N; on rare branches the reference and the cache move by far more, which balances the mean. Only the mean energy and the joint entropy are refunded. The theorem is a rate law: it bounds neither the capital nor the tails of the heat, refunds the controller's energy only on average, uses no efficient or local hardware, and assumes carriers whose dimension and Hamiltonian do not change. The clock, the cache and the heat baths must be very large. The delayed channels and the correlated process of Section 7 are explicit instances with closed-form values.

7 Examples

7.1 Channels

By Proposition 5, a process with one slot is a channel and K_β is its real-time cost κ_β, and n uses of a time-covariant channel cost $n\kappa_\beta$. So Theorem 1 contains the real-time work law for channels of [4], and Theorem 2 contains its characterization of the free channels. What is new for processes is that the task itself can force the device to keep correlations with what it has released.

7.2 Order costs work

Proof of Proposition 3. Free as one channel. The map $(a, q) \mapsto (a, V^{a\oplus r}q)$, with A dephased, is a controlled-V gate followed by a V applied with probability p. Both act with two Gibbs coin qubits and no battery, and both conserve energy: controlled-Z is diagonal for any diagonal Hamiltonians, and controlled-X needs Q to be degenerate.

Lower bound. In every episode feed A a uniformly random bit a, keeping a copy, and feed Q half of a fresh maximally entangled pair (for $V = Z$) or a fresh uniform bit (for $V = X$). Let e_1 be the probability of reading a wrongly from the first output and e_2 the probability that the value of $a \oplus r$ read from the second output and the kept reference differs from a. Three facts hold in each episode, whatever the device's correlations with earlier ones. Releasing the first output lowers the device's entropy by at most $h(e_1)$, on average over a. After the release the device must still hold a record of a: the second slot, together with the fresh probe, is a channel acting on the device alone, so Holevo's bound and Fano's inequality give $\chi(a{:}D) \geq 1 - h(e_2)$. And the second slot cannot lower the device's entropy, because its input is locally maximally mixed. Together,

$$S(D_{\text{after}}) - S(D_{\text{before}}) \;\geq\; 1 - h(e_1) - h(e_2),$$

with $e_1 \leq \varepsilon$ and $e_2 \leq p + \varepsilon$. Summing over episodes and using (3), with the energy change bounded by ε times the spans of the Hamiltonians, gives the rate $1 - h(p)$ as $\varepsilon \to 0$. For $V = X$ every input and test is classical.

Construction. Prepare, before the episode, a bit M that equals 1 with probability p. At the first slot, add a into M and release a; at the second, apply V^M to Q. Then $M = a \oplus r$ is a record of a masked by the coin, and its correlation with the released a is $1 - h(p)$ bits. Forming N such biased bits from Gibbs supplies costs $N[1 - h(p)] + o(N)$ when they are formed collectively; forming them one episode at a time would cost more. The used bits are reset after the last episode. If a could be withheld until Q arrives, the device could apply the whole map at once, and the rate would be 0. □

The mechanism is the Landauer–Bennett price of a record whose original has already left [23, 13]. What the proposition adds is the exact rate in a process that is free as a channel, against a coherent adaptive observer, with a converse that does not depend on the size of the device. The advance-committed bit M is the point of the construction: it stores a only up to the coin, which is all the second slot needs.

7.3 The delay process

Let both slots carry a qubit with $H = \operatorname{diag}(0, E)$. The delay process releases the ground state in slot 1, releases the first input in slot 2, and discards the second input. It commutes with the common time translation but not with the translation of each slot separately. A ground-state memory qubit and two swaps implement it: swap the first input into the memory and release the memory's ground state; then swap the second input into the memory and release the first input. For every tester this costs

$$S(\rho_{I_2}) - \tfrac{\beta E}{\ln 2}\,\langle 1|\rho_{I_2}|1\rangle,$$

where ρ_{I_2} is the second input: the device absorbs that input and must later reset it. The maximum over ρ_{I_2} is attained at the Gibbs state and equals $\log_2(1 + e^{-\beta E})$. Conversely, a tester that sends the ground state first and a purified Gibbs state second forces this value on every extension, since both outputs must then be close to pure. So

$$K_\beta = \log_2(1 + e^{-\beta E}), \qquad \text{work per episode} = k_{\mathrm{B}} T \ln(1 + e^{-\beta E}).$$

This is the free energy needed to prepare a ground state from a thermal one [32, 33], the cost Zambon and Adesso use in their example of work investment [21]. Here it is the exact real-time cost of a process that moves coherence from one time to another.

7.4 Delayed channels

The delay process generalizes to any channel. Let Φ be a channel on the port, and let $\mathcal{P}_\Phi$ release the ground state in slot 1, release Φ applied to the first input in slot 2, and discard the second input; take the ground energy to be 0. Then

$$K_\beta(\mathcal{P}_\Phi) = \kappa_\beta(\Phi) + \log_2 Z_H,$$

where κ_β is the real-time cost of [4]. For the lower bound, feed a purified input ρ in slot 1 and, independently, a purified Gibbs state in slot 2. The causal constraints make the second reference independent of everything that exists before slot 2, so it can join the side output F; and the first output is nearly pure. What remains is a one-slot extension of Φ, whose score is at least the channel cost at ρ, while the second input contributes $\log_2 Z_H$. For the upper bound, swap the first input into a ground-state cell, apply a dilation of Φ before the second slot, and swap the second input in. By subadditivity the score is at most the channel score at the first input plus $S(\rho_{I_2}) - \frac{\beta}{\ln 2}\operatorname{Tr} H\rho_{I_2} \le \log_2 Z_H$, for every tester. The two terms add, because correlating the inputs can only lower this sum.

If Φ is time-covariant, $\mathcal{P}_\Phi$ commutes with the common time translation, and Theorem 1 gives the work $k_\mathrm{B}T\ln 2\,\kappa_\beta(\Phi) + k_\mathrm{B}T\ln Z_H$ per episode, with the strong converse. If Φ is not time-covariant, no device whose controller returns exactly and whose battery is stationary can run $\mathcal{P}_\Phi$ (Remark 1). With a controller that returns approximately, the same work is optimal. The lower bound of Section 3 does not use covariance, and it is attained by a device that runs Φ on the stored input with the returning clock of [4] and resets the absorbed second inputs together at the end. For the Hadamard gate on a qubit with $H = E|1\rangle\langle 1|$, this is $E/\sqrt{2} + k_\mathrm{B}T\ln(1+e^{-\beta E})$ per episode. To our knowledge, these are the first exact real-time costs of processes with memory that do not commute with time translation. If slot 1 releases the Gibbs state instead of the ground state, the cost is $\kappa_\beta(\Phi)$ alone, and the process is free exactly when Φ is a limit of thermal operations; the thermal delay process of Section 5 is the case $\Phi = \mathrm{id}$, and the process of Fig. 2(a) of Zambon and Adesso, a thermal qubit swapped in, flipped and swapped out, is the case $\Phi = X(\cdot)X$ [21].

7.5 A correlated process

A process can also correlate its slots through hidden memory. At the start of each episode let the process draw a hidden uniformly random bit b. In slot 1 it applies U_b, with $U_0 = I$ and $U_1 = e^{\mathrm{i}\vartheta Z}$; in slot 2 it applies V_b, with $V_0 = X$ and V_1 the Hadamard gate; both slots carry $H = E|1\rangle\langle 1|$, and $0 < \vartheta < \pi/8$. The process does not commute with time translation, and its memory, the bit b, ties the two slots together. Its cost is $K_\beta = \beta v_\vartheta/\ln 2$ with

$$v_\vartheta = \frac{E}{2}\sqrt{\frac{3}{2} + \cos 2\vartheta + \sin 2\vartheta}.$$

The calculation has two independent restrictions: the first output limits how well the observer can distinguish the two values of b, while causality limits how much of b can be hidden in F. The first fixes the energy maximum and the second rules out a negative entropy credit. The pairs (ρ_0, ρ_1) of second-slot inputs an observer can produce for the two values of b are exactly those with root fidelity $\|\sqrt{\rho_0}\sqrt{\rho_1}\|_1 \ge \cos\vartheta$. Maximizing the energy the second slot must supply over these pairs gives v_ϑ.

The entropy term drops out: because the second-slot gates of the two branches differ by more than a phase, the causal constraint at slot 2 allows F only a classical correlation with b, so $S(Z|F) \geq 0$, with equality for a classical copy of b. Appendix F gives the three steps. For $\pi/8 \leq \vartheta \leq 7\pi/8$ the fidelity constraint no longer binds and the cost is $(E/2)(1 + 1/\sqrt{2})$; the process depends on ϑ only modulo π.

A device that runs each branch with a returning clock pays each branch's own worst case, $(E/2)(1 + 1/\sqrt{2}) > v_\vartheta$: the observer cannot produce both worst cases at once, and a device that ignores this overpays. Paying exactly v_ϑ needs an energy correction that depends on the observer's whole strategy, through the first input, the first output and the second input together. A device can obtain it without touching any released output. On a rare branch, of probability $\delta = N^{-3}$, it keeps the first input and emits half of a maximally entangled pair instead; the state it keeps is then exactly the observer's strategy operator. A heat effect of norm at most $8E/\delta$ on that state, routed to a thermal qubit of gap of order EN^8, makes the average energy exactly v_ϑ per episode. With a clock that returns approximately, the error is $O(N^{-2})$, the battery reaches its target on every branch, and the work is $Nv_\vartheta + O(E/N^2)$. As for channels, the controller's energy comes back only on average, and the clock's capital is large, of order EN^{20}. It is an explicit instance of Theorem 9, with a closed-form value.

7.6 Classical agents

A classical agent that answers online is a process whose slots measure and prepare in an energy eigenbasis. Thompson et al. compare such an agent with batch execution, writing outputs onto a tape that keeps the inputs [7]. Their minimal online cost is proved for agents whose memory is a function of the observed past [7]: unifilar, or history-synchronized, agents. Our law places no such restriction.

Take the classical case of Proposition 3: $Y_1 = A$, then $Y_2 = B \oplus A \oplus R$ with a fresh coin R equal to 1 with probability $p < 1/2$. A classical agent whose memory is a function of the past must keep the causal state, which here is A itself. An unrestricted agent may instead keep $M = A \oplus R$, committing its randomness in advance; M is not a function of the observed past: it is randomness sampled before the episode, and the identity $M = Y_2 \oplus B$ holds only afterwards, giving the device no advance knowledge of B or Y_2. Under two accounting conventions, the exact costs per episode, in units of $k_{\mathrm{B}}T \ln 2$ and with h the binary entropy, are as follows.

Interface	Unrestricted memory	Classical memory that is a function of the past
In place (each output replaces its input)	$1 - h(p)$	1
Copy-out (inputs kept, outputs on fresh mixed cells)	$2 - h(p)$	2

The values in the left column are lower bounds against arbitrary quantum memory and correlated episodes; those in the right column are lower bounds for classical memory that is a function of the past. All four are attained by classical constructions with all supplies and resets counted. In the copy-out interface the agent writes Y_1 on a fresh blank and, at the second slot, releases its memory register itself as the cell that carries Y_2, after adding B (and, for the agent whose memory is a function of the past, a fresh coin), keeping the incoming mixed cells; keeping the memory instead would cost one bit more. One-time-pad phase coins, which cost nothing, dephase the kept inputs

against coherent observers. On their tape, the copy-out row, Thompson et al. show that quantum memory that is a function of the past can lower the right-column entry [7, Result 2]; we do not compare our agent with their quantum agents. In both interfaces the classical agent whose memory is a function of the past pays $h(p)$ more. In the copy-out interface, which matches Thompson et al.'s tape, the unrestricted cost $2 - h(p)$ equals the batch cost, so the excess $h(p)$ over batch of the agent whose memory is a function of the past is exactly their online cost, Eq. (2) of [7], evaluated episode by episode on this strategy with uniform independent inputs; an agent that commits its randomness in advance pays nothing for answering online. The in-place interface is that of the information ratchets of Boyd, Mandal and Crutchfield, whose modularity dissipation for arbitrary hidden states is [13, Eq. (14)]; there the batch cost is 0, and $1 - h(p)$ is the online price.

This does not contradict the theorem of Thompson et al., who assume throughout that agents are causal in their sense, that is, history-synchronized, with memory a function of the observed past. One sentence of their appendix, in the arXiv version, states that ε-transducers are minimal among all classical agents [7]; this example shows that the minimality holds only within that class. Both agents obey the same deadlines and use no future input. Boyd, Mandal and Crutchfield found the same kind of gap for information ratchets: pattern generators whose hidden states are retrodictive avoid the modularity dissipation of predictive ones [13], and retrodictive generators also minimize transient costs [25]; Loomis and Crutchfield proved that a classical generator is efficient exactly when it is retrodictive [14]. The advance-committed memory $M = Y_2 \oplus B$ is retrodictive in the same sense, for a transducer whose outputs depend on its inputs. In the copy-out interface it removes the online cost entirely; in the in-place interface, which is that of the ratchets, it still dissipates $1 - h(p)$, so for a transducer with informative inputs retrodiction alone does not give efficiency. Two caveats apply to the comparison. Thompson et al. price a given input process, while K_β is a worst case over observers; the numbers above coincide for uniform independent inputs. And at $p = 1/2$ the predictive states merge and the gap disappears.

8 Discussion

A prediction. Theorem 8 and the remark after Theorem 1 make a sharp, testable distinction. Run the delay process N times on a device whose controller and reservoirs start in independent Gibbs states, with deterministic work below $k_{\mathrm{B}}T \ln(1 + e^{-\beta E})$ per episode: its strategy error tends to one, so an adaptive observer can tell it from the ideal process with probability tending to one. With the same *average* work, a device that runs the optimal implementation with probability λ and does nothing otherwise keeps its error at most about $1 - \lambda$. Deterministic and average work therefore carry different guarantees: a zero–one law for the first and a smooth law for the second.

Clocks. Our constructions use a single energy bank shared by all episodes, and tracing it out averages the whole history over one common phase. A device that kept a separate energy record at each slot would average each slot over its own phase, and could not implement the delay process at vanishing error. More generally, a device whose clock returns exactly, and whose battery holds no phase reference, cannot implement a process that fails (2) at all (Remark 1). For single channels, a clock that need only return approximately runs every channel at the real-time rate, though the clock and the baths must be large [4]. The same holds for every finite process with memory (Theorem 9); the rare record of the observer's strategy is what makes this possible.

Open problems. Our constructions use reservoirs of astronomical size; efficient constructions, and the rate at which the error can vanish, are open. Ports whose input and output Hamiltonians

differ need a model of how the device changes the Hamiltonian, and then the rate can be negative. The constructions assume that the number of episodes is known in advance. For processes that do not respect time translation we have a rate law only: a strong converse, and versions with bounded clock capital or a controller that returns its energy branch by branch, are open. Finally, the cost of processes at exactly zero error is open; for extreme channels it is at least the real-time rate, even in batch [4].

How this work was done

This paper is the work of its two authors together with AI models from more than one developer. The authors set the direction of the research throughout: they chose the questions, decided which lines to pursue and which to drop, judged and selected the results, and take full responsibility for every claim. The models did most of the exploration, derivation, computation and drafting.

Much of the exploration took place in long working sessions, in which many model instances worked on a problem in parallel, communicated with one another, and built on and challenged each other's results. We also set the models against one another. Before a result was kept, an instance that had not seen its proof, usually from a different developer, first tried to refute the statement and then audited the proof line by line, with exact computation wherever possible. Each paper then went through several rounds of adversarial review, on its own and together with its companion papers, and AI-assisted literature searches were used to check priority and how we report the work of others. Errors found at any of these stages are corrected in the text. The proofs are given in full: the checks support them but do not replace them.

Machine checks accompany the paper. Programs in its repository, https://github.com/Apsiape/answering-in-real-time, check the numbers, worked examples and finite constructions listed in its verification notes, with exact arithmetic where possible, and test the stated inequalities on many random and near-extremal instances. The general theorems, and the construction and continuity arguments the notes do not list, are proved in the text; results quoted from the companion papers are checked in their own repositories. The finite and algebraic cores of the arguments are proved in Lean with Mathlib, and every formal proof is audited to use only the standard axioms. Here the formal proofs cover the Gibbs variational step and the value of the delay process, the order-costs-work device and the classical-agents table, the bound on the correlated process, the whole-history phase lift and its energy conservation, and the arithmetic of the strong converse, the repair of small gain and the finite-source factory. The verification notes in the repository list which statements are checked by machine and which are proved only in the text. Reports of the independent checks of this version are available from the authors.

A Testers, dilations and port bounds

Testers. A deterministic tester for n slots is described by a positive operator $t = I_{O_n} \otimes C_n$, whose dual prefixes C_k on $I_1 O_1 \cdots O_{k-1} I_k$ satisfy $\mathrm{Tr}_{I_k} C_k = I_{O_{k-1}} \otimes C_{k-1}$ for $k \geq 2$ and $\mathrm{Tr}\, C_1 = 1$ [5, 6]. Put $x = t^T$. Run against an extension with Choi prefixes R_j^F, the tester's purified final state is

$$\nu_{ZF} = (\sqrt{x} \otimes I_F)\, R_n^F\, (\sqrt{x} \otimes I_F) \tag{10}$$

on $Z = \mathrm{supp}(x)$, which has dimension at most d^2. Any further purification or waste of the tester can be compressed isometrically into this Z; it may not be discarded, since that would change $S(Z|F)$. The tester normalizations form a compact convex set.

Dilations. A dilation is a bath B in a state τ and unitaries U_j on I_jB; nothing is assumed about energy. Purify τ by F and run a purified tester. Then ZBF is pure at the end, so

$$S(Z|F) = S(\omega_t) - S(\tau),$$

where ω_t is the final bath. Conversely, every finite extension has a causal dilation that leaves F untouched. Allocating every later ancilla as initial padding in one bath, and completing each isometry to a unitary on I_jB (possible because the ports are square), turns it into a dilation. So extensions and dilations describe the same objects, with no bound on the size of the bath.

Port bounds. Since $\operatorname{Tr} t = d$, we have $x \le d\, I$. For the extended Choi operator R_n^B the final bath on tester t is $\omega_t = \operatorname{Tr}_Z[(x \otimes I)R_n^B]$, and the flat tester $t_0 = I/d$ gives $\omega_0 = \operatorname{Tr}_Z R_n^B/d$. Positivity gives

$$\omega_t \;\le\; d^2\,\omega_0 \qquad \text{for every tester } t. \tag{11}$$

For normalized Choi operators $J = R/d$, the strategy distance satisfies

$$d_{\mathrm{c}}(R, R') \;\le\; d^2 \cdot \tfrac{1}{2}\|J - J'\|_1, \tag{12}$$

since $\sqrt{x}(R - R')\sqrt{x}$ has trace norm at most $d\|R - R'\|_1$. These bounds involve only the ports. Every score $S(Z|F)_t + e_\beta(t)$ is bounded in absolute value by $2\log_2 d + \frac{\beta}{\ln 2}\sum_j \operatorname{span}(H_j)$, and the scores are equicontinuous in t uniformly in the bath, because conditional entropy continuity uses only $\dim Z \le d^2$. This compactness is what the minimax of Lemma 6 uses.

B The construction

Fix N. All constants C below depend only on d, the Hamiltonians and β. Let $\omega_1, \dots, \omega_r$ be a basis of the group generated by all port energies, and write each energy as $a_e \cdot \omega$ with an integer vector a_e.

B.1 One finite dilation for each number of episodes. Choose a dilation with $d_{\mathrm{c}}(\Psi, \mathcal{P}) \le N^{-3}$ and gain at most $K_\beta + N^{-1/8}$. Mix a weight $\rho = N^{-12}$ of the maximally mixed state into its source and round the source's logarithmic energies down to a mesh $h = N^{-12}$; call the result p. Because a tester's final register has dimension at most d^2, conditional entropy continuity changes every score by $O(\sqrt{\rho} + \sqrt{h})$, uniformly in the bath, and energies change by $O(\rho + h)$. Averaging over a finite grid of common phases, controlled by a flag in the bath, makes the dilation covariant without increasing its gain (Section 3.4).

B.2 The joint tangent. Let t_* maximize $S(\omega_t) + e_\beta(t)$, put $\zeta = \omega_{t_*}$, and let v be the maximum. The support of ζ contains that of every ω_t: otherwise mixing in outside weight would raise the entropy at first order faster than the energy changes. The directional derivative gives

$$\operatorname{Tr}\omega_t(-\log_2 \zeta) + e_\beta(t) \;\le\; v \qquad \text{for every tester } t.$$

When ζ is singular this is read on its support, by regularizing and letting the regularization vanish. Put $\zeta_s = (1 - s)\zeta + s\,\omega_0$ and round its logarithmic energies to the mesh h, normalizing to a faithful state q. Then $e^{-h}\zeta_s \le q \le e^h\zeta_s$, so by (11) $\omega_t \le (e^h d^2/s)\, q$, and

$$\operatorname{Tr}\omega_t(-\log_2 q) - S(p) + e_\beta(t) \;\le\; g + c_s + h/\ln 2, \qquad c_s = \log_2 \tfrac{1}{1-s}. \tag{13}$$

The relative-entropy mismatch between ω_t and q is controlled by this single tangent, for entropy and energy together.

B.3 A strip estimate for a whole episode. For complex z replace each contact by $U_j(z) = e^{z\beta H_j/2} U_j e^{-z\beta H_j/2}$, leaving the observer's isometries untouched, and call the resulting map of the whole episode W_z. For $x = \operatorname{Re} z$ the imaginary filters are unitary and the real ones have singular values bounded by the port energies, so $e^{-C|x|} I \le W_z^* W_z \le e^{C|x|} I$. Positivity and (13) give $W_z p W_z^* \le A_x (I_Z \otimes q)$ with $A_x = e^{h+C|x|} d^4/s$; no dimension of the bath enters. Define

$$M(z) = \operatorname{Tr}\big[(I_Z \otimes q)^{-z}\, W_z\, p^{1+z}\, W_{\bar z}^*\big].$$

On the line $x = 1$, $\operatorname{Tr}[(I \otimes q)^{-1} W_z p^2 W_z^*] \le e^{h+3C} d^4/s$; on $x = -1$, $\operatorname{Tr}[(I \otimes q) W_z W_z^*] \le e^C d^2$. By Cauchy–Schwarz these bound $|M|$ on both boundary lines; M is bounded on the strip for each finite dilation, and the three-lines theorem, with an $e^{\varepsilon z^2}$ barrier, gives $|M(z)| \le K_s := C e^h/s$ for $|x| \le 1$. Cauchy's formula then gives, for real $|t| \le 1/4$,

$$\begin{aligned} M(t) &\le \exp\big[t M'(0) + C K_s t^2\big], \\ M'(0) = \ln 2 \big[\operatorname{Tr} \omega_t(-\log_2 q) - S(p) + e_\beta(t)\big] &\le a := \ln 2\,(g + c_s) + h, \end{aligned}$$

by (13). Here and in B.4, t is a real parameter, while the subscript of ω_t and the argument of $e_\beta(t)$ refer to the fixed tester. Each energy derivative is the actual energy before or after a contact; the observer's gates do not enter. A bound $M'(0) \ge -a_-$ with a fixed a_- controls negative tails.

B.4 One bank for all episodes. Let R be the observer's memory, of any size, and fix its whole protocol for one episode. For a common phase θ of the ports define the positive operator on R

$$X_t(\theta) = \operatorname{Tr}_B\big[p^{1+t}\, W_t(\theta)^*\, (I_Z \otimes q^{-t})\, W_t(\theta)\big].$$

Every pure state of R defines a tester, so B.3 gives the operator bound $X_t(\theta) \le e^{ta + CK_s t^2} I_R$ for $0 \le t \le 1/4$, uniformly in θ, the protocol and the size of R.

The bank has coordinates $(b_0, b) \in \mathbb{Z} \times \mathbb{Z}^r$ and energy $E_{\mathrm{cap}} + (h/\beta) b_0 + b \cdot \omega$; it starts in the sharp state $(0, 0)$ with a large counted capital E_{cap}. Each contact is lifted by translating b by the charge it moves and b_0 by the controller's mesh energy; the lift conserves energy exactly. Let $Q_t = e^{-t\beta(H_{\mathrm{bank}} - E_{\mathrm{cap}})/2}$ and let δ_l be the episode's change of controller offsets and log-partition functions. In the Fourier representation of the bank,

$$\mathbb{E}_p\big[W_{\mathrm{lift}}^* (I \otimes Q_t^2) W_{\mathrm{lift}}\big] = e^{t\delta_l}\, Q_t\, \mathrm{Mult}_\theta\big[X_t(\theta)\big]\, Q_t, \tag{14}$$

where Mult_θ is multiplication by the θ-dependent operator on the payload torus. The identity can be checked by expanding all bath labels and charge histories: initial and final labels agree in bra and ket, no source label is copied, and no payload energy is measured. The pointwise bound on $X_t(\theta)$ makes (14) at most $e^{t\delta_l + ta + CK_s t^2} Q_t^2$ as an operator on R, the bank and any spectator, even when they are entangled. Only the current source needs to be fresh. Iterating over the N episodes, with $Y = \beta(E_{\mathrm{cap}} - H_{\mathrm{bank}})$ and $\Delta = \sum_l \delta_l$,

$$\mathbb{E}\, e^{tY} \le \exp\big[t\Delta + N(ta + CK_s t^2)\big].$$

With $s = \gamma = N^{-1/8}$, optimizing t gives $\Pr\{Y > \Delta + N(a + \gamma \ln 2)\} \le e^{-cN^{5/8}} =: q_{\mathrm{cut}}$, and the positive and negative bounds give $\mathbb{E}(|Y| + CN)^2 \mathbf{1}_{\mathrm{bad}} \le CN^2 \sqrt{q_{\mathrm{cut}}}$: truncated second moments that do not depend on the bank's span.

B.5 A fixed Hamiltonian and a finite bank. The controller holds all N bath cells and a clock with $nN+1$ states. In clock block k, used cells carry the Hamiltonian of q, unused cells that of p, and the active cell the appropriate one at its episode boundaries and zero at intermediate contacts; with offsets $c_k = h\lfloor \ln Z_k/h \rfloor$ this is one fixed block-diagonal Hamiltonian. The Gibbs state restricted to clock 0 is $p^{\otimes N}$ and to the last clock value $q^{\otimes N}$, and $|\Delta| \le h$. The bank lives in the finite box $|b_0| \le 2D_E$, $\|b\|_\infty \le 4(A+1)(nN+1)$, where D_E is the span of the controller's mesh indices and $A = \max \|a_e\|_\infty$. Lifted columns are specified on the inner box, which every actual history stays in; their images are orthonormal within each energy shell, and the remainder of each shell is completed arbitrarily. Tracing out the bank averages the whole episode, all n contacts together, over one common phase, which reproduces Ψ exactly because it is covariant; the N cells give N episodes of Ψ in every adaptive experiment.

B.6 Forming the sources. An energy-conserving converter between subspaces P and Q of Gibbs volumes Z_P, Z_Q moves the battery by any signed work below $\beta^{-1}\ln(Z_Q/Z_P)$. It is built from a reservoir of tiles with energies $j\cdot\omega'$ and integer degeneracies $\lfloor K_0 e^{\beta j\cdot\omega'}\rfloor$, by isometric injection and exact swaps within each energy shell, succeeds with an input-independent probability $1-\delta$, and leaves every other column untouched. Use it to form the controller's clock-0 block from its Gibbs state, at work W_f, with tile rounding θ_f and failure δ_f at most N^{-40}, and twirl each occupied target shell with a finite one-design. The formed controller then satisfies

$$\rho_{\text{good}} \;\le\; \kappa\, p^{\otimes N}, \qquad \kappa = \big[(1-\theta_f)^2(1-\delta_f)\big]^{-1} = 1 + O(N^{-40}).$$

This multiplicative domination transfers every positive statistic of the ideal fresh-source comparison to the actual, correlated source, with no additive error multiplied by the bank's span. On formation failure the device applies the identity, which conserves energy because each input and its output have the same Hamiltonian.

B.7 Reset, rate and endpoints. After the last episode, convert the subspace P_X (last clock value, bank accepted at $Y \le d_N$, with d_N the mesh rounding of $N(a+\gamma\ln 2)+h$) to Q_X (bank at $(0,0)$). Summing the geometric tails, $Z_{P_X} \le Z_M e^{-\beta E_{\text{cap}}} e^{d_N} L_N/(1-e^{-h})$, where $L_N = O(N^r)$ counts the payload coordinates, and $Z_{Q_X} = Z_C e^{-\beta E_{\text{cap}}}$. The capital and the bath spectrum cancel. Finally swap the controller with a fresh Gibbs copy, so it returns exactly and independently on every branch. The nominal debit satisfies

$$\frac{\beta W_{\text{nom}}}{N\ln 2} \;\le\; g + c_s + \gamma + O\Big(\frac{\log N}{N} + h\Big) \;\le\; K_\beta(\mathcal{P}) + o(1).$$

The battery fails to reach its endpoint with probability at most $\delta_f + \kappa q_{\text{cut}} + \delta_r$, and on every failed branch its energy deviates by at most $C(|Y|+N)$, never by the capital. So its absolute energy variance tends to zero. Its entropy also tends to zero: the bank's coordinate second moments vanish, and an integer-valued variable with second moment μ has entropy at most $\lambda\mu + \ln\sum_b e^{-\lambda b^2}$, which tends to zero with $\lambda = \frac{1}{2}\ln(1/\mu)$. The observer's error is at most

$$N\cdot N^{-3} + N(\rho + e^h - 1) + 2(\kappa-1) + \delta_f \;\to\; 0.$$

Every bound is uniform in the witness chosen at each N, whose hardware may depend on N in a non-effective way. Rank deficiency is handled directly: the bank pays for changes of relative likelihood, so no exact implementation of the target is ever needed. At zero Hamiltonian the bank has only the coordinate b_0.

C The strong converse

Throughout, t is a tester, $x = t^T$, and $h_\varepsilon(t)$ is the infimum of $S(Z|F)_t + e_\beta(t)$ over ε-extensions, so that $K_\beta = \sup_t \lim_{\varepsilon\to 0} h_\varepsilon(t)$ by Lemma 6.

C.1 An observer that conserves energy. For a process satisfying (2), $h_\varepsilon(t)$ is invariant under a common time translation of t and concave in t. Averaging t over a finite grid of the torus generated by the gaps therefore does not lower its limit, and mixing in a small weight of the flat tester makes it full rank at an arbitrarily small loss, since all scores are bounded. Fix, before N or the device is given, a full-rank invariant tester with $\lim_\varepsilon h_\varepsilon(t) \geq K_\beta - \eta$. Then x commutes with the charge

$$H_Z = \sum_j (H_{O_j} - H_{I_j}^T).$$

Realize t causally. Its dual prefixes C_k are Choi operators of channels from the outputs $O_1, \dots, O_{k-1}$ to the inputs $I_1, \dots, I_k$, and the minimal dilation $V_k = (I \otimes \sqrt{C_k^T})(|I\rangle\!\rangle \otimes I)$ has memory $M_k = \operatorname{supp}(C_k^T)$. Common invariance of C_k, inherited from C_n, makes V_k intertwine $\sum_{l\leq k} H_{I_l} + H_{M_k}$ with $\sum_{l<k} H_{O_l}$, where H_{M_k} is the compression of $\sum_{l<k} H_{O_l} - \sum_{l\leq k} H_{I_l}^T$. By uniqueness of minimal dilations there are stage isometries $F_k : M_{k-1} \otimes O_{k-1} \to I_k \otimes M_k$ with $V_k = (I \otimes F_k)(V_{k-1} \otimes I)$. These intertwine the stage energies on the whole memory: the two factorizations transformed by a common phase agree on the range of $V_{k-1} \otimes I$, and that range spans $M_{k-1} \otimes O_{k-1}$ by minimality. No later input is used to release an earlier output. Keeping O_n and M_n as Z gives $\nu_Z = \sqrt{x} R_{\mathcal{P}} \sqrt{x}$ on a space of dimension at most d^2 with physical Hamiltonian H_Z, and telescoping the intertwiners gives

$$\operatorname{Tr} H_Z \nu_Z = \sum_j (\text{energy delivered in } O_j - \text{energy injected in } I_j).$$

The sign matters: for one slot the reference carries $-H_I^T$, and the physical charge is $H_O - H_I^T$. All observer waste is kept in Z.

C.2 A tradeoff valid for every device. Let $\nu_* = \sqrt{x} R_{\mathcal{P}} \sqrt{x}$. Because x has full rank, $\nu \mapsto$ the associated process is a bijection on the affine space of deterministic processes. For ν in its image let $g_x(\nu)$ be the infimum of $S(Z|F)$ over finite extensions of the associated process, with all prefix constraints, and $g_x = +\infty$ elsewhere. Flagged mixtures make g_x convex, and $|g_x| \leq 2 \log_2 d$. Its lower semicontinuous hull at ν_* is the $\varepsilon \to 0$ cost for this tester, by (12). A separating hyperplane gives, for every $\zeta > 0$, an affine $\ell \leq g_x$ with

$$\ell(\nu_*) + \tfrac{\beta}{\ln 2} \operatorname{Tr} H_Z \nu_* \;\geq\; K_\beta - \eta - \zeta.$$

This holds even when ν_* is rank-deficient, and needs no optimal extension of bounded size. Average ℓ over the charge H_Z: the domain and the constraints are invariant and ν_* commutes with H_Z, so ℓ may be taken to depend only on the energy-block pinching of ν.

In each energy block of dimension m, the vectors e_i, $e_i + e_k$ and $e_i + \mathrm{i} e_k$ $(i < k)$, normalized by their frame operator, give m^2 rank-one effects that are informationally complete on the block. Over all blocks this gives $L \leq d^4$ effects $\Lambda_x = |v_x\rangle\langle v_x|$ from whose probabilities the pinching is recovered linearly. Fix $0 < \gamma < 1/4$. In each episode the observer applies to Z the instrument with Kraus operators $\sqrt{1-\gamma}\, I$ and $\sqrt{\gamma}\, |u\rangle\langle v_x|$, where u has the same energy as v_x; it has a charge-conserving dilation that writes two copies of the outcome at zero energy. One copy stays with Z, forming the register A; the other is counted with the device's records below. For any state of the device, purified

by an inert F, the fresh tester is independent of F and the induced extension satisfies all prefix constraints, so $S(Z|F) \geq g_x(\nu) \geq \ell(\nu)$. Given a tested outcome, Z is pure, so $S(A|F) \geq (1-\gamma)\ell(\nu)$. The outcome distribution, $q_x = \gamma \operatorname{Tr} \Lambda_x \nu$ and $q_0 = 1-\gamma$, determines the pinching, so

$$f(q) = (1-\gamma)\,\ell\big(\text{pinching recovered from } q_x/\gamma\big)$$

is an affine min-tradeoff function that is valid for every state of the device, with no dimension of the device in its coefficients.

C.3 Accumulation and an energy event. Run this tester afresh in every episode; the device may keep arbitrary correlations between episodes, and the rounds of entropy accumulation are whole episodes. Let E purify all initial thermal hardware. The Markov conditions of [29] hold trivially, since there is no side register passed between rounds. Let q_* be the ideal outcome distribution, fix $\zeta_s > 0$, put $h = f(q_*) - \zeta_s$, and let G be the event that the empirical score is at least h and at most $2\gamma N$ episodes are tested. Then $|G| \leq 2^{N\theta_\gamma}$ with $\theta_\gamma = h(2\gamma) + 2\gamma \log_2 L$, and by [29, Thm. 4.4], for $p_G = \Pr(G) > 0$,

$$H^{\delta}_{\min}(A^N|E)_{|G} \;>\; Nh - C_f\sqrt{N\big[1 - 2\log_2(\delta p_G)\big]}, \tag{15}$$

with C_f depending only on d, L and f. Let $\Delta E = \operatorname{Tr} H_Z \nu_*$. The instrument conserves charge, so in the ideal experiment the observer's total charge H_A is a sum of N independent bounded terms with mean $N\Delta E$, and the projector P_A onto $H_A \geq N(\Delta E - u)$ rejects with probability $e_N \leq ce^{-bN}$ for any fixed $u > 0$. If p is the probability that G, P_A and P_B all pass in the real experiment, then $p \geq 1 - \varepsilon_N - e_N - q_N$.

C.4 A Gibbs-weighted cap.

Lemma 10. *Let ω_{AEK} be pure with $(H_A + H_K - H_E)|\omega\rangle = E_{\mathrm{i}}|\omega\rangle$, suppose $H_A \geq L_E$ on the support, $\omega_E \leq e^{-\beta H_E}/(pZ_E)$, and K is supported in a subspace that reduces H_K and has Gibbs volume Z_K. Then for $0 < s < 1$,*

$$H^{s}_{\min}(A|E)_{\omega} \;\leq\; \log_2 \frac{e^{\beta(E_{\mathrm{i}} - L_E)} Z_K}{pZ_E} + 2\log_2 \frac{2}{1-s}.$$

Proof. The charge condition makes ω_{AE} commute with $H_A - H_E$, and its block paired with K-energy k has rank at most the dimension d_k of the K-support at that energy. Put $a_0 = (1-s)/(2Z_K)$ and keep the eigenvectors with eigenvalue at least $a_0 e^{-\beta k}$; the rejected weight is at most $a_0 \sum_k d_k e^{-\beta k} = (1-s)/2$. For the accepted projector Π, $\omega_{AE} \geq a_0 e^{\beta(H_A - H_E - E_{\mathrm{i}})}\Pi$, and conjugating with $e^{-\beta H_A/2}$ gives $\operatorname{Tr}_A \Pi \leq e^{\beta(E_{\mathrm{i}} - L_E)}/(a_0 p Z_E)\, I_E$. A state within purified distance s has accepted weight at least $(1-s)/2$; if it is at most $2^{-h_0} I_A \otimes \zeta_E$, this weight is at most 2^{-h_0} times the bound on $\operatorname{Tr}_A \Pi$. □

Apply this after the three projections. The initial thermofield state has charge zero with $H_E = H_T^T$, the battery has energy E_{init}, and the observer and instrument conserve charge, so the charge shell is exact. After conditioning, $\omega_E \leq \gamma_E/p$. The controller, all reservoirs, the battery and the second outcome copy are in K, so $Z_E = Z_T$ and $Z_K = Z_T \operatorname{Tr}(P_B e^{-\beta H_B})|G|$: the hardware's partition function cancels, whatever its size. With equal input and output Hamiltonians, swapping a carrier in or out does not change it. For $s = 1/2$ and $L_E = N(\Delta E - u)$,

$$H^{1/2}_{\min}(A^N|E)_{\mathrm{proj}} \;\leq\; b_\beta(P_B) - \tfrac{\beta N(\Delta E - u)}{\ln 2} + N\theta_\gamma + \log_2(1/p) + 4. \tag{16}$$

C.5 Assembly. If a pure state is projected by $P_A \otimes P_K$ with probability $r > 0$ and $\delta \le r/64$, then $H_{\min}^{1/2}(A|E)$ of the projected state is at least $H_{\min}^{\delta}(A|E) + \log_2 r$: lift a near-optimal smoothing state by Uhlmann's theorem, filter it, and renormalize; its squared purified distance to the projected state is at most $4\delta/r \le 1/16$. Use this with $\delta = p/64$ after G, where the conditional pass probability is $r = p/p_G \ge p$. Combining with (15) and (16),

$$b_\beta(P_B) \;>\; N\Big[h + \tfrac{\beta(\Delta E - u)}{\ln 2} - \theta_\gamma\Big] - C_f\sqrt{N\big[13 + 4\log_2(1/p)\big]} - 2\log_2(1/p) - 4.$$

Given δ in Theorem 8, choose $\eta, \zeta, \gamma, \zeta_s, u$ before N so that the bracket is at least $K_\beta - \delta/2$; entropy and energy here belong to the same tester. Put $a = \min\{\delta/16, \delta^2/(256C_f^2)\}$. If $b_\beta(P_B) \le N(K_\beta - \delta)$ and $\log_2(1/p) \le aN$, the right-hand side exceeds the left for large N. So $p < 2^{-aN}$, and with $p \ge 1 - \varepsilon_N - q_N - ce^{-bN}$ this proves Theorem 8.

C.6 Remarks. The constants are not uniform as $\delta \to 0$. Covariance enters only through the choice of the tester and its energy-conserving realization in C.1. A degenerate battery target pays $\log_2$ of its rank, as noted after Theorem 8. For $n = 1$ the argument reduces to the fixed-error converse for channels in [4].

D Repairing small gain

We use the notation of Section 5: a dilation with gain g and distance ε to $\mathcal{P}$, the enlarged source σ, the gauged contacts W_j of (7), and $\eta_j = G_j \xi_j G_j^*$, the gauged bath after j Gibbs-driven slots, with $\eta_0 = \sigma$. By (6) and unitary invariance, $D(\sigma\|\eta_j) = \sum_{l \le j} a_l \le g$ for every j, so $\|\sqrt{\eta_j} - \sqrt{\sigma}\|_2 \le \sqrt{\ln 2\, g}$. Under $\gamma_j \otimes \eta_{j-1}$ the one-contact output is related by G_j to $U_j(\gamma_j \otimes \tau_{j-1})U_j^*$, so its relative entropy to $\gamma_j \otimes \eta_j$ is $a_j \le g$. The triangle inequality for square roots then gives (8).

Step 1: regularization. Let D_B be the dimension of the enlarged bath and $\sigma_u = (1-u)\sigma + uI/D_B$ with $u = g$. This changes the process by at most u. The gauged Gibbs-driven baths η_j^u are mixtures of η_j and another state with weight u, so

$$\|\sqrt{\eta_j^u} - \sqrt{\sigma_u}\|_2 \le \sqrt{\ln 2\, g} + 2\sqrt{2u} =: k_g, \qquad \big\|[W_j, \sqrt{\gamma_j \otimes \sigma_u}]\big\|_2 \le 3\sqrt{\ln 2\, g} + 2\sqrt{2u} =: b_g.$$

No entropy continuity on the growing bath is used.

Step 2: commensurate Hamiltonians. Replace each H_j by a Hamiltonian H_{0j}, diagonal in the same basis, whose gaps are all integer multiples of a mesh h_E; let $\delta = \max_j \|H_{0j} - H_j\|$ and γ_{0j} its Gibbs state. For fixed port dimensions,

$$\sqrt{q_j} := \big\|[W_j, \sqrt{\gamma_{0j} \otimes \sigma_u}]\big\|_2 \le b_g + C_H\delta.$$

Step 3: one offset for all contacts. Put $h = \beta h_E$ and write the eigenvalues of σ_u as $e^{-\xi_b}$. For an offset $s \in [0, h)$ round $\xi_b' = h\lfloor(\xi_b + s)/h\rfloor - s$ and normalize $\sigma' \propto e^{-\xi'}$. Then $e^{-h}\sigma_u \le \sigma' \le e^h\sigma_u$ up to normalization, and $\frac{1}{2}\|\sigma' - \sigma_u\|_1 \le (e^h - 1)/2$. At slot j the blocks of the rounded total energy $\beta H_{0j} + \xi'$ coincide with bins of $\beta H_{0j} + \xi$, because every gap of H_{0j} is a multiple of h. Let X_j be W_j pinched into these blocks and $L_j(s) = \|(W_j - X_j)\sqrt{\gamma_{0j} \otimes \sigma_u}\|_2^2$. Two eigenvalues λ_a, λ_b

fall into different bins with probability $\min\{1, |\ln(\lambda_a/\lambda_b)|/h\} \le |\sqrt{\lambda_a/\lambda_b} - 1|/(1 - e^{-h/2})$, so by Cauchy–Schwarz $\mathbb{E}_s L_j(s) \le \sqrt{q_j}/(1 - e^{-h/2})$. Choose one offset with

$$\sum_j L_j(s) \;\le\; \frac{n(b_g + C_H\delta)}{1 - e^{-h/2}}. \tag{17}$$

Separately good offsets for separate contacts would not give one bath Hamiltonian.

Step 4: polar completion. Complete each block of X_j to a unitary V_j by its polar decomposition, completing null subspaces within the same block. On each block $\gamma_{0j} \otimes \sigma'$ is a multiple of the identity, and $1 - s \le 1 - s^2$ for singular values $s \le 1$, so

$$\|(W_j - V_j)\sqrt{\gamma_{0j} \otimes \sigma_u}\|_2^2 \;\le\; 2e^{2h} L_j.$$

Every V_j conserves $H_{0j} + H'_B$ exactly, with $H'_B = -\beta^{-1} \ln \sigma'$.

Step 5: the process telescope. Feed fresh purifications of the true γ_j, keep all references and outputs, and compare the original process with the one whose contacts from j on are repaired. At contact j the bath of the original prefix is η^u_{j-1}. Since $\gamma_j \le c_\delta \gamma_{0j}$ with $c_\delta = e^{O(\delta)}$ and $\|W_j - V_j\| \le 2$,

$$\|(W_j - V_j)\sqrt{\gamma_j \otimes \eta^u_{j-1}}\|_2 \;\le\; \sqrt{c_\delta}\big[\sqrt{2e^{2h} L_j} + 2k_g\big],$$

first replacing $\sqrt{\eta^u_{j-1}}$ by $\sqrt{\sigma_u}$ at cost $2k_g$. This holds for every purification of that bath, so the correlations of earlier outputs are kept. Summing the purified vectors over j, and converting from Gibbs inputs to the flat probe with the fixed factor $R_H = \prod_j (d_j \lambda_{\min}(\gamma_j))^{-1}$ and (12), the repaired process is within

$$u + \frac{e^h - 1}{2} + d^2\sqrt{R_H c_\delta}\Big[\sqrt{\frac{2e^{2h} n^2 (b_g + C_H\delta)}{1 - e^{-h/2}}} + 2nk_g\Big] \tag{18}$$

of the original. No constant depends on the dimension, spectral range or smallest eigenvalue of the bath.

Step 6: irrational gaps. Choose a rationally independent basis $\omega = \Delta(1, v_2, \dots, v_r)$ of the group generated by all gaps of all ports, and write each energy as $E_{j,*} + a_{j,e} \cdot \omega$ with integer vectors $a_{j,e}$. For $r \ge 2$ and every integer $L \ge 2$, Dirichlet's theorem gives $1 \le q \le L^{r-1}$ and integers p_k with $|qv_k - p_k| \le 1/L$. Let $\omega_0 = \Delta(1, p_2/q, \dots, p_r/q)$, define H_{0j} with the same integer labels, and take $h_E = \Delta/(mq)$ with $m = \lceil L^{1/3} \rceil$. Then $\delta \le C/(Lq)$, and with $u = g$ the bound (18) is a fixed multiple of $L^{-1/3} + g^{1/4} L^{r/2 - 1/3} + \sqrt{g} + g$. Taking $L = \lceil g^{-1/(2r)} \rceil$ gives $O(g^{1/(6r)})$. For $r \le 1$ take $H_0 = H$ and a mesh of order $g^{1/6}$ dividing its gap, which gives $O(g^{1/6})$.

Step 7: the original Hamiltonians. The repaired process conserves H_{0j}, not H_j. Append one Gibbs lattice register with labels $\ell \in [-M, M]^r$ and Hamiltonian $H_L = E^0_L + \ell \cdot (\omega - \omega_0)$, shared by all slots and retained throughout. For each label k at distance at least $A = \max \|a_{j,e}\|_\infty$ from the boundary, the isometries $J_{j,k}|e, b\rangle = |e, b, k - a_{j,e}\rangle$ have orthogonal images, and $J^*_{j,k}(H_j + H'_B + H_L)J_{j,k} = H_{0j} + H'_B + \text{const}$. So

$$\widetilde{V}_j = \bigoplus_k J_{j,k} V_j J^*_{j,k} \oplus I$$

conserves the original $H_j + H'_B + H_L$ exactly; the identity on the complement is allowed because the complement reduces the energy. A history of inputs and outputs shifts the lattice label by its total charge $\sum_j (a_{j,\mathrm{in}} - a_{j,\mathrm{out}})$. After the lattice is traced out, two histories keep their coherence exactly when their total charges agree, so for labels far from the boundary the process is the average of the repaired process over one common phase of the whole episode. Since $\mathcal{P}$ satisfies (2), this average is no farther from $\mathcal{P}$. No separate phase flag is needed. Labels near the boundary have Gibbs weight at most $C_{r,n,A}/M + (e^{2\beta M\|\omega-\omega_0\|_1} - 1)/2$, which is $O(L^{-1/3})$ for M of order $L^{1/3}$. This gives (9).

Step 8: covariance is forced. Each repaired process before Step 7 commutes with the common time translation for H_0. If $g \to 0$ and the dilations converge to $\mathcal{P}$, the repaired processes converge to $\mathcal{P}$ and $H_0 \to H$, and at every fixed time the translations for H_0 converge to those for H. So $\mathcal{P}$ satisfies (2).

Step 9: a clock for the episodes. The zero-rate construction of Section 5 needs a controller register that indexes the episodes. Let it have dimension $D_c = nN + 1$, zero Hamiltonian, and start maximally mixed. A two-level battery with gap $W_c = \beta^{-1} \ln D_c$ starts excited, and a Gibbs ladder has levels jW_c with degeneracy D_c^j for $j = 0, \ldots, L$, each of Gibbs probability $1/(L+1)$. Energy-conserving swaps map (clock, battery excited, ladder level j) to (clock $|0\rangle$, battery ground, ladder level $j+1$) for $j < L$. The top branch fails, with probability $1/(L+1)$. At the end the clock is swapped with a fresh maximally mixed clock, so the controller returns exactly. With $L \geq N^8$ the work is $O(\log N)$, and the battery's entropy and variance tend to zero.

E The compiler without time covariance

This appendix proves Theorem 9; Theorem 11 restates it in energy units, with every hypothesis explicit. Section E.1 states the contract. Section E.2 proves the lower bound and a minimax step. Section E.3 keeps, on a rare branch, a record of the observer's strategy, and Section E.4 turns the gap between the most demanding observer and the actual one into a positive heat effect on that record, which Section E.5 pays with a hot thermal qubit. Section E.6 forms the required source from Gibbs supplies (Lemma 12). Sections E.7 and E.8 return the cache and lift the whole history to energy-conserving gates, Section E.9 fixes the order in which the parameters are chosen, and Section E.10 states the scope.

E.1 The contract and the theorem

All entropies in this appendix are natural-log entropies. Put $F_\beta(\rho) = \mathrm{Tr}\, H\rho - S(\rho)/\beta$, with $\beta > 0$. Trace distance and complete strategy distance use the half-trace convention. A fixed finite process $\mathcal{P}$ has ℓ slots $I_j \longrightarrow O_j$, with $I_j = O_j$ as physical carriers, dimension d_j, and the same finite Hamiltonian H_j on entry and exit. Write $d = \prod_j d_j$, $c = \prod_{j<\ell} d_j$, and $R_H = \sum_j \mathrm{span}(H_j)$. An episode releases O_j before I_{j+1} arrives. No released output is recalled.

An implementation of N episodes has a controller C, a mechanical battery W (written B in Section 2), and declared finite Gibbs stocks G. Initially their state is $\alpha_C \otimes |w_\mathrm{i}\rangle\langle w_\mathrm{i}| \otimes \gamma_G$, independent of the observer. The controller includes every phase register, cache, program, counter, seed and retained working cell. Its arbitrary initial athermal state is invested capital, with price $\beta^{-1} D_{\mathrm{nat}}(\alpha_C \| \gamma_C)$; no bound on that price is assumed. Every gate is a full square unitary conserving one fixed additive Hamiltonian exactly. All Gibbs stocks start independently in their true Gibbs states. Retired stocks and physical records remain in the physical rest, denoted R.

Uniformly over all adaptive observers, including cross-episode adaptation and reference entanglement, the return contract requires

$$\varepsilon_N \to 0, \qquad \tfrac{1}{2}\|\Omega_{C\mathsf{R}} - \alpha_C \otimes \Omega_{\mathsf{R}}\|_1 \le \eta_N \to 0, \qquad |\Delta E_C| = o(N), \quad |\Delta S_C| = o(N).$$

The energy and entropy here are actual unconditional quantities, not inferences from trace closeness. The battery has a fixed nominal rank-one endpoint of energy E_{nom}, with failure probability $q_W \to 0$, $S(W_{\mathrm{f}}) = o(N)$, and $\mathbb{E}|H_{W,\mathrm{f}} - E_{\mathrm{nom}}| = o(N)$, where the last expectation is the spectral expectation of the absolute operator. Nominal debit is $W_{\mathrm{nom}} = E_{W,\mathrm{i}} - E_{\mathrm{nom}}$. The endpoint is uniform in the observer.

An extension emits an inert F before the first input and satisfies all causal prefix constraints. A purified tester retains all its waste and outputs in Z; it does not read F. Define the energy-valued score

$$G_\varepsilon(\mathcal{P}) = \inf_{\substack{\mathcal{E} \text{ finite} \\ d_{\mathrm{c}}(\mathrm{Tr}_F \mathcal{E}, \mathcal{P}) \le \varepsilon}} \sup_t \left\{ \frac{S(Z|F)_t}{\beta} + E_{\mathrm{pay}}(t) \right\}, \quad k(\mathcal{P}) = \lim_{\varepsilon \downarrow 0} G_\varepsilon(\mathcal{P}).$$

Here E_{pay} is the sum of output minus input energies, on the same run and tester. In the bits of Section 2, $k = (\ln 2/\beta) K_\beta$.

Theorem 11 (Work law for a fixed process without covariance)**.** *For every fixed finite process just described, at arbitrary Choi rank and without time covariance, the infimum of limiting nominal operating work per episode among families that meet the contract is $k(\mathcal{P})$. In particular there is one finite implementation at each horizon N, serving every observer, with*

$$\begin{aligned} &\varepsilon_N = O(N^{-2}), \qquad \eta_N \to 0, \\ &|\Delta E_C| + |\Delta S_C| = o(1), \\ &S(W_{\mathrm{f}}) + \mathbb{E}|H_{W,\mathrm{f}} - E_{\mathrm{nom}}| = o(1), \\ &W_{\mathrm{nom}}/N \le k + o(1). \end{aligned}$$

Hardware and invested capital may be arbitrarily large and non-effective functions of N. The theorem is a rate law at vanishing error; Section E.10 states what it does not give.

E.2 Finite dilations, minimax and the converse

A finite extension has a causal Stinespring realization leaving F inert. Purifying its first-output state by a bath, allocating all future ancillas in that bath initially, and completing the stage isometries on their orthogonal complements gives a fixed bath B, source τ, and square contacts $U_j : I_j B \to O_j B$. The dimensions of the two carrier spaces agree, so these completions are possible. Conversely any such dilation gives an extension by its inert source purifier. For a purified tester the terminal ZBF state is pure, hence

$$S(Z|F) = S(\omega_t) - S(\tau).$$

The bath state and payload energy are affine functions of the deterministic tester operator. A tester's effective final span is at most d^2: in the Choi representation its final state is a sandwich by the square root of its deterministic tester on the d^2-dimensional port space. Any extra pure spectator is an isometric redundant extension. Thus every energy-valued score has absolute value at most

$$C_0 = \frac{2 \ln d}{\beta} + R_H. \tag{E1}$$

This bound holds for enlarged dilations too. The gain of every equal-Hamiltonian dilation is nonnegative. Feed fresh Gibbs inputs in all slots. For the bath marginal τ_{j-1} entering contact

j, entropy conservation gives its score increment as $D_{\mathrm{nat}}(U_j(\gamma_j \otimes \tau_{j-1})U_j^* \| \gamma_j \otimes \tau_j)/\beta \geq 0$. These increments telescope.

For completeness, $G_\varepsilon = \sup_t \inf_{\mathcal{E}} f_{\mathcal{E}}(t)$. Each $f_{\mathcal{E}}$ is continuous and concave on the compact tester set. Flagging a finite convex mixture of extensions in F makes its score exactly the corresponding mixture of scores; flag entropy cancels. Let $b = \sup_t \inf_{\mathcal{E}} f_{\mathcal{E}}(t)$. For any $a > 0$, continuity and compactness give a finite set $f_1, \ldots, f_L$ with $\min_i f_i(t) < b + a$ for every tester. The set of vectors dominated by $(f_1(t), \ldots, f_L(t))$ for some t is closed, convex and downward closed. It excludes $(b + a, \ldots, b + a)$. Separation supplies nonnegative weights summing to one whose flagged mixture has maximum at most $b + a$. Let $a \downarrow 0$.

Let $D = CGW$ include all internal stocks, retired ones included. Starting each episode with a fresh independent copy of one fixed tester, purify the current D by an inert F_i, containing earlier outputs and references. The complete error ε_N implies that this episode's marginal process is an ε_N-extension: use the chosen histories before it, an arbitrary tester in it, and contract the whole experiment afterwards. Global purity gives

$$S(D_{i+1}) - S(D_i) = S(Z_i|F_i).$$

Internal inter-episode unitaries do not change $S(D)$. Minimax therefore gives, for any $a > 0$, a single repeated tester with

$$\frac{\Delta S_D}{\beta} + \Delta E_{\mathrm{pay}} \geq N(G_{\varepsilon_N} - a). \tag{E2}$$

No covariance or independence from earlier outputs is used.

The exact Gibbs accounting identity is

$$W_{\mathrm{mean}} - \Delta F_\beta(C) + \frac{S(W_{\mathrm{f}})}{\beta} = \frac{\Delta S_D}{\beta} + \Delta E_{\mathrm{pay}} + \frac{D_{\mathrm{nat}}(\Omega_G \| \gamma_G) + I(C{:}W{:}G)_{\mathrm{f}}}{\beta}. \tag{E3}$$

Indeed initial stocks are product Gibbs, and $\beta \Delta E_G = \Delta S_G + D_{\mathrm{nat}}(\Omega_G \| \gamma_G)$; final multi-information equals the sum of marginal entropies minus $S(D_{\mathrm{f}})$. Equations (E2)–(E3), the actual refunds required by the contract and the battery absolute nominal-deviation bound yield $\liminf W_{\mathrm{nom}}/N \geq k$.

E.3 Recording the observer's strategy

Fix one finite dilation of gain $g = \max_t[(S(\omega_t) - S(\tau))/\beta + E_{\mathrm{pay}}(t)]$. Set a classical flag f of weight δ, and add pairs $A_j R_j$, $j < \ell$, of dimensions d_j^2. Let K denote their product. The enlarged source is

$$\tau' = \tau_B \otimes [(1 - \delta)|0\rangle\langle 0|_f \otimes \pi_K + \delta |1\rangle\langle 1|_f \otimes \Omega_K], \tag{E4}$$

where Ω_K is the product of normalized EPR states. On $f = 0$, run the original contacts and leave K idle. On $f = 1$, at each $j < \ell$ SWAP the arriving I_j with A_j, releasing the EPR half originally there; leave B idle. At the last slot use the identity. These are full square unitaries. The branch is a coherent part of the implementation; nothing is postselected. It changes the complete one-episode process by at most δ.

Immediately after input j arrives, before its contact, trace the observer's private memory and ignore unused future EPR pairs. On $A_{<j} R_{<j} I_j$, the retained state is Z_j/c_j, where $c_j = \prod_{a<j} d_a$, and

$$Z_j \geq 0, \quad \mathrm{Tr}_{I_j} Z_j = I_{R_{j-1}} \otimes Z_{j-1} \ (j > 1), \quad \mathrm{Tr}\, Z_1 = 1, \quad \mathrm{Tr}\, Z_\ell = c. \tag{E5}$$

To verify (E5), write the initial input–memory density as $\rho_{a_1 r_1; a_1' r_1'}$. Let the observer's isometry after output $j - 1$ have entries $L_j(a_j, r_j; r_{j-1}, o_{j-1})$, where r_j includes all observer waste. With

$v = (a_1, o_1, \ldots, a_{j-1}, o_{j-1}, a_j)$, the Gram matrix is explicitly

$$(Z_j)_{v;v'} = \sum_{\substack{r_1,\ldots,r_j \\ r'_1,\ldots,r'_j}} \rho_{a_1 r_1; a'_1 r'_1}\, \delta_{r_j r'_j} \prod_{k=2}^{j} L_k(a_k, r_k; r_{k-1}, o_{k-1}) \overline{L_k(a'_k, r'_k; r'_{k-1}, o'_{k-1})}.$$

For $j = 1$ this means $Z_1 = \mathrm{Tr}_{r_1}\, \rho$. Each EPR insertion contributes $d_j^{-1/2}$ in bra and ket; contraction of the observer memory gives exactly the deterministic dual Choi matrix. Tracing the last input uses the isometry normalization and produces $I_{R_{j-1}}$, with the stated factor c_j. There is no partial transpose in this physical matrix convention. Conversely every positive chain (E5) is realized by a tester: purify Z_1, and at each step apply the isometry supplied by equality of the two purifications of the required prefix, using inverses only on supports. Retain all dilation waste. Thus the physical records fill precisely this cone, including quantum-memory testers.

These rare-branch inputs are the tester's responses to the EPR outputs, not the inputs on the ordinary branch. Main-branch states and energies are affine link products of the same Z_ℓ with the dilation. Write $\chi_t = \mathrm{Tr}_{I_\ell}\, Z_\ell / c$. Flag entropy cancels, so

$$S(\omega'_t) - S(\tau') = (1-\delta)[S(\omega_t) - S(\tau)] + \delta S(\chi_t), \qquad g' \le g + \delta C_0. \tag{E6}$$

E.4 A joint tangent and a positive heat effect

Maximize $S(\omega'_t)/\beta + E'_{\mathrm{pay}}(t)$, and call the terminal bath of a maximizing tester ζ. Its support contains all terminal outputs: any outside weight in a convex mixture contributes $-x \ln x$, dominating the bounded first-order energy change. The directional entropy derivative then gives

$$E'_{\mathrm{pay}}(t) - \frac{\mathrm{Tr}\, \omega'_t \ln \zeta + S(\tau')}{\beta} \le g'. \tag{E7}$$

Let $R = \dim B'$, $q = (1-s)\zeta + s I_R/R$, and $T = -\ln q/\beta$. It is faithful, with partition function one. The flag blocks remain separate in q and T. The source τ' need not commute with T. Set

$$F = D_{\mathrm{nat}}(\tau' \| q)/\beta, \quad J(t) = E'_{\mathrm{pay}}(t) + \mathrm{Tr}\, T(\omega'_t - \tau'), \quad w = \max_t J(t), \quad a_s = -\ln(1-s)/\beta.$$

Since $q \ge (1-s)\zeta$, (E7) implies

$$F + w \le g' + a_s. \tag{E8}$$

The signed service pump w may be negative.

Let $\mathcal{V}$ be the real linear space of homogeneous tester prefix constraints, $n(Z) = \mathrm{Tr}\, Z/c$. The slack $L(Z) = wn(Z) - J(Z)$ is linear and nonnegative on $\mathcal{V} \cap \mathrm{PSD}$. The flat normalized tester is $Z_0 = I/d$. The orthogonal projection of the PSD cone into $\mathcal{V}$ is closed: if its images converge, pairing them with Z_0 bounds the trace of the positive preimages, so a bounded subsequence converges. Its dual inside $\mathcal{V}$ is $\mathcal{V} \cap \mathrm{PSD}$; finite-dimensional bipolar duality gives a positive representative D_{hist} with

$$\mathrm{Tr}\, D_{\mathrm{hist}} Z = L(Z) \quad (Z \in \mathcal{V}), \qquad \|D_{\mathrm{hist}}\| \le \mathrm{Tr}\, D_{\mathrm{hist}} = d[w - J(Z_0)] \le d(2C_0 + a_s). \tag{E9}$$

The last step uses $g' \le C_0$ directly from (E1), and $J(Z_0) + F \ge (S(\omega'_{Z_0}) - S(\tau'))/\beta + E'_{\mathrm{pay}}(Z_0) \ge -C_0$.

Before the last output, the accessible effect

$$D = |1\rangle\langle 1|_f \otimes I_B \otimes (c/\delta) D_{\mathrm{hist}}, \qquad b := \|D\| \le cd(2C_0 + a_s)/\delta \tag{E10}$$

has expectation $w - J(t)$ on every tester. It acts only on genuinely retained records and the current last input. It touches no released output, source purifier or observer memory.

E.5 The hot-qubit contact

Append a fresh independent Gibbs qubit G, gap $M \geq 2b$, excited probability $z = (1 + e^{\beta M})^{-1}$. Put

$$C_D = \sqrt{I - D/M}, \quad F_D = \sqrt{D/M}, \qquad J_D = \begin{pmatrix} C_D & -F_D \\ F_D & C_D \end{pmatrix}.$$

The whole contact is $(U'_\ell \otimes I_G)J_D$. Because C_D, F_D commute and $C_D^2 + F_D^2 = I$, this is square unitary on every column, excited columns included. Tracing G for either initial eigenstate gives $X \mapsto C_D X C_D + F_D X F_D$. The initialized heat effect is exactly $(1 - 2z)D$.

Let $Q = U'^*_\ell (H_\ell + T) U'_\ell$, and $h = \mathrm{span}(H_\ell) + \mathrm{span}(T)$. The remainder

$$R_M = C_D Q C_D + F_D Q F_D - Q - 2zD, \qquad \|R_M\| \leq r_M := 3hb/M + 2zb \tag{E11}$$

follows after subtracting a scalar from Q. All earlier bath energies telescope with the same fixed T. The total initialized episode energy gain is therefore $w + \langle R_M \rangle$ for every tester. The source τ' need not commute with T. Since D is zero on $f = 0$, the entire ordinary branch still implements the original dilation exactly; for error it suffices to charge the rare probability δ, not the norm b/M.

E.6 Forming the source from Gibbs supplies

Lemma 12 (Finite-source module)**.** *Let τ be a fixed finite source on a retained body, q a faithful state with $T = -\beta^{-1} \ln q$, and consider a finite sequence of full square gates on held systems. If a terminal contact makes the initialized total external energy gain $w + r(t)$ uniformly in testers, then source formation, an invested one-source cyclic cache, a finite conserving whole-history phase lift, sharp pumps and a terminal Gibbs swap reset give nominal work $\beta^{-1} D_{\mathrm{nat}}(\tau \| q) + w + o(1)$, with complete error $o(1)$, joint return of the whole controller with all physical rest within $o(1)$, actual controller energy and entropy changes $o(1)$, and battery entropy and absolute nominal deviation $o(1)$, as $r \to 0$ and the parameters below are chosen. The proof is the rest of this subsection and the next two; it does not use square slots.*

This section proves a pointwise module for any fixed finite source τ on a Hamiltonian T, not necessarily commuting with it. Let $R = \dim B$, $\gamma = e^{-\beta T}/Z_T$, and $F_\tau = D_{\mathrm{nat}}(\tau \| \gamma)/\beta$. For each sufficiently large integer m, it produces a state ρ_m with

$$\tfrac{1}{2} \|\rho_m - \tau^{\otimes m}\|_1 \leq e_m \to 0, \qquad a_m \to F_\tau, \qquad W_{\mathrm{factory,nom}} = m a_m. \tag{E12}$$

All its nonphase work deviations are bounded by mA_τ on a failure event of probability f_m, where $A_\tau < \infty$ is fixed for this source and f_m can be chosen arbitrarily small. The factory's output may stay correlated with the coding and thermal residues. The joint return of Lemma 12 does not need them to be independent; it is obtained through the cyclic cache of Section E.7.

Energy-shell support conversion

For commuting projectors P, Q on finite X, H_X, write $Z_P = \mathrm{Tr}\, P e^{-\beta H_X}$, $Z_Q = \mathrm{Tr}\, Q e^{-\beta H_X}$. For any extracted work $x < \beta^{-1} \ln(Z_Q/Z_P)$ and $f > 0$, a finite Gibbs stock and full conserving square take every P-supported state, even reference-entangled, to Q with sharp credit x on a stock event of probability at least $1 - f$, independent of that state.

Here is the exact shell construction. Choose an integer basis ω of the finite group generated by the energies and x. Let reservoir tiles j in a finite box have energy $E_0 + j \cdot \omega$ and integer degeneracy

$g_j = \lfloor Ke^{\beta j\cdot\omega} \rfloor$. Choose K so $(1-\theta)Ke^{\beta j\cdot\omega} \le g_j \le Ke^{\beta j\cdot\omega}$ on that finite box. Tile Gibbs weights lie between common factors $1-\theta$ and 1. A margin larger than every source–target energy-coordinate difference and the work shift makes the inner tiles have probability $1-f$ when the box is enlarged.

In shell coordinate u, the good source dimension is at most $Ke^{\beta u\cdot\omega}Z_P$; the available target dimension is at least $(1-\theta)Ke^{\beta(u\cdot\omega-x)}Z_Q$. Choose $(1-\theta)e^{\beta[\ln(Z_Q/Z_P)/\beta-x]} > 1$. Inject every source basis column isometrically into that target shell. Source and image carry orthogonal work labels. Exchange these two subspaces and use identity on their orthogonal complements. This is a full conserving square, including every boundary and unused column. Bad initial tiles leave the initialized work label unchanged. For negative x this is a debit; at zero gap the labels remain orthogonal. Origins E_0 are fixed explicitly to make finite Hamiltonians nonnegative. This is a support statement, not a universal output-Gibbs-state statement.

Exact flat-rank formation

For zero-Hamiltonian dimension D and target rank $K \le D$, the debit is $\ln(D/K)/\beta$. Tiles $j = 0, \ldots, J$ have that gap and degeneracy $D^j K^{J-j}$, hence exactly equal Gibbs probability. The full initialized flat-D block at $j < J$ and the full target rank-K block at $j+1$ have equal dimension. Exchange them with orthogonal work labels; cap columns stay unchanged. Conditional output is exactly flat rank K, since the target block is fully occupied by a flat source. Failure is $1/(J+1)$. This distributional conclusion uses the fixed flat source, whereas the preceding support converter has an all-state quantifier.

Four-work-qubit factory and the residue

Let $\sigma = S(\tau)/\ln 2$, $\gamma_m = m^{-1/4}$,

$$r = \lceil 3m\gamma_m \rceil + 4, \quad D = R^m 2^r, \quad K = 2^{\lceil m(\sigma+2\gamma_m)\rceil}.$$

The controller coding body is zero-Hamiltonian A^m of dimension R^m and a zero-Hamiltonian residue of dimension 2^r, initially flat on dimension D. Four sharp work qubits record these four operations.

1. Form physical $\gamma^{\otimes m}$ into one ground vector $|g\rangle^{\otimes m}$ by the support converter, with debit $m(h_{\min} + \ln Z_T/\beta) + s_0$, where $s_0 > 0$ is fixed. Its failure can be at most $f_m/3$; on success the source is exactly pure and hence product with everything.

2. Form flat rank K on coding body plus residue, debit $\ln(D/K)/\beta$, failure at most $f_m/3$. In the τ-eigenstring basis assign $\lfloor Kp_x \rfloor$ distinct flat labels to each typical x, with distinct residue labels. The typical bounds are $2^{-m(\sigma+\gamma_m)} \le p_x \le 2^{-m(\sigma-\gamma_m)}$. Each fiber fits in the residue, and the assignment completes to a square zero-Hamiltonian unitary. The marginal coding error is bounded by
$$e_{\text{code}} \le 2e^{-c_\tau\sqrt{m}} + 2^{-m\gamma_m}, \tag{E13}$$
on the nonzero spectrum of τ; if its log-eigenvalue variable is constant, the atypical term is zero. The first term follows from bounded iid log-likelihood concentration, the second from at most $2^{m(\sigma+\gamma_m)}$ rounded labels. Leftovers are placed arbitrarily. Correlation with the residue is retained.

3. Lift the whole SWAP $A^m \leftrightarrow B^m$ with a phase register, then pump the known $m(\operatorname{Tr}\tau T - h_{\min})$. The physical source receives the coding marginal. On ground-formation success, the coding output is exactly $|g\rangle^{\otimes m}$, even in the lifted SWAP: each output term places the old energy eigenvector in zero-Hamiltonian A^m.

4. Extract $m \ln R/\beta - s_0$ from the pure coding support, using the support converter on $|g\rangle\langle g|^{\otimes m} \otimes I_{\text{residue}}$ to the whole coding/residue space. The entire residue participates, so no correlation is omitted. Then SWAP coding body plus residue with an independent flat Gibbs stock of dimension D. They return exactly to their initial flat state, independently of all physical rest. Extraction failure from impure coding is already included in ground-formation failure.

The sum of the four nominal debits/credits is

$$ma_m = m[\operatorname{Tr}\tau T + \ln Z_T/\beta - S(\tau)/\beta] + O(m^{3/4}) + O(1) = mF_\tau + o(m). \tag{E14}$$

One may choose positive overhead, hence $a_m \geq 0$. All four work gaps have magnitude at most mA_τ for large m. The nonphase failures contribute f_m and can be made arbitrarily small by choosing finite tile boxes and caps.

If the phase-vector error is d_f, exact SWAP conservation and the known pump, compared only with the touched physical B^m energies, give

$$|\Delta E_{P_f}| \leq Cm \operatorname{span}(T)(e_m + d_f), \quad e_m \leq e_{\text{code}} + f_m. \tag{E15}$$

There is no multiplication by the phase register's large energy span. The phase construction and entropy bound are proved below.

E.7 A cyclic cache and its joint return

Fix this source and factory, then choose $m > N$. Use $m-1$ cache cells with Hamiltonian T, and a classical clock $k = 0, \ldots, m-1$ of fixed energy $(m-1-k)a_m$. Invest the correlated controller state

$$\sigma_m = \frac{1}{m} \sum_{k=0}^{m-1} |k\rangle\langle k| \otimes \tau^{\otimes k} \otimes \gamma^{\otimes(m-1-k)}. \tag{E16}$$

Each call receives a fresh independent γ carrier. For $k > 0$, SWAP out the first cached τ, permute cache cells, and decrement k. Raising the clock costs exactly a_m, supplied by the main sharp battery. The equal-dimensional product-support identities give full-square conserving permutations. For $k = 0$, the clock drops to $m-1$, releasing $(m-1)a_m$; with the same a_m main debit it charges a two-label auxiliary work register by ma_m. Apply the m-copy factory to the cache cells and fresh carrier. Every path advances the clock. The auxiliary register and the four factory work qubits are part of W, not unpriced controller records. On a successful boundary their total energy equals its no-boundary value: the factory debit ma_m cancels the auxiliary credit.

The N calls cross a boundary for exactly N initial clock values, so $p = N/m$, independently of the future observer. On every no-boundary sector the N emitted carriers are exactly $\tau^{\otimes N}$, product with cache and all pre-existing physical rest. The resulting cache sectors agree with the corresponding sectors of (E16). Boundary sectors have total weight p, and may carry all source–residue correlations. Comparing both joint states on these sectors, and then replacing the comparison rest marginal by the actual one, gives the conservative uniform bound

$$\tfrac{1}{2}\|\Omega_{\text{cache,rest}} - \sigma_m \otimes \Omega_{\text{rest}}\|_1 \leq 4p + 2pe_m, \quad \tfrac{1}{2}\|\Omega_{\text{cache}} - \sigma_m\|_1 \leq pe_m, \quad \tfrac{1}{2}\|\rho_{\text{emitted}} - \tau^{\otimes N}\|_1 \leq pe_m. \tag{E17}$$

The first inequality is weaker than the $2p$ estimate available for an exact factory, but it proves return jointly with the whole physical rest without assuming that source and residue are independent. Subsequent service cannot increase it because it never acts on the cache. Equations (E17)–(E18) initially describe the bare factory. For the physical factory let $\widetilde{e}_m = e_m + d_f$, where d_f is its

conditional full-vector error on any boundary sector. On no-boundary sectors the factory phase register is never operated. Consequently the actual cache and emitted-source marginal errors are at most $p\widetilde{e}_m$, not an unweighted d_f. In the following refund bounds e_m is to be replaced by $\widetilde{e}_m$ for the physical implementation. The all-rest estimate remains bounded by the displayed conservative expression with that replacement, and the phase factor itself is returned by (E23).

Put $D_{\mathrm{cache}} = mR^{m-1}$. Separate actual refunds follow from the marginal bound, not from the joint bound:

$$|\Delta S_{\mathrm{cache}}| \le h(pe_m) + pe_m \ln(D_{\mathrm{cache}} - 1), \qquad |\Delta E_{\mathrm{cache}}| \le Ne_m[\mathrm{span}(T) + a_m]. \tag{E18}$$

Here h is binary entropy in nats. Formula (E15) is weighted by boundary probability p, giving $O(N\,\mathrm{span}(T)(e_m + d_f))$ for factory-phase mean drift. The combined five helper work registers have at most 32 labels. They differ from their nominal no-boundary vector only on a boundary, so

$$q_{W,\mathrm{helpers}} \le p, \quad S(W_{\mathrm{helpers}}) \le h(p) + p\ln 32, \quad \mathbb{E}|H_{W,\mathrm{f}} - E_{\mathrm{nom}}| \le N(4A_\tau + a_m)f_m. \tag{E19}$$

Success-boundary helpers need not be in the nominal vector, but their total energy is nominal. Failed boundaries are retained, not deleted. The work estimate is an absolute deviation bound.

The invested cache capital is explicitly

$$\frac{D_{\mathrm{nat}}(\sigma_m \| \gamma_{\mathrm{cache}} \otimes \gamma_{\mathrm{clock}})}{\beta} = \frac{m-1}{2}(F_\tau + a_m) + \frac{\ln(Z_{\mathrm{clock}}/m)}{\beta}, \qquad Z_{\mathrm{clock}} = \sum_{k=0}^{m-1} e^{-\beta(m-1-k)a_m}. \tag{E20}$$

This cache is an invested Duan-type correlated state, not a free Gibbs supply. All N calls are made before the first observer input. The factory convergence is pointwise for this fixed source; m is chosen only after the horizon's dilation and every source constant.

E.8 Whole-history phase lift and finite completion

Every finite set of gaps and signed pump energies generates a finite-rank free integer group. Choose positive generators $\omega_1, \ldots, \omega_r$, flipping signs when necessary, and integer labels a_e for all exact energy differences. On an analytic bilateral lattice $T_v|b\rangle = |b + v\rangle$, define

$$\mathcal{L}(X) = \sum_{e,f} P_e X P_f \otimes T_{a_f - a_e}. \tag{E21}$$

Intermediate energy projectors show $\mathcal{L}(XY) = \mathcal{L}(X)\mathcal{L}(Y)$, $\mathcal{L}(X^*) = \mathcal{L}(X)^*$, and $\mathcal{L}(I) = I$. Thus whole square contacts lift to whole unitaries, preserving all coherent branch Grams and exactly conserving external energy plus $H_P|b\rangle = (b \cdot \omega)|b\rangle$. A sharp work-label transition losing $x = k \cdot \omega$ adds k to the phase translation. This allows signed w too.

The physical phase register is a finite positive-energy box. Initialize a product sine vector of width L, offset by a buffer B_* in each coordinate,

$$|\xi_L\rangle = \bigotimes_{a=1}^{r} \sqrt{\frac{2}{L+1}} \sum_{u=0}^{L-1} \sin\frac{\pi(u+1)}{L+1} |B_* + u\rangle.$$

Choose B_* larger than the sum of absolute coordinate excursions on every possible path, including excited Gibbs inputs, hot quanta, factory failures and all work branches. The box side is $D_P = L + 2B_* + 1$. On all reachable columns the analytic lift is an isometry in each exact total-energy shell. Domain and image have equal rank there; complete orthonormal bases of their complements

in the same finite shell. This gives a full finite conserving square, with no boundary error on any physically reachable path. A pure zero-Hamiltonian scheduling counter is invested in C, selects the finitely many gate blocks, advances on every path, and returns cyclically to its initial label. The controlled direct sum and counter shift are full conserving squares. This is not an autonomous/local circuit efficiency claim.

Purify every stock and input mathematically, retaining all observer waste. Fourier transformation of the analytic lattice gives the whole contact fibre $D_\theta U D_\theta^*$, with pumps as scalar phases. The observer's isometries are θ-independent. If the sum of finite fibre Lipschitz constants is K_*, the entire adaptive purified circuit has $\|V_\theta - V_0\| \le K_* \|\theta\|_2$, uniformly in all observer dimensions and correlations. The sine nearest-neighbor overlap is $\cos(\pi/(L+1))$, and $\mathbb{E}\|\theta\|_2^2 \le r\pi^4/[4(L+1)^2]$. Parseval yields

$$\|\Psi_{\text{physical}} - \xi_L \otimes \Psi_{\text{bare}}\| \le d_P := \frac{\pi^2 \sqrt{r}\, K_*}{2(L+1)}. \tag{E22}$$

It holds at every prefix with the prefix Lipschitz sum. In particular

$$q_P = 1 - \langle \xi_L | \Omega_P | \xi_L \rangle \le d_P^2, \quad \tfrac{1}{2}\|\Omega_{P,\text{rest}} - |\xi_L\rangle\langle\xi_L| \otimes \Omega_{\text{rest}}\|_1 \le 2d_P, \quad S(P) \le h(d_P^2) + d_P^2 \ln(D_P^r - 1). \tag{E23}$$

The first projection and normalization cost at most $\sqrt{q_P}$; replacing its rest marginal by the actual marginal costs another $\sqrt{q_P}$. Inert purifiers are never acted on. Use separate factory and service phase registers if convenient.

Actual service mean drift requires another argument. After all cache calls, trace retired factory residues for the comparison, not for physical accounting. The emitted marginal is within pe_m, or conservatively p, of fresh $\tau'^{\otimes N}$. Uniform prefix comparison by (E22) costs d_s. Let $Q_N < \infty$ bound the energy span of one touched service system, including H_j, T and the hot gap M, and all loading/reset swaps. Comparison of pre- and post-contact energies, not of H_P, gives

$$|\Delta E_{P_s}| \le N r_M + CNQ_N(p + d_s). \tag{E24}$$

Exact conserving lifts and pumps imply this equality/bound directly. This is valid even when τ' fails to commute with T; one must not move τ' through an energy lift under a partial trace. The ideal comparison has fresh sources before their first use, even when the observer is entangled with future inputs. Past conditional records cannot steer an untouched independent source. After each episode's last output leaves, SWAP its service body with a fresh q Gibbs cell of the same T. This exactly returns those working controller cells to their original q's, independently of physical rest. Loading sources uses equal-Hamiltonian swaps before inputs. No released output participates.

E.9 Order of choices and the ledger

For each N, choose one finite dilation within N^{-3} of $\mathcal{P}$, with $g \le k + N^{-1/8}$. This is justified by the definition of the right-error limit and needs no exactification. Choose $\delta = N^{-3}$, $s = N^{-1/8}$, then form (E4)–(E10). Now the finite R, T, F, w, b are known. Choose M so $Nr_M \le N^{-3}$, and then record finite Q_N. Only now choose the source factory and cache length $m > N$ so

$$\begin{aligned} &a_m - F \le N^{-1/8}, \qquad Ne_m(1 + \ln R + \operatorname{span}(T) + a_m + A_\tau) \le N^{-3}, \\ &N(4A_\tau + a_m) f_m \le N^{-3}, \qquad N^2 Q_N/m \le N^{-3}, \qquad h(N/m) + (N/m)\ln 32 \le N^{-3}. \end{aligned} \tag{E25}$$

Pointwise convergence (E12)–(E14) allows these simultaneous choices for this already fixed source, even if its dimensions or spectrum change arbitrarily with N. Choose all stocks, energy bases and all-path buffers first; choose phase widths last so their entropy bounds and $N(1+\ln R+\operatorname{span}(T)+a_m)d_f$,

$NQ_N d_s$ are at most N^{-3}. This also makes $h(p\widetilde{e}_m) + p\widetilde{e}_m \ln(D_{\text{cache}} - 1) = o(1)$; indeed $\ln D_{\text{cache}} = \ln m + (m-1)\ln R$.

On N ordinary episodes, the target error is at most $N \cdot N^{-3}$; rare branches add $N\delta$. Factory/source and phase comparisons give

$$\varepsilon_N \le N^{-2} + N\delta + p + d_f + d_s = O(N^{-2}), \qquad \eta_N \le 4p + 2p\widetilde{e}_m + 2d_f + 2d_s \to 0. \tag{E26}$$

The marginal cache bound, (E18), controls its actual entropy and energy. Factory and service phase actual energies are controlled by (E15), (E24), (E25), not their spans. Their entropies are controlled by (E23); Araki–Lieb bounds the change in joint controller entropy by the cache entropy change plus the sum of phase entropies. All other coding/service factors are swapped exactly to their initial Gibbs/flat states and the counter returns exactly. Therefore $|\Delta E_C| + |\Delta S_C| = o(1)$, uniformly over observers, and (E26) refers to the whole controller versus all physical stocks, records, the battery, the delivered outputs and the observer's references; (E17) and (E22) hold for the purified state of all of these.

The main battery has fixed service/source-counter labels with signed increments $a_m + w$, ending sharply on every initialized column. The boundary auxiliary and four factory helpers are included in the same mechanical ledger. Equations (E19), (E25) give

$$q_W \le p, \qquad S(W_{\text{f}}) \le h(p) + p\ln 32 = o(1), \qquad \mathbb{E}|H_{W,\text{f}} - E_{\text{nom}}| \le N(4A_\tau + a_m) f_m = o(1).$$

A conservative finite battery dimension is $32(N+1)^2$. Its initial charged helper capital and energy span can be enormous; both are declared. The nominal total debit and its rate obey

$$W_{\text{nom}} = N(a_m + w), \qquad \frac{W_{\text{nom}}}{N} \le g + \delta C_0 + a_s + N^{-1/8} \le k + o(1). \tag{E27}$$

With (E3), this proves Theorem 11.

A full controller inventory is: cache/clock of dimension mR^{m-1}; N service bodies of dimension R initially q; zero-Hamiltonian coding body/residue of dimension $R^m 2^r$, initially flat; two finite phase boxes; and a pure cyclic counter. Its cache capital is (E20). Each phase contributes

$$\langle H_P\rangle_{\xi_L} + \ln Z_P/\beta, \qquad \langle H_P\rangle_{\xi_L} = [B_* + (L-1)/2]\sum_a \omega_a, \quad Z_P = \prod_a \sum_{b=0}^{D_P-1} e^{-\beta b\omega_a}.$$

The counter costs $\ln D_{\text{ctr}}/\beta$; flat coding and thermal service factors cost zero relative to their own Gibbs states. These are startup investments, not operating debits or thermally prepared free clocks.

Gibbs supplies are fresh source carriers, hot qubits, support-converter tiles, exact flat-rank tiles, coding/residue reset stocks, and service q-reset stocks. All are finite and initially independent. Every spent tile and old controller state remains in physical rest. Rare hot stocks and cache boundary events can carry input-dependent energy far larger than order N, with tiny probabilities, and the cache and service reference can be off by order N on common branches while the unconditional mean is refunded.

E.10 Scope

The construction is a finite-hardware existence proof for fixed processes whose ports carry equal Hamiltonians. Its source and phase modules accept any fixed finite source and any finite sequence of square gates on systems the device actually holds; another access schedule must supply its own carrier inventory, record of the observer's strategy and positive dual, as Appendix B of [26] does

for general schedules. Rectangular ports, free switches of the Hamiltonian and recall of released outputs are not covered. The proof uses causal Stinespring purification and finite-dimensional cone separation, both reconstructed above; it uses no Gibbs-preserving map as a black box, no erasure against released outputs and no action on a purifier. The theorem is a rate law at vanishing error. It is not a strong converse, and it gives no exact or branch-by-branch return of the controller, no bound on capital or on the tails of the heat, no efficient circuit, and no construction for an unknown horizon or for local or autonomous hardware.

F The correlated process

This appendix derives the value $K_\beta = \beta v_\vartheta / \ln 2$ of the correlated process of Section 7. Both slots carry a qubit with $H = E|1\rangle\langle 1|$; write $Z = |0\rangle\langle 0| - |1\rangle\langle 1|$, and let branch b apply U_b in slot 1 and V_b in slot 2. The energy the second gate adds on an input ρ is $\operatorname{Tr}\rho D_b$ with $D_b = V_b^* H V_b - H$, so $D_0 = EZ$ and $D_1 = \frac{E}{2}(Z - X)$. The first gates commute with H and add no energy. If ρ_b is the second input on branch b, the energy an episode adds is therefore $\frac{1}{2}\sum_b \operatorname{Tr}\rho_b D_b$, and $e_\beta(t)$ is $\beta/\ln 2$ times it.

Step 1: the reachable pairs. The observer prepares the first input A together with a memory in a pure state ψ, receives $U_b A$, and applies an operation, which cannot depend on the hidden b, to produce the second input. Before that operation the two branch states are pure, with overlap $|\langle\psi|U_1 \otimes I|\psi\rangle| = |\cos\vartheta + \mathrm{i}\sin\vartheta\ \operatorname{Tr}\sigma Z| \geq |\cos\vartheta|$, where σ is the state of A. Root fidelity does not decrease under channels and partial traces, so every reachable pair has $F(\rho_0,\rho_1) = \|\sqrt{\rho_0}\sqrt{\rho_1}\|_1 \geq |\cos\vartheta|$. Conversely, let $F(\rho_0,\rho_1) = f \geq \cos\vartheta > 0$. By Uhlmann's theorem the two states have purifications $|\phi_0\rangle, |\phi_1\rangle$ with $\langle\phi_0|\phi_1\rangle = f$. Adjoin an ancilla in states $|a_0\rangle, |a_1\rangle$ with $\langle a_0|a_1\rangle = \cos\vartheta/f$. The pair $|\phi_b\rangle|a_b\rangle$ then has overlap $\cos\vartheta$, as does the pair $U_b|+\rangle$, so an isometry maps the second pair onto the first. An observer that sends $|+\rangle$, applies this isometry to the first output and passes on the system of ρ_b produces exactly (ρ_0,ρ_1).

Step 2: the energy maximum. Write a qubit state as $\frac{1}{2}(I + r\cdot\vec{\sigma})$ with $|r| \leq 1$, and lift r to the unit vector $R = (r, \sqrt{1-|r|^2})$ in $\mathbb{R}^4$. For qubits $F(\rho,\rho')^2 = \operatorname{Tr}\rho\rho' + 2\sqrt{\det\rho\det\rho'} = \frac{1}{2}(1 + R\cdot R')$, so the constraint $F(\rho_0,\rho_1) \geq \cos\vartheta$ says that the angle between R_0 and R_1 is at most 2ϑ. The energy is

$$\tfrac{1}{2}\sum_b \operatorname{Tr}\rho_b D_b = \tfrac{E}{2}\Big(R_0\cdot u + \tfrac{1}{\sqrt{2}}\,R_1\cdot w\Big),$$

where $u = (0,0,1,0)$ and $w = (-1,0,1,0)/\sqrt{2}$ are the Bloch directions of the top eigenvectors of D_0 and D_1, at angle $\pi/4$. Let $0 < \vartheta < \pi/8$ and $\gamma = \pi/4 - 2\vartheta > 0$. By the triangle inequality for angles on the sphere, the angles α of R_0 from u and β of R_1 from w satisfy $\alpha + \beta \geq \gamma$. Since the cosine decreases on $[0,\pi]$, the energy is at most $\frac{E}{2}\max\{\cos\alpha + \frac{1}{\sqrt{2}}\cos\beta : \alpha,\beta \geq 0,\ \alpha+\beta = \gamma\}$, and this maximum is the length of the sum of two plane vectors of lengths 1 and $1/\sqrt{2}$ at angle γ:

$$\tfrac{E}{2}\sqrt{1 + \tfrac{1}{2} + \sqrt{2}\cos\gamma} = \tfrac{E}{2}\sqrt{\tfrac{3}{2} + \cos 2\vartheta + \sin 2\vartheta} = v_\vartheta.$$

It is attained by two pure states in the plane of u and w, at angle 2ϑ, which Step 1 makes reachable. Mixed states do no better, because the lift turns them into points of the same sphere. When $|\cos\vartheta| \leq \cos(\pi/8)$, the two top eigenvectors, whose root fidelity is $\cos(\pi/8)$, are themselves reachable, and the value is each branch's own maximum, $\frac{E}{2}(1 + 1/\sqrt{2})$.

Step 3: no entropy credit. Let $|C_b\rangle\!\rangle$ be the Choi vector of branch b, the two-slot comb that applies U_b and then V_b with no memory. The Choi operator of $\mathcal{P}$ is $\frac{1}{2}\sum_b |C_b\rangle\!\rangle\langle\!\langle C_b|$, and that of an exact extension $\widetilde{\mathcal{P}}$ is positive with this marginal, so it is supported on the span of the two vectors tensored with F:

$$\widetilde{R} = \sum_{b,b'} |C_b\rangle\!\rangle\langle\!\langle C_{b'}| \otimes G_{bb'},$$

with operators $G_{bb'}$ on F. Tracing out the second output gives $\mathrm{Tr}_{O_2} |C_b\rangle\!\rangle\langle\!\langle C_{b'}| = |U_b\rangle\!\rangle\langle\!\langle U_{b'}| \otimes (V_{b'}^* V_b)^T$. The causal constraint at slot 2 requires $\mathrm{Tr}_{O_2} \widetilde{R}$ to be the identity on the second input tensored with an operator on everything else. The diagonal terms have this form. In the off-diagonal terms, $V_1^* V_0 = \frac{1}{\sqrt{2}}\left(\begin{smallmatrix} 1 & 1 \\ -1 & 1 \end{smallmatrix}\right)$ is not a multiple of the identity, and $|U_0\rangle\!\rangle\langle\!\langle U_1|$ and $|U_1\rangle\!\rangle\langle\!\langle U_0|$ are linearly independent because U_1 is not a multiple of U_0; so the traceless part of $(V_1^* V_0)^T$ can be cancelled only if $G_{01} = G_{10} = 0$. Hence $\widetilde{R} = \sum_b |C_b\rangle\!\rangle\langle\!\langle C_b| \otimes G_{bb}$: F is only classically correlated with b. Against any purified tester, the final state of Z and F is then $\sum_b |\psi_b\rangle\langle\psi_b| \otimes g_b$, with $|\psi_b\rangle$ the tester's pure final state on branch b and $g_b \geq 0$ on F. This state is separable, so $S(Z|F) \geq 0$. A classical copy of b, with $g_b = \frac{1}{2}|b\rangle\langle b|$, gives $S(Z|F) = 0$ for every tester, so that extension scores exactly the energy, and its best tester scores $\beta v_\vartheta / \ln 2$. For an ε-extension the block structure holds up to an error controlled by Uhlmann's theorem, and the continuity bound for conditional entropy, which uses only $\dim Z \leq 16$ and not the size of F, gives $S(Z|F) \geq -b(\varepsilon)$ with $b(\varepsilon) \to 0$ as $\varepsilon \to 0$ at fixed ϑ; its constants grow like $\csc\vartheta$ as $\vartheta \to 0$, where the two branches merge. Together with Steps 1 and 2, this gives $K_\beta = \beta v_\vartheta / \ln 2$.

G Results from the companion papers

Table 2 lists every result of the companion papers and of our earlier papers on which an argument here depends, and where it is used. Results mentioned only for comparison are cited in the text. Numbers refer to the versions cited in the bibliography.

Result	Source	Used in
device model; accounting identity	[4, Sec. 2, Lemma 4]	Section 2, Eq. (3), Appendix C
real-time entropy cost of a channel and its rate κ_β	[4, Def. 1]	Proposition 5(a), Section 7
invariance argument for an exactly returned controller	[4, Prop. 10]	Remark 1
fixed-error converse for channels	[4, Thm. 6, App. B]	Theorem 8, Appendix C
$\|\sqrt{\rho} - \sqrt{\sigma}\|_2^2 \leq D_{\mathrm{nat}}(\rho\|\sigma)$	[4, Eq. (15)]	Section 5
irrational gaps: commensurate approximation and a thermal lattice	[4, Thm. 9, App. C]	Section 5
a clock that returns approximately runs every channel at κ_β	[4, Thm. 11]	delayed channels, Section 7

Table 2: Results of the companion papers and of our earlier papers on which an argument here depends.

Appendix E, including Lemma 12, which [4] imports, uses none of these results, so the dependency between the two papers runs one way at the level of lemmas.

References

[1] Rolf Landauer. Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development*, 5:183–191, 1961.

[2] Philippe Faist, Mario Berta, and Fernando G. S. L. Brandão. Thermodynamic capacity of quantum processes. *Physical Review Letters*, 122:200601, 2019. arXiv:1807.05610.

[3] Philippe Faist, Mario Berta, and Fernando G. S. L. Brandão. Thermodynamic implementations of quantum processes. *Communications in Mathematical Physics*, 384:1709–1750, 2021. arXiv:1911.05563.

[4] Seth Douglas and Nidhal Mghirbi. Only quantum channels save work by waiting. Preprint, version 1.0.2, Zenodo, 2026. doi:10.5281/zenodo.23310624.

[5] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Memory effects in quantum channel discrimination. *Physical Review Letters*, 101:180501, 2008.

[6] Gus Gutoski and John Watrous. Toward a general theory of quantum games. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing*, pages 565–574, 2007.

[7] Jayne Thompson, Paul M. Riechers, Andrew J. P. Garner, Thomas J. Elliott, and Mile Gu. Energetic advantages for quantum agents in online execution of complex strategies. *Physical Review Letters*, 135:160402, 2025. arXiv:2503.19896. doi:10.1103/dryx-l5rt.

[8] Simon Milz, Kyrylo Simonov, Zoltán Zimborás, Tamal Guha, Saptarshi Roy, and Giulio Chiribella. Higher-order quantum thermodynamics: equilibrium and causal structure, 2026. arXiv:2609.15829.

[9] Philippe Faist, Jonathan Oppenheim, and Renato Renner. Gibbs-preserving maps outperform thermal operations in the quantum regime. *New Journal of Physics*, 17:043003, 2015. arXiv:1406.3618.

[10] Uffe Haagerup and Magdalena Musat. Factorization and dilation problems for completely positive maps on von Neumann algebras. *Communications in Mathematical Physics*, 303:555–594, 2011. arXiv:1009.0778.

[11] Yuqiang Ding, Feng Ding, and Xueyuan Hu. Exploring the gap between thermal operations and enhanced thermal operations. *Physical Review A*, 103:052214, 2021. arXiv:2101.05690.

[12] Lukas J. Fiderer, Paul C. Barth, Isaac D. Smith, and Hans J. Briegel. Information thermodynamics of agents: The work capacity of channels with memory. *Physical Review X*, 16:041004, 2026. arXiv:2504.06209. doi:10.1103/7nds-tjr8.

[13] Alexander B. Boyd, Dibyendu Mandal, and James P. Crutchfield. Thermodynamics of modularity: structural costs beyond the Landauer bound. *Physical Review X*, 8:031036, 2018. arXiv:1708.03030.

[14] Samuel P. Loomis and James P. Crutchfield. Thermodynamically-efficient local computation and the inefficiency of quantum memory compression. *Physical Review Research*, 2:023039, 2020. arXiv:2001.02258.

[15] Seth Douglas and Nidhal Mghirbi. The price of keeping time. Preprint, version 1.0.2, Zenodo, 2026. doi:10.5281/zenodo.23310632.

[16] Seok Hyung Lie, Jeongrak Son, Paul Boes, Nelly H. Y. Ng, and Henrik Wilming. Thermal operations from informational equilibrium. *Physical Review Letters*, 137:030403, 2026. arXiv:2507.16637.

[17] Yutong Luo, Jakub Czartowski, Felix Hubmann, Simon Milz, and Felix C. Binder. The trinity of Markovian quantum thermodynamics: unifying the axiomatic, microscopic, and operational paradigms, 2026. arXiv:2607.15376.

[18] Seth Douglas. Bath dimension and initial entropy for closed repeated use of a quantum channel, 2026. arXiv:2609.18267.

[19] Seth Douglas and Nidhal Mghirbi. Causal quantum-channel simulation: memory beyond entropy. Zenodo, version 1.1.2, 2026. doi:10.5281/zenodo.23070870.

[20] Seth Douglas and Nidhal Mghirbi. No algorithm can tell how much memory a repeated quantum channel needs. Zenodo, version 1.1.0, 2026. doi:10.5281/zenodo.23204398.

[21] Guilherme Zambon and Gerardo Adesso. Quantum processes as thermodynamic resources: the role of non-Markovianity. *Physical Review Letters*, 134:200401, 2025. arXiv:2411.05559.

[22] Ruo Cheng Huang, Isha Singh Le Xue, Yuxuan Qu, Paul M. Riechers, Varun Narasimhachar, and Mile Gu. Time-ordered free energy in correlated quantum systems: An agentic approach, 2026. arXiv:2608.12942.

[23] Charles H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development*, 17:525–532, 1973.

[24] Charles H. Bennett. The thermodynamics of computation—a review. *International Journal of Theoretical Physics*, 21:905–940, 1982.

[25] Alexander B. Boyd, Dibyendu Mandal, Paul M. Riechers, and James P. Crutchfield. Transient dissipation and structural costs of physical information transduction. *Physical Review Letters*, 118:220602, 2017. arXiv:1612.08616.

[26] Seth Douglas and Nidhal Mghirbi. The price of time order. Preprint, version 1.0.2, Zenodo, 2026. doi:10.5281/zenodo.23310622.

[27] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Theoretical framework for quantum networks. *Physical Review A*, 80:022339, 2009. arXiv:0904.4483.

[28] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Optimal covariant quantum networks. In *Quantum Communication, Measurement and Computing (QCMC 2008)*, volume 1110 of *AIP Conference Proceedings*, 2009. arXiv:0812.3922.

[29] Frédéric Dupuis, Omar Fawzi, and Renato Renner. Entropy accumulation. *Communications in Mathematical Physics*, 379:867–913, 2020.

[30] Runyao Duan, Yuan Feng, Xin Li, and Mingsheng Ying. Multiple-copy entanglement transformation and entanglement catalysis. *Physical Review A*, 71:042319, 2005. arXiv:quant-ph/0404148.

[31] Naoto Shiraishi and Takahiro Sagawa. Quantum thermodynamics of correlated-catalytic state conversion at small scale. *Physical Review Letters*, 126:150502, 2021. arXiv:2010.11036.

[32] Michał Horodecki and Jonathan Oppenheim. Fundamental limitations for quantum and nanoscale thermodynamics. *Nature Communications*, 4:2059, 2013.

[33] Fernando G. S. L. Brandão, Michał Horodecki, Jonathan Oppenheim, Joseph M. Renes, and Robert W. Spekkens. Resource theory of quantum states out of thermal equilibrium. *Physical Review Letters*, 111:250404, 2013.