The price of keeping time

Seth Douglas Nidhal Mghirbi

October 2026

Abstract

A device that runs a quantum operation which does not commute with time translation needs its own clock: a phase reference that supplies the missing asymmetry. If the clock must come back exactly, and the battery holds no phase reference, no such operation runs at vanishing error; if it need only come back approximately, every channel runs in real time at the rate given by the time-covariant formula. For a gate, the clock is paid for in capital, in the energy range of the heat baths, and in an energy that comes back only as a lottery, each growing faster than the number of uses as the error vanishes. In batch, the work per use of every channel with a returned clock is a single-letter quantity: the largest energy-plus-entropy cost that two inputs can impose while leaving the channel's environment in the same state. It equals the thermodynamic capacity of Faist, Berta and Brandão whenever the environment can carry the channel's energy change, and can exceed it otherwise, even for channels of capacity zero. The same condition decides how large the clock's capital spread must be at fixed error: a spread growing like the square root of the number of uses suffices when the environment can carry the energy change, by a construction that does not attain the optimal work; otherwise the channel can read its own clock at the Heisenberg limit, and the spread must grow at least linearly. Counting its free energy, no time reference beats the returned clock in batch: one that comes back only in state, like an embezzling state, pays for any battery work it saves. Finally, in batch with a returned clock, a Gibbs-preserving qutrit channel costs $0.1547\,k_{\mathrm{B}}T$ per use while its thermodynamic time reverse costs nothing. Fixed full-rank thermal noise removes this asymmetry, but noise shrinking with the number of uses removes it only when the expected number of noisy uses diverges.

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 **6**

3 The dichotomy **9**

4 What the clock costs **9**

5 The clock's energy is a lottery **10**

6 Batch with a returned clock **12**

6.1 Inputs the environment cannot tell apart 12

6.2 The batch law 13

6.3 Examples 16

7 A process and its time reverse **18**

8 A controller that returns only in state **20**

9 Discussion **21**

How this work was done **22**

A Proofs **22**
A.1 The range of the baths . . . 22
A.2 What the clock costs . . . 23
A.3 The clock's energy . . . 24
A.4 A strong converse with a reusable controller . . . 25
A.5 The single-letter bound . . . 27
A.6 A returned clock reaches the capacity . . . 28
A.7 The batch law . . . 28
A.8 Unbalanced channels need a broad clock . . . 28
A.9 How large the clock must be . . . 29
A.10 The two families . . . 32
A.11 A channel that reaches its capacity . . . 32
A.12 A process and its reverse . . . 33
A.13 The noise horizon . . . 34
A.14 A controller that returns only in state . . . 34

B Physical construction for the batch law and the noise horizon **35**
B.1 The operational contract . . . 36
B.2 The lower bound with explicit budgets . . . 36
B.3 The device for a fixed channel . . . 37
B.3.1 Faithful auxiliary balance . . . 38
B.3.2 Uniform test and two pinchings . . . 38
B.3.3 Conserving position decoder . . . 39
B.3.4 Exact mechanical factory . . . 39
B.3.5 Rare excitation and a finite conserving bank . . . 40
B.3.6 Whole-controller and battery refunds . . . 41
B.4 Uniform noise construction . . . 42
B.5 Finite horizons . . . 43
B.5.1 Full-horizon padding and strict positivity . . . 43
B.5.2 Envelopes . . . 44

C Results from the companion papers **44**

1 Introduction

A device that must run a gate which does not commute with its Hamiltonian needs a clock, and if the clock is part of the device it must come back. This paper asks what owning such a clock costs: in work, in capital, and in the size of the clock.

Energy-conserving dynamics with heat baths can implement only operations that commute with time evolution, the time-covariant ones [1, 2]. Anything else, a Hadamard gate on a qubit with an energy gap for instance, needs a phase reference: a system whose own time evolution supplies the missing asymmetry. Åberg showed that such a reference can be reused without losing its power to enable coherent operations [3], and Korzekwa, Lostaglio, Oppenheim and Jennings showed that, repumped with part of the extracted work, it lets thermal machines extract work from coherence arbitrarily well, and accounted for the cost of the repumping [4]. Chiribella, Yang and

Renner, Tajima, Shiraishi and Saito, and Tajima, Yamaguchi, Takagi and Kuramochi showed that an approximate unitary needs a reference whose energy, coherence or free energy grows as the error shrinks [5, 6, 7]. Faist, Berta and Brandão introduced the thermodynamic capacity of a channel, the optimal work rate of a universal implementation of many uses when Gibbs-preserving maps are free, and reached it with thermal operations for time-covariant channels [8, 2]. They left open whether thermal operations can reach it for every channel, and expected this to need a very large amount of coherence, as embezzling states are needed in the quantum reverse Shannon theorem, because the environment must coherently compensate the energy difference between superposed inputs [8], [2, Sec. 9].

In a companion paper we priced the real-time implementation of channels, where each output must leave before the next input arrives [9]. For a clock carried by the device, it gives a sharp dichotomy (Section 3). A clock that must return exactly, with a battery that holds no phase reference of its own, cannot run any non-covariant channel: Marvian and Spekkens, and Lostaglio and Müller, showed that a finite reference used to implement an operation that is not covariant is necessarily degraded, as a consequence of the no-broadcasting theorem for asymmetry [10, 11], and [9] gives the resulting error bound for our devices. A clock that need only return approximately runs every channel at the real-time rate κ_β, the same formula as for covariant channels. So a clock adds no term to the operating work beyond κ_β, which can itself be positive. This paper asks what it costs instead.

Results. Throughout, n uses are implemented with error ε by a device whose controller returns approximately, with return error η. Three resource quantities appear below and should not be confused. *Capital* is the controller's initial free energy above its Gibbs state, $\beta^{-1}D(\alpha_C\|\gamma_C)$ with the relative entropy in natural logarithms; it is an energy. Its *spread* s_n, defined in Proposition 4, is dimensionless: a mean absolute deviation of the controller's initial log-likelihood ratio, which for a pure controller is β times the mean absolute deviation of its initial energy from a median. The *lottery* concerns the energy the controller exchanges during a run, not its initial spread.

What a clock costs. For a unitary U that does not commute with H:

(i) *Capital* (Proposition 2(a)). The controller's initial free energy above its Gibbs state is at least of order $n/(\varepsilon+\eta)$. The return error enters linearly because a clock that returns can be reused. At errors of order $1/\log n$ this is achievable up to a factor $\log n/\log\log n$.

(ii) *Baths and interactions* (Proposition 2(b)). The heat absorbed by the baths is nearly independent of the input, yet its mean must depend on it. So rare, large heat transfers are unavoidable: the baths' energy range, and the integrated strength of their interactions of bounded energy span, grow like $n/\sqrt{\varepsilon}$.

(iii) *The lottery* (Proposition 3). The controller's energy cannot come back branch by branch. On some input it moves by order n with probability at least about $1/2$, and if it returns on average, its standard deviation is of order $n/\sqrt{\varepsilon}$.

These hold in batch as well as in real time, and none is a work surcharge. Capital also decides reliability, through its spread rather than its size: a controller whose spread is $o(n/\sqrt{\eta})$, where η is the return error, as it is for linear capital or for a sharp energy state of any size, forces a strong converse at the capacity for every channel, while a bank of spread of order $n/\sqrt{\eta}$ breaks it for resets (Proposition 4); what such a bank lends is energy variance, not energy or entropy. In real time, the operating work of every channel is the covariant formula κ_β, while capital comes back and the baths and fluctuations are separate ledgers; in batch, a returned clock can cost more than the capacity

(Section 6). Channels that are not gates can need far smaller clocks: a Gibbs-preserving qubit channel runs in batch at zero work with a clock whose spread grows like $\sqrt{n}$ (Proposition 13).

The batch law. For every channel Φ:

(iv) *The law* (Theorem 7). With a returned clock, the optimal batch work per use is

$$B(\Phi) = \max\{\operatorname{Tr}\rho\,(\Phi^*(H) - H) + k_{\mathrm{B}}T\,[S(\sigma) - S(\Phi(\sigma))] \;:\; \Psi(\rho) = \Psi(\sigma)\},$$

with entropies in nats, where Ψ is a complementary channel. Two inputs can be charged separately, one for energy and one for entropy, exactly when the environment cannot tell them apart. The quantity is single-letter and additive, and for a channel given exactly it can be computed to any precision; it is not continuous in the channel.

(v) *When the capacity is reached* (Theorem 6). $B(\Phi)$ equals the thermodynamic capacity whenever the energy change $\Phi^*(H) - H$ lies in the span of the products $A_i^* A_j$ of Kraus operators: the environment can carry it. For Gibbs-preserving channels this is necessary and sufficient for zero cost, and every Gibbs-preserving qubit channel costs nothing. This answers a channel version of the question of Faist, Oppenheim and Renner whether thermal operations with a time reference can match Gibbs-preserving maps [12], for a reference that must return: on a qubit, including their own coherence-creating example, they always can; on a qutrit they can fail.

(vi) *When it is not* (Theorem 10, Propositions 11 and 12). For qubit dephasing in the X basis, on a qubit with $H = E|1\rangle\langle 1|$, batch and real time cost exactly the same, $E/2$ per use, above the capacity $\beta^{-1} \ln\cosh(\beta E/2)$. Two families of non-covariant channels are classified exactly. A qutrit channel that preserves the Gibbs state, and so has zero capacity, costs work of order n.

A process and its reverse.

(vii) (Proposition 14 and Proposition 15.) A Gibbs-preserving qutrit channel costs $c\,k_{\mathrm{B}}T$ per use in batch with a returned clock, $c = 0.1547103002\ldots$, while its Petz reverse, the reverse process of quantum fluctuation theorems, costs nothing. Which direction is free is decided by two linear tests on the Kraus operators. The asymmetry is not generic: channels of full Kraus rank pass both tests, so any fixed amount of full-rank thermal noise removes it, but for a finite run what matters is the expected number of replacements, $N\delta_N$ (Proposition 16). If it tends to zero, the original work rate remains; if it diverges, sublinear work is possible; if it tends to a positive finite value, the work rate stays positive, between two envelopes whose equality is open.

A controller that returns only in state.

(viii) (Corollary 17.) Counting the controller's free energy, no batch implementation of n uses, with any time reference, returned or not, costs less than $nB(\Phi) - o(n)$ on its most demanding input. A controller that comes back only in state, as an embezzling state does, can lower the battery's work to the capacity or below it, but only by paying the difference from its own free energy. To pay a fixed amount of work per use within return error η, its energy range plus $k_{\mathrm{B}}T$ times the logarithm of its dimension must be at least of order n/η, and a classical ladder of that size suffices. This answers the case of the question of Faist, Berta and Brandão that the batch law leaves aside, under both ways of counting: the battery's work alone has no meaningful optimum, and the total is $B(\Phi)$.

Prior work. Clocks that time energy-preserving dynamics, studied by Malabarba, Short and Kammerlander, Woods, Silva and Oppenheim, Erker et al. and Woods and Horodecki [13, 14, 15, 16], are a different object: they schedule operations that already commute with the total Hamiltonian, rather than supply asymmetry. Takagi and Shiraishi showed that catalysts which become correlated among themselves, each returned exactly as a marginal, let covariant operations perform any transformation of coherence [17]; such catalysts are not returned jointly, as our controller must be. The degradation of reference frames under repeated use has been modelled in [18, 19].

The capital bound uses the reuse argument of Chiribella, Yang and Renner, who bound the energy of the battery that powers a gate by running the gate and its inverse alternately on the same battery [5]. Errors alone already force capital of order $n/\sqrt{\varepsilon}$ in our error measure: Tajima, Yamaguchi, Takagi and Kuramochi bound the free energy of the resource directly, and Chiribella, Yang and Renner and Tajima, Shiraishi and Saito its energy and coherence [5, 6, 7]. What is new is that the return error enters as a separate parameter, linearly, so that a clock that returns more accurately than $\sqrt{\varepsilon}$ must hold more capital, and that the bound falls on the controller of a device whose work comes from a separate battery and whose heat goes to Gibbs reservoirs. The bound on the baths' range is an argument of Wigner–Araki–Yanase type, of the kind Chiribella, Yang and Renner use to bound the mean energy and the energy range of the battery that implements a gate [5]; what is new is that it falls on the heat baths, because the work is deterministic and the controller returns.

Work beyond the Landauer bound. Bedingham and Maroney showed that a device that must be correct on each of several non-orthogonal inputs, rather than only on their average, generally dissipates more heat than the Landauer bound of the average state [20]. The batch law is a counterpart of that effect for many uses and worst-case inputs: correctness on inputs that the environment cannot tell apart forces work beyond the capacity, even with a returned clock. The two differ in origin: their excess is entropic and occurs with degenerate Hamiltonians, while at zero Hamiltonian $B(\Phi)$ equals the capacity, and the excess in $B(\Phi)$ comes from energy the environment cannot carry. Tajima and Takagi showed that a Gibbs-preserving map with two orthogonal inputs whose outputs remain orthogonal needs coherence growing as the inverse of the error, whenever $H - \Phi^*(H)$ has a nonzero matrix element between the two inputs [21]; Tajima, Yamaguchi, Takagi and Kuramochi extended such divergences to free energy and to work drawn from a work bit [7]. For a channel on one system their condition implies ours: orthogonal outputs force $\langle\psi_1|A_i^* A_j|\psi_2\rangle = 0$ for all i, j, so $\Phi^*(H) - H$ lies outside the span of the $A_i^* A_j$.

Coherence and Gibbs-preserving maps. Faist, Oppenheim and Renner showed that Gibbs-preserving maps can create coherence between energy levels, which thermal operations cannot; they asked whether thermal operations combined with some form of time reference achieve the same transformations, and warned that a reference returned only approximately can embezzle work [12]. Brandão, Horodecki, Ng, Oppenheim and Wehner found that how nearly cyclic a process must be decides which second laws apply, and that in one regime work can be embezzled from a large system that stays arbitrarily close to its initial state [22]; Ng, Mančinska, Cîrstoiu, Eisert and Wehner found strong restrictions on such catalysis when the catalyst's dimension or mean energy is bounded [23]. Corollary 17 gives the exact ledger of this effect for implementations of channels. Faist's universal implementation with a variable work cost, optimal for each i.i.d. input, treats time-covariant channels, up to a decoherence that reveals the work consumed [24]; the batch law and its decoder extend the thermal-operation construction of [2, Sec. 7] from time-covariant channels to every channel with a returned clock, at the rate $B(\Phi)$. The span of the operators $A_i^* A_j$ that

controls our criterion is the non-commutative confusability graph of Duan, Severini and Winter [25].

Reversal. The Petz map is the reverse process of quantum fluctuation theorems [26, 27]. Aw, Zaw, Balanzó-Juandó and Scarani, and Song, Kwon and Scarani, ask when the Petz map of a channel can be implemented with the same devices as the channel itself, and find that exact reversal in general needs time-sensitive control of an ancilla [28, 29]; they do not price work or require a clock to return. The interaction and lottery bounds, the cheaper clock, the batch law and its criterion, the classification of Proposition 11 and the asymmetry in work between a process and its reverse appear to be new; the criterion contains the condition of Tajima and Takagi, and the asymmetry is the work counterpart of their asymmetry between detecting and creating coherence [21].

Outline. Section 2 recalls the model and Section 3 the dichotomy. Sections 4 and 5 price the clock. Section 6 gives the batch law and its examples, Section 7 the asymmetry between a process and its reverse, Section 8 what a controller that comes back only in state can buy, and Section 9 open problems. A companion paper [30] places these results among general schedules of arrivals, releases and loans. Proofs are in the appendices, as follows.

Result	Proof
Theorem 1	recalled from [9]
Proposition 2	Appendices A.2 (capital) and A.1 (baths)
Proposition 3	Appendix A.3
Proposition 4	Appendix A.4
Propositions 5 and 12	Appendix A.5
Theorem 6	Appendix A.6
Theorem 7	Appendix A.7; construction in Appendix B
Proposition 8	Appendix A.8
Proposition 9	Appendix A.9
Theorem 10	Section 6, from Theorem 7
Propositions 11 and 13	Appendices A.10 and A.11
Propositions 14 and 15	Appendix A.12
Proposition 16	Appendix A.13; construction in Appendix B
Corollary 17	Appendix A.14

2 Setting

We use the model of [9, Sec. 2]. There $\kappa_\beta(\Phi)$ denotes the optimal real-time work rate, in bits; its general formula is not needed here, and for a unitary U it equals $\beta\lambda_{\max}(U^*HU - H)/\ln 2$. A device consists of a controller C, a battery B and Gibbs reservoirs, all declared in advance; every operation is a unitary that conserves the total energy of the systems it acts on, and there is no external coherence. The controller holds all hardware, including the clock. The battery is *stationary* if it starts in a state that commutes with its Hamiltonian, as an energy eigenstate does. Throughout this paper the battery is stationary, so that every phase reference the device holds is in its controller: a battery in a broad superposition of energy levels would itself be a clock, which the device could spend without returning it [9]. Errors lie in $[0, 1]$: the error ε is the complete strategy distance to n ideal uses, the largest half trace distance over adaptive observers who keep quantum references; in batch it is the half diamond distance $\frac{1}{2}\|\widetilde{\Phi}_n - \Phi^{\otimes n}\|_\diamond$ of the implemented channel $\widetilde{\Phi}_n$, with $U^{\otimes n}$ for a gate. Work is the mean energy the battery loses. It is *deterministic* when the battery starts in an energy eigenstate, ends in a fixed eigenstate with probability tending to one, and its mean energy differs from that eigenvalue by $o(n)$. The last clause matters when the controller returns

only approximately. Without it, a broad reference could lend the device its energy and be refilled from the battery in one rare, large event, leaving the battery's typical endpoint untouched.

A controller *returns approximately* if, uniformly over observers, its final joint state with all other systems is within $\eta \to 0$ in trace distance of α_C tensored with their final state, and its mean energy and its entropy change by $o(n)$. Trace-distance closeness alone would not suffice. As Faist, Oppenheim and Renner warned, a catalyst returned only approximately can embezzle work [12]: here a broad reference could hand over energy of order n while staying close to its initial state, and would then act as a second battery. Section 8 gives the exact free-energy ledger of such a controller. The controller may start with any declared amount of phase reference. Entropies are in nats when they appear as $k_{\mathrm{B}}T\,S$ and in bits when they appear as $k_{\mathrm{B}}T\ln 2\,S$. Write $D = U^{*}HU - H$, $\mathrm{span}(X)$ for the difference between the largest and smallest eigenvalues of X, and g for the spread of the diagonal of D in an energy eigenbasis. If U does not commute with H, then $g > 0$ (Appendix A). The thermodynamic capacity $\mathcal{T}(\Phi)$ is an energy here; it equals $k_{\mathrm{B}}T\ln 2$ times the capacity in bits used in the companion papers [9, 30]. The batch cost $B(\Phi)$ is also an energy, and $\beta B(\Phi)/\ln 2$ is its rate in bits. Time evolutions are written with $\hbar = 1$; the interaction strength of Proposition 2 restores it. Table 1 lists the conditions that each group of results uses.

Results	Controller starts	Controller returns	Battery	Error and conclusion
Batch law (Theorems 6 and 7, Proposition 5)	in any declared state; its free energy is invested capital	the whole controller, jointly with everything else, with actual energy and entropy changes $o(n)$	starts in an energy eigenstate; ends with entropy $o(n)$ and mean energy within $o(n)$ of its nominal endpoint	error tending to zero against entangled inputs; optimal work $B(\Phi)$ per use, the capacity when D lies in the span of the $A_i^* A_j$
Clock spread (Propositions 4, 8 and 9)	in any state; the spread s_n of its capital is what is bounded	not required at fixed error; for vanishing error, independent of everything else within η_n, uniformly over inputs	starts in a state commuting with its Hamiltonian; any work, or work below the capacity in Proposition 4	below the spreads of Table 2, error tending to one
Gate resources (Propositions 2 and 3)	in any state	approximately, within η, for the capital and bath bounds; the lottery needs no return, and its variance bound only a mean refund	sharp for the capital and bath bounds; in the lottery every system but the controller starts stationary	error ε; capital of order $n/(\varepsilon+\eta)$, bath range and interaction strength of order $n/\sqrt{\varepsilon}$, controller energy moving by order n
Free-energy ledger (Corollary 17)	in any declared state	no condition for (a); for (b), its own state within η_n of α_C on every input, with no refund of energy or entropy; in (c), jointly within $\eta_n + o(1)$, entropy refunded, energy not	as for the batch law	error tending to zero; battery and controller together supply at least $nB(\Phi) - o(n)$ on the most demanding input

Table 1: The contracts of the four groups of results; all hold in batch as well as in real time, except the batch law and the free-energy ledger, which are batch statements. Every error is a complete strategy distance, a half diamond distance in batch, against observers that keep quantum references. Startup capital is declared and is not part of the work. Constructions are built for a known number n of uses, and each channel has the same Hamiltonian at input and output.

3 The dichotomy

Theorem 1 (Exact and approximate return [9]; (a) after [10, 11]). *(a) If the controller returns exactly and the battery is stationary, every implementation of a channel Φ has error at least $\frac{1}{4}\sup_t \|J_\Phi - e^{-itG} J_\Phi e^{itG}\|_1$, with J_Φ the Choi state and $G = H_O - H_I^T$; this is positive exactly when Φ is not time-covariant.*

(b) If the controller returns approximately, the optimal real-time work per use at vanishing error is $k_B T \ln 2\, \kappa_\beta(\Phi)$ for every channel Φ. For a unitary it is exactly $\lambda_{\max}(D)$, with the battery reaching its target on every branch.

The construction behind (b) rests on one primitive. With $w = \lambda_{\max}(D)$ and $B = wI - D \geq 0$, a thermal qubit with a large gap M absorbs the input-dependent energy $-D$ by a rare excitation: a rotation between payload and qubit by the angle whose sine is $\sqrt{B/M}$ changes the payload channel by only $O(\|B\|/M)$, yet takes up the energy B on average. The constant w is supplied by a sharp battery, and lifting the whole contact with a phase reference, as in the construction of Chiribella, Yang and Renner [5] after Åberg [3], makes it conserve energy exactly. For the Hadamard gate, a reference of two sine-shaped packets, the reference state of [5], with widths $(n+1)^8$ and $(n+1)^4$ gives work exactly $nE/\sqrt{2}$ with errors $O(n^{-3})$; the rest of this paper shows that such large resources are not an accident of the construction.

4 What the clock costs

The operating work is the formula of Theorem 1, but the clock and the baths are not cheap.

Proposition 2 (What the clock costs). *Let U be a unitary that does not commute with H, let $D = U^*HU - H$, and let $g > 0$ be the spread of the diagonal of D in an energy eigenbasis. Consider any implementation of n uses whose controller returns approximately, with error ε and return error η, $\varepsilon + \eta \leq 1/16$, Gibbs reservoirs and a sharp battery.*

(a) Capital. *For all large n, the controller's initial free energy above its Gibbs state satisfies*

$$\beta^{-1} D(\alpha_C \| \gamma_C) \;\geq\; \frac{g\, n}{2048(\varepsilon+\eta)} - \frac{\ln 2}{\beta};$$

for the Hadamard gate, $nE/(256(\varepsilon+\eta)) - \ln 2/\beta$ for every n.

(b) Baths. *On the two products of extreme eigenvectors of D, the laws of the heat Q absorbed by the baths are within $2\sqrt{2\varepsilon}$ in total variation, while their means differ by $L_n = n(\lambda_{\max} - \lambda_{\min})(D) - 2n\varepsilon\, \mathrm{span}(H) - o(n)$. So the baths' total energy range is at least $L_n/(2\sqrt{2\varepsilon})$. If every interaction that changes the baths' energy acts on cells of total energy span at most q_0, there are at least $L_n/(4q_0\sqrt{2\varepsilon})$ such interactions, and, if $\varepsilon \leq 1/(8e^2)$, their integrated strength $A_n = \int \sum_l \|h_l(t)\|\, dt/\hbar$ satisfies*

$$A_n \;\geq\; \frac{L_n/(8q_0\sqrt{2\varepsilon}) - \ln(1/(2\sqrt{2\varepsilon}))}{2(\sqrt{e}-1)},$$

which grows like $n/\sqrt{\varepsilon}$.

Both parts hold for batch implementations too.

The ideas are simple (Appendix A). For (a), an observer feeds K uses half of a GHZ state whose partner carries compensating energy, so that everything except the controller starts stationary. The ideal outputs then carry coherence $1/2$ at frequency KE, and only the controller can supply it. As in the argument of Chiribella, Yang and Renner [5], a clock that returns approximately can be reused, here for about $1/(8(\varepsilon+\eta))$ blocks of n uses, so K grows as $n/(\varepsilon+\eta)$. For (b), accuracy forces the baths' final state to be nearly independent of the input, while energy conservation forces their mean heat to depend on it. A bounded-span interaction of total strength A gives the heat exponential tails of scale A, which forces A to be large.

The capital bound is nearly tight at slowly vanishing error. For the Hadamard gate there are implementations with work exactly $nE/\sqrt{2}$, errors $O(1/\log n)$ and capital $O(E\,n\log^2 n/\log\log n)$, and the rare high-energy qubits can be replaced by blocks of ordinary qubits of gap E, flipped together (Appendix A). Errors alone force capital of order $n/\sqrt{\varepsilon}$, by the energy, coherence and free-energy cost laws of Chiribella, Yang and Renner, of Tajima, Shiraishi and Saito, and of Tajima, Yamaguchi, Takagi and Kuramochi [5, 6, 7], applied to n uses. Chiribella, Yang and Renner show that for a fixed register a battery recycled through a circuit needs energy independent of the number of gates; here each use acts on a fresh input, so the register, and with it the capital, grows with n. In (a) the return error enters linearly because a clock that returns can be reused; in our constructions η is of order $\sqrt{\varepsilon}$, and the two scalings agree. In real time, asymmetry is paid in invested free energy, which comes back, in the range of the baths and in interaction effort, never in operating work; in batch, Section 6 shows that it can also cost operating work.

5 The clock's energy is a lottery

In the construction of [9], for any input other than the top eigenvector of D, the controller typically ends with energy higher by order n; rare branches in which a high-gap thermal qubit absorbs a large quantum balance this on average. This is not an accident of the construction.

Proposition 3 (The clock's energy is a lottery)**.** *Let U be a unitary that is not time-covariant, let $d_i = \langle i|(U^*HU - H)|i\rangle$ in an energy eigenbasis of H, and let $g = \max_i d_i - \min_i d_i$; then $g > 0$. Take any energy-conserving implementation of n uses of U with error ε in which every system other than the controller starts in a state that commutes with its Hamiltonian. Feed $|i\rangle^{\otimes n}$, with i maximizing or minimizing d_i, and let q be the energy gained by the controller, read off from energy measurements of everything else at the start and at the end. Let $p_{\max}$ and $p_{\min}$ be the probabilities that $|q| > an$ on the two feeds. Then for every $0 < a < g/2$,*

$$p_{\max} + p_{\min} \;\geq\; 1 - 2\sqrt{2\varepsilon} - 2\varepsilon - 4e^{-n(g-2a)^2/(2\,\mathrm{span}(H)^2)}.$$

If in addition the controller's mean energy changes by at most u on both feeds, then on one of them the standard deviation of q is at least

$$\frac{(ng - 2n\varepsilon\,\mathrm{span}(H) - 2u)\sqrt{1-16\varepsilon}}{8\sqrt{\varepsilon}} - \Big(\frac{\sqrt{n}}{2} + \sqrt{\varepsilon}\,n\Big)\mathrm{span}(H).$$

The measurement that defines q never touches the controller. Everything else starts in a stationary state, so measuring its energy first changes nothing. Each outcome then shifts the controller's energy by exactly q, whatever its coherence, and the first two moments of q are those of the operator $V^*H_CV - H_C$, where V is the whole circuit. The proof combines two facts (Appendix A). Accurate service makes the final state of the heat baths and the battery nearly independent of the input, while energy conservation requires the input-dependent part of the energy, of size about ng,

to go somewhere. If the controller's energy change were small on both feeds, two nearly identical heat distributions would have to concentrate on intervals about ng apart. So no accurate clock returns its energy branch by branch. One that returns it on average does so through a lottery, in which rare, large transfers balance a typical change of order n.

The bounds are sharp in order. The construction of [9] has a standard deviation of order $n/\sqrt{\varepsilon}$, and the value $1/2$ for the larger of the two probabilities cannot be improved; only the size of q is forced, not its sign. On a coherent input, the variance of $V^*H_CV - H_C$ is at least $n^2[\mathrm{span}(U^*HU - H) - 4\varepsilon\,\mathrm{span}(H)]^2/4$, with no condition on its mean. Chiribella, Yang and Renner bound the clock's initial mean energy, energy range and energy spread, and Tajima, Shiraishi and Saito its initial coherence [5, 6]; Proposition 3 bounds its energy change.

Capital and reliability. Capital also decides whether a device that saves work must fail. With an approximately returning controller, a broad energy bank can make loans that depend on the branch and cancel in the mean, and so run a covariant reset at half the vanishing-error rate with error exactly $1/2$, even with a sharp battery. Only a broad bank can do this: what matters is the spread of the controller's capital, not its size.

Proposition 4 (Strong converse with a reusable controller)**.** *Let Φ be a channel on a finite-dimensional system with Hamiltonian H and thermodynamic capacity $\mathcal{T}(\Phi) = \max_\rho[F(\Phi(\rho)) - F(\rho)]$, where $F(\rho) = \mathrm{Tr}\,H\rho - \beta^{-1}S(\rho)$ with S in nats, and let ρ be a maximizer. Consider implementations of n uses of Φ, in batch or in real time, in which the battery starts in an energy eigenstate $|i\rangle$ and ends in a fixed energy eigenstate $|f\rangle$ except with probability q_n, with $E_i - E_f \le n\,(\mathcal{T}(\Phi) - \Delta)$ for a fixed $\Delta > 0$. Let the controller start in $\alpha_C = \sum_a \lambda_a |a\rangle\langle a|$, with Gibbs state $\gamma_C = \sum_b \mu_b |b\rangle\langle b|$, and call*

$$s_n = \min_c \sum_{a,b} \lambda_a |\langle a|b\rangle|^2 \left|\ln \frac{\lambda_a}{\mu_b} - c\right|$$

the spread of its capital. There are $b > 0$ and $f_n \to 0$, depending only on Φ, β and Δ, such that the error ε_n satisfies:

(a) $1 - \varepsilon_n - q_n \le e^{D_{\max}(\alpha_C\|\gamma_C) - bn} + f_n$;

(b) if on the input $\rho^{\otimes n}$ the controller returns independently of everything else, $\frac{1}{2}\|\Omega_{C,\mathrm{rest}} - \alpha_C \otimes \Omega_{\mathrm{rest}}\|_1 \le \eta_n$, then

$$1 - \varepsilon_n - q_n \;\le\; 51\sqrt{\eta_n}\,\max\Big\{1, \Big(\frac{8s_n}{\beta\Delta n}\Big)^2\Big\} + f_n;$$

(c) under the return condition of (b), also

$$1 - \varepsilon_n - q_n \;\le\; C\,\eta_n\,\max\Big\{1, \Big(\frac{8s_n}{\beta\Delta n}\Big)^2\Big\} + f_n,$$

with $C = 2^{17}$, and $C = 36$ if α_C commutes with H_C and the maximizer ρ commutes with H. A maximizer that commutes with H exists whenever Φ is time-covariant; for other channels the bound with $C = 36$ holds with $\mathcal{T}(\Phi)$ replaced by the largest value of $[D(\Phi(\rho)\|\gamma) - D(\rho\|\gamma)]/\beta$ over states ρ that commute with H.

In particular, if $q_n \to 0$, $\eta_n \to 0$ and $s_n = o(n/\sqrt{\eta_n})$, then $\varepsilon_n \to 1$. No time covariance is assumed.

Part (a) needs no return: capital below bn cannot help, which extends the converses for a controller that starts in its Gibbs state [9]. Parts (b) and (c) are the reusable case, and they charge

the spread of the capital, not its size. Mean capital of order n is enough, since $s_n \le D(\alpha_C\|\gamma_C)+2/e$, and so is every controller with $\ln\dim C = O(n)$ and an energy span of order n. A controller in an energy eigenstate has $s_n = 0$ however much energy it holds, and for a pure controller s_n is β times the mean absolute deviation of its energy from a median. The threshold in the last sentence of the proposition is sharp up to constants: a flat energy bank of spread about $0.159\,n\eta_n^{-1/2}$ runs n resets at half the vanishing-error rate with error exactly $1/2$, while its mean energy and entropy return exactly (Appendix A.4). What the bank lends is energy variance, not energy or entropy. Approximately returned catalysts are usually controlled through the product of their return error and their size, $\eta_n K$ with $K = \beta\,\mathrm{span}(H_C) + \ln\dim C$ [22, 23]. For reliability this is not enough: for the flat bank $\eta_n K$ tends to zero when its number of levels grows faster than n, and the quantity that decides reliability is $\eta_n s_n^2/n^2$. The proof uses the spread only to keep the characteristic function of the controller's log-likelihood law at least $1/2$ in modulus near zero, so what matters is how concentrated that law is: a controller that puts weight at least $3/4$ on a single value of the log-likelihood is covered whatever its spread. Reliability can be broken only by a broad bank, but breadth alone is not enough. The converse is at the capacity; Proposition 8 extends it to the batch cost $B(\Phi)$ of Section 6. For a gate, Proposition 2 shows that vanishing error already needs capital of order $n/(\varepsilon+\eta)$, more than linear, so the proposition bears on covariant channels and on noisy channels whose clocks are cheap, such as the channel of Proposition 13 at slowly vanishing error. The proof is in Appendix A.

6 Batch with a returned clock

A batch device may hold all its outputs until the end and erase its baths using them as side information; this is how Faist, Berta and Brandão reach the thermodynamic capacity of time-covariant channels with thermal operations [2, Sec. 7]; with Gibbs-preserving maps they reach it for every channel [2, Sec. 6]. For channels that are not time-covariant, whether thermal operations reach the capacity is their open question [8], [2, Sec. 9], and Faist, Oppenheim and Renner asked more generally whether thermal operations with a time reference can match Gibbs-preserving maps [12]. For a clock that must return, this section answers both questions at the level of work per use in batch: the batch cost of every channel is a single-letter quantity, which equals the capacity whenever the environment can carry the channel's energy change, and for a Gibbs-preserving channel only then. The expectation of Faist, Berta and Brandão that a very large amount of coherence is needed is borne out for gates (Proposition 2); and when the coherence must be returned, even unlimited coherence does not reach the capacity of some channels (Theorem 10). A reference that is not returned, or that comes back only in state, such as an embezzling state, can lower the battery's share of the work, even below the capacity, but not the total: it pays the difference from its own free energy (Section 8).

6.1 Inputs the environment cannot tell apart

If two inputs ρ and σ of a channel Φ leave the same state in its environment, then every implementation of n uses whose controller returns approximately spends at least

$$n\Big[\mathrm{Tr}\,H\big(\Phi(\rho)-\rho\big) + k_{\mathrm{B}}T\ln 2\,\big(S(\sigma)-S(\Phi(\sigma))\big)\Big] - o(n),$$

with entropies in bits. The idea is this. The inert purifier of the device's entire initial internal state does not depend on the input. If the environment cannot tell ρ from σ, accurate service forces the outputs on ρ to look, to that purifier, like the outputs on σ. Continuity of conditional entropy on

the outputs then forbids the device's internal entropy from falling by more than σ allows, so the energy that ρ requires cannot be drawn from the baths' entropy beyond that amount and must come from the battery. Nothing depends on the dimension of the controller or on the tails of the heat. For a unitary every input leaves the environment unchanged, and the bound is the capacity.

The best bound of this kind is single-letter. Write Ψ for a complementary channel of Φ, $D = \Phi^*(H) - H$, and

$$B(\Phi) = \max\{\operatorname{Tr}\rho D + k_{\mathrm{B}}T\,[S(\sigma) - S(\Phi(\sigma))] \;:\; \Psi(\rho) = \Psi(\sigma)\},$$

with entropies in nats. For example, for qubit dephasing in the X basis with $H = E|1\rangle\langle 1|$, the ground state and the maximally mixed state leave the environment in the same state. The first requires the channel to supply energy $E/2$; the second offers no entropy decrease to pay for it. So $B(\Phi_X) \geq E/2$, above the capacity $\beta^{-1}\ln\cosh(\beta E/2)$, and Theorem 10 shows that this is the cost.

Proposition 5 (A single-letter batch bound)**.** *Consider implementations of n uses whose controller returns approximately, with its actual energy and entropy changing by $o(n)$, and whose battery starts in an energy eigenstate and ends with entropy $o(n)$.*

(a) Every implementation spends at least $nB(\Phi) - o(n)$, and $B(\Phi^{\otimes N}) = NB(\Phi)$: neither entangled inputs nor pairs whose environment states differ by a vanishing amount give more.

(b) $B(\Phi)$ is at least the thermodynamic capacity, with equality whenever D lies in the range of Ψ^*, in particular for every covariant channel and every channel whose complementary channel is injective.

(c) $B(\Phi) = 0$ if and only if $\Phi(\gamma) = \gamma$ and D lies in the range of Ψ^*. For a qubit this holds exactly when Φ is Gibbs-preserving.

The range of Ψ^* is the span of the products $A_i^* A_j$ of Kraus operators, and D lies in it exactly when some dilation conserves mean energy for a suitable environment Hamiltonian: the environment can carry the energy change, input by input on average. We call such a channel *balanced*; balance concerns the mean energy and does not make the channel time-covariant. Covariant channels always satisfy this. The criterion makes precise the obstruction that Faist, Berta and Brandão anticipated: for superposed inputs of different energy, the environment must compensate the energy difference without disturbing the process [8].

6.2 The batch law

Theorem 6 (A returned clock reaches the capacity)**.** *Suppose $\Phi^*(H) + \Psi^*(K) = H$ for some Hermitian K on the environment of a Stinespring dilation of Φ; equivalently, D lies in the span of the products $A_i^* A_j$. Then, with a returned clock, the optimal batch work per use is the thermodynamic capacity, attained with deterministic work. In particular, a Gibbs-preserving channel costs no work per use in batch with a returned clock if and only if $H - \Phi^*(H)$ lies in the span of the $A_i^* A_j$, and every Gibbs-preserving qubit channel costs none. For time-covariant channels this is [2, Thm. 7.1].*

The construction prepares the environment in a zero-energy eigenstate of K, runs the dilation with the clock absorbing the energy of each branch, which on average is zero for every input, and then erases the environment using the held outputs as side information, recovering the preparation cost up to the capacity. A returned clock can supply phase, but it cannot provide net extensive energy: for every observer its actual mean energy change must be sublinear, although it can still make large loans that depend on the branch, as long as their mean is refunded. So thermal operations with a

returned time reference implement a Gibbs-preserving channel at zero work per use in batch exactly when its energy change lies in the span of its Kraus products.

When the environment cannot carry the whole energy change, the remainder must be paid, and the bound of Proposition 5 says exactly how much.

Theorem 7 (The batch law)**.** *For every channel Φ on a finite-dimensional system with Hamiltonian H at input and output, and every $0 < T < \infty$, consider batch implementations of n uses that hold all outputs until cleanup, with independently supplied Gibbs reservoirs, a controller in any declared initial state whose free energy is counted as invested capital, a battery that starts in an energy eigenstate, and finite energy-conserving gates. Require that the error, against inputs entangled across uses and with a reference, tend to zero, that the whole controller, not only the clock, return jointly with everything else, with actual energy and entropy changes $o(n)$, and that the battery end with entropy $o(n)$ and with mean energy within $o(n)$ of its nominal endpoint. The optimal work per use, the drop of the battery's nominal energy, is $B(\Phi)$. A single family attains it with a common nominal battery endpoint, reached except with vanishing probability, and with battery entropy and mean absolute deviation from that endpoint $o(n)$; its reservoirs, capital and running time grow with n. For time-covariant channels, where $B(\Phi)$ is the capacity, this is [2, Thm. 7.1].*

The dual form of B shows how. For every Hermitian Y on the environment,

$$B(\Phi) \le \lambda_{\max}\big(D - \Psi^*(Y)\big) + \max_{\sigma}\big[k_{\mathrm{B}}T\big(S(\sigma) - S(\Phi(\sigma))\big) + \operatorname{Tr}\sigma\Psi^*(Y)\big],$$

with one σ serving both terms of the inner maximum, and the infimum over Y is $B(\Phi)$. Write t_Y and g_Y for the two terms on the right, so that $B(\Phi) = \inf_Y(t_Y + g_Y)$. Each Y is a device. Give the environment the Hamiltonian $\lambda_{\max}(Y)I - Y$ and prepare it in its ground state; the environment then carries the part $\Psi^*(Y)$ of the energy change. Pay the part it cannot carry, $D - \Psi^*(Y)$, at its top eigenvalue by rare excitations of a thermal qubit of large gap; run the dilation with a returned clock; and erase the environment using the held outputs as side information. The erasure step uses an energy-conserving decoder. Every input leaves the environment with conditional free energy $\operatorname{Tr}\Psi(\rho)H_E - k_{\mathrm{B}}T[S(\rho) - S(\Phi(\rho))] \ge \lambda_{\max}(Y) - g_Y$, directly from the definition of g_Y. So one energy-conserving test, invariant under the energy and accepting every code state uniformly in operator norm, separates the code from a thermal environment at that rate; with Gibbs-weighted rejection and the thermal-operation form of position-based decoding [31, 2] it erases the environment, and an exact mechanical factory converts the resulting information label into work. When D lies in the range of Ψ^*, the rare excitations are unnecessary and the device is that of Theorem 6. The complete construction, with its physical ledger, is in Appendix B. It uses large capital and reservoirs; it shows what is possible in principle, not an efficient engine.

How large the clock must be. The balanced case is also the case of small clocks.

Proposition 8 (Unbalanced channels need a broad clock)**.** *Let Φ be a channel whose energy change $D = \Phi^*(H) - H$ lies outside the span of the products $A_i^*A_j$ of its Kraus operators. Consider implementations of n uses of Φ, at any work, whose battery, like the Gibbs reservoirs, starts in a state that commutes with its Hamiltonian, and let s_n be the spread of the controller's capital, as in Proposition 4. If $s_n = o(n)$, the error tends to one. If the controller returns independently of everything else, uniformly over all inputs and references, with return error η_n, the error also tends to one whenever $\eta_n \max\{1, (s_n/n)^2\} \to 0$; and if $\eta_n > 0$, there are $b, c, N > 0$ depending only on Φ and H such that $s_n \ge c\,\beta n/\sqrt{\eta_n(\varepsilon_n + \eta_n)}$ whenever $n \ge N$ and $\varepsilon_n + \eta_n \le b$. Consequently, for every channel, nominal work below $n[B(\Phi) - \Delta]$ with a fixed $\Delta > 0$ and the battery as in Proposition 4, together with $q_n \to 0$, $\eta_n \to 0$ and $s_n = o(n/\sqrt{\eta_n})$, forces the error to one.*

The second statement combines the first with Proposition 4, since $B(\Phi)$ is the capacity when D lies in the span. The reason is metrological. The generator $i\sum_i A_i^* \dot{A}_i$ of the channel's orbit under time evolution is exactly D, and when it lies outside the span of the $A_i^* A_j$, that orbit can be estimated at the Heisenberg limit, with phase resolution of order $1/n$ rather than the $1/\sqrt{n}$ of the standard quantum limit, using error correction [32, 33, 34]: inputs that the environment cannot tell apart differ in the energy the channel adds, $\operatorname{Tr}\rho D$, and a code built from them turns n uses into a probe of phase with resolution $1/n$. Related metrological arguments limit codes that are covariant under a continuous symmetry [35, 36]. Rotating the controller rotates the implemented channel along that orbit, so a device that implements Φ reliably reads its own clock to within $1/n$. A clock of spread $o(n)$ cannot be read that finely at all, and a clock that returns within η_n needs spread of order $n/\sqrt{\eta_n}$ even at fixed error. To make the error vanish it must be broader still: passing one logical qubit through about $1/(\varepsilon_n + \eta_n)$ consecutive runs before reading its phase multiplies the signal while the noise grows only at second order, so every reliable implementation has $s_n\sqrt{\eta_n}/n \to \infty$, with spread at least of order n/η_n when $\varepsilon_n = O(\eta_n)$. Since $s_n \le D(\alpha_C\|\gamma_C) + 2/e$, the same bound holds for the capital; for a gate it improves Proposition 2(a) whenever $\eta_n = o(\varepsilon_n)$. When D lies in the span there is no such probe, and the clock can be small: the channel of Proposition 13 needs spread only of order $\sqrt{n}$ at fixed error. The balanced criterion of Theorem 6 is therefore also the line between clocks that may be small and clocks that must be read at the Heisenberg limit. X-dephasing lies on the far side: every reliable implementation, at any work, needs a clock whose spread grows faster than $n/\sqrt{\eta_n}$. The proof is in Appendix A.

The exponent of the clock. The span criterion also sets how the clock must grow with n. Proposition 8 gives the linear lower bound for unbalanced channels; the next proposition gives the rest. Its parts (a) and (b) use only the time orbit of the channel and the standard quantum limit, part (c) a lift whose errors cancel between uses, and part (d) the code of Proposition 8 together with the no-programming theorem [37, 38].

Proposition 9 (How large the clock must be)**.** *Let Φ be a channel that is not time-covariant, consider implementations whose battery, like the Gibbs reservoirs, starts in a state that commutes with its Hamiltonian, and let s_n, ε_n and η_n be as in Proposition 8. There are constants $c, C, N > 0$, depending only on Φ, H and β, such that the following hold, in batch and in real time.*

(a) At any work, and without any return of the controller, $s_n = o(\sqrt{n})$ forces $\varepsilon_n \to 1$.

(b) If the controller returns independently of everything else, uniformly over all inputs and references, then $s_n \ge c\sqrt{n/(\varepsilon_n + \eta_n)}$ whenever $n \ge N$ and $\varepsilon_n + \eta_n \le 1/32$.

(c) If D lies in the span of the $A_i^* A_j$, then for every $a \in (0, 1]$ and every n there is a batch implementation of n uses with error at most $a/2$, whose whole controller returns jointly with everything else within $2a$, with actual changes $o(1)$ of its mean energy and $O(\log n)$ of its entropy, whose battery starts in an energy eigenstate and ends in a fixed one except with probability at most $1/n$, and with $s_n \le C\sqrt{n}/a$.

(d) For X-dephasing, $\Phi_X(\rho) = (\rho + X\rho X)/2$ with $H = E|1\rangle\langle 1|$: at any work and without return, $s_n \ge \beta E n/(512\sqrt{\varepsilon_n})$ whenever $n \ge 4096$ and $0 < \varepsilon_n \le 1/4096$. Conversely, for all $\varepsilon, \eta \in (0, 1]$ there is a real-time implementation with error below ε, return error below η, actual changes $o(n)$ of the controller's mean energy and entropy, work exactly $E/2$ per use on every branch, and $s_n = O(\beta E n/\min\{\eta, \sqrt{\varepsilon}\})$.

Table 2 collects the bounds. For X-dephasing with $\varepsilon_n = O(\eta_n)$ they match: the optimal spread is of order $\beta E n \max\{\eta_n^{-1}, \varepsilon_n^{-1/2}\}$, and it is reached at the optimal work $E/2$ per use. When $\varepsilon_n \ll \eta_n^2$ the bound of (d) decides it, and when $\eta_n^2 \lesssim \varepsilon_n \lesssim \eta_n$ the reliable bound of Proposition 8; when $\eta_n \ll \varepsilon_n$ the two sides differ by a factor $\sqrt{\varepsilon_n/\eta_n}$. The construction in (c) spends work of order $n \log(n/a)$, which is not the batch rate, and for irrational gaps its controller is astronomically large, although its spread is not. Without return, vanishing error need not cost a balanced channel any extra spread: the replacement channel $\rho \mapsto |+\rangle\langle+|$, which is balanced, is run exactly by swapping in a stock of $|+\rangle$ states, whose spread is about $\beta E\sqrt{n/(2\pi)}$. For X-dephasing, which is unbalanced, it must: by Proposition 8 and (d), its spread is at least of order n, and of order $n/\sqrt{\varepsilon_n}$ as the error vanishes.

Channel	Fixed error	Vanishing error, controller returned
Time-covariant	no clock	no clock
Balanced, not covariant	$s_n \asymp \sqrt{n}$	$\sqrt{n/(\varepsilon_n+\eta_n)} \lesssim s_n \lesssim \sqrt{n}\,\max\{\eta_n^{-1}, \varepsilon_n^{-1}\}$
Unbalanced	$s_n \gtrsim n$	$s_n \gtrsim n/\sqrt{\eta_n(\varepsilon_n+\eta_n)}$
X-dephasing	$s_n \asymp n$	$s_n \asymp n \max\{\eta_n^{-1}, \varepsilon_n^{-1/2}\}$ if $\varepsilon_n = O(\eta_n)$

Table 2: The spread of the clock needed for n uses of a fixed channel (Propositions 8 and 9), with constants depending on Φ, H and β. At fixed error $0 < \varepsilon < 1$ the lower bounds are strong converses and need no return. The balanced upper bound is a batch construction whose work is of order $n \log(n/a)$, not the batch rate; the X-dephasing upper bound runs in real time at the optimal work. All lower bounds assume a battery that starts in a state commuting with its Hamiltonian.

6.3 Examples

Theorem 10 (For X-dephasing, batch costs as much as real time)**.** *Let $\Phi_X(\rho) = (\rho + X\rho X)/2$ on a qubit with $H = E|1\rangle\langle 1|$; its thermodynamic capacity is $\mathcal{T}(\Phi_X) = \beta^{-1} \ln\cosh(\beta E/2) < E/2$. Among implementations of n uses at vanishing error whose controller returns approximately, with its actual energy and entropy changing by $o(n)$, the optimal work per use is $E/2$, in batch and in real time alike. The lower bound allows any controller, however much entropy it holds, reservoirs of any spectrum, any amount of phase reference, and cleanup assisted by the outputs; it uses only that the controller's free energy comes back up to $o(n)$ and that the battery's entropy changes by $o(n)$, so it covers catalysts returned only as marginals too. The upper bound is attained by a finite construction with deterministic work.*

Here $B(\Phi_X) = E/2$, with $\rho = |0\rangle\langle 0|$ and $\sigma = I/2$. The batch cost therefore exceeds the capacity by $\beta^{-1} \ln[2/(1+e^{-\beta E})]$, and waiting saves nothing. On the ground-state input, Φ_X must deliver the maximally mixed state while preserving the eigenstates of X on the other inputs, and the environment cannot tell the ground state from the maximally mixed state.

The necessary condition settles two families of non-covariant channels exactly. In the first, σ is the input with its coherences removed; in the second, it is the input with its discarded part decorrelated.

Proposition 11 (Which channels reach the capacity in batch)**.** *Consider implementations as in Theorem 10.*

(a) Phase noise. *If Φ fixes every state of an orthonormal basis, every implementation of n uses spends at least $nw - o(n)$, where $w = \lambda_{\max}(\Phi^*(H) - H)$. If $\Phi = \sum_j p_j V_j(\cdot)V_j^*$ with commuting*

unitaries V_j and $p_j > 0$, then w is the optimal work per use, in batch and in real time, and Φ reaches its thermodynamic capacity if and only if the top eigenspace of $\Phi^*(H) - H$ contains a common eigenvector of all $V_j^* V_k$. In particular, a projective pinching reaches its capacity if and only if it is time-covariant, while a Hadamard gate on one qubit tensored with energy dephasing on another, with additive H, reaches it although it is neither covariant nor unitary.

(b) A protected subsystem. *Let $\Phi = \mathrm{id}_A \otimes R_B$, where R_B replaces B by its maximally mixed state, and let H act on AB. Every implementation spends at least $nL - o(n)$, where*

$$L = \max_{\rho} \Big\{ \mathrm{Tr}\big[(\Phi^*(H) - H)\rho\big] + k_{\mathrm{B}} T \big[S(\rho_B) - \ln d_B\big] \Big\},$$

with entropies in nats, and L is also the optimal real-time rate of [9]. The capacity equals L if and only if $H = H_A \otimes I + I \otimes H_B$, equivalently if and only if Φ is time-covariant.

For instance, let $H = E(I - |\Phi^+\rangle\langle\Phi^+|)$ on two qubits and $\Phi = \mathrm{id}_A \otimes R_B$. Every output is mixed. The batch and real-time work is $3E/4$ per use at every temperature, attained by a finite construction with deterministic work, while the capacity is $-E/4 + \beta^{-1} \ln[(e^{\beta E} + 3)/4]$, about $0.072E$ at $\beta E = \ln 2$. So the obstruction does not come from pure outputs. In both families, the capacity is reached exactly when an energetically optimal input gives the environment nothing beyond what its decohered or decorrelated version gives: a clock that must return can lend asymmetry, but it cannot recover information the channel throws away. The gaps are strict at every finite temperature and close as $T \to 0$.

Proposition 12 (Gibbs preservation is not enough)**.** *On a qutrit with $H = \mathrm{diag}(0, E_1, E_2)$, let $K_1 = x|0\rangle\langle 1| + y|1\rangle\langle 0|$ and $K_0 = \mathrm{diag}(\sqrt{1-y^2}, \sqrt{1-x^2}, 1)$ with $x^2 = y^2 e^{\beta E_1} \le 1$. This channel is Gibbs-preserving, so its capacity is zero, yet with a returned clock every implementation spends at least $n y^2 E_1 - o(n)$.*

The environment learns that a jump happened but not its direction. So the ground state, which must be excited, looks to the environment exactly like a stationary mixture with the third level, from which no entropy can be drawn. Gibbs preservation, which makes a channel free in the framework of [2], therefore does not make it free for a device whose clock must return. Tajima and Takagi showed that some Gibbs-preserving maps need unbounded coherence, and Tajima, Yamaguchi, Takagi and Kuramochi that some need free energy and work diverging as the error vanishes when the resource is consumed [21, 7], which leaves the corresponding work cost open. Here the coherence is unlimited and returned, and the cost is a work rate that does not depend on the error.

Channels that only create coherence, like the example of Faist, Oppenheim and Renner [12], are the opposite case, and they show that a returned clock can be cheap.

Proposition 13 (A Gibbs-preserving channel that reaches its capacity)**.** *Let $H = E|1\rangle\langle 1|$ with $\beta E = \ln 3$, so that $\gamma = \mathrm{diag}(3/4, 1/4)$, and let $\Phi(\rho) = \langle 0|\rho|0\rangle\sigma_0 + \langle 1|\rho|1\rangle\sigma_1$ with*

$$\sigma_0 = \begin{pmatrix} 3/4 & \sqrt{3}/12 \\ \sqrt{3}/12 & 1/4 \end{pmatrix}, \qquad \sigma_1 = \begin{pmatrix} 3/4 & -\sqrt{3}/4 \\ -\sqrt{3}/4 & 1/4 \end{pmatrix}.$$

This channel, the map of Faist, Oppenheim and Renner [12] for a pure target with the populations of γ, preserves γ, is not time-covariant, and has thermodynamic capacity zero. With a returned clock its optimal batch work per use is zero, attained with deterministic work. In the implementation, the clock's energy changes by a sum of n independent increments in $\{-E, 0, E\}$ with mean zero and variance $3E^2/8$ each, whatever the input, and a clock in a uniform superposition of $L \ge 2\sqrt{3n/8}/\varepsilon^2$ consecutive levels, with n spare levels on each side, returns within ε plus the error of the rest of the construction, jointly with everything else.

The channel reads only populations and writes coherence: Φ ignores the coherences of its input, and the populations of its output are γ whatever the input. Tajima and Takagi found that maps of this coherence-creating kind have finite coherence cost [21]. So the clock lends coherence but never has to absorb an energy that depends on the input, and its energy spreads like $\sqrt{n}$ rather than n. The costs of Propositions 2 and 3 are therefore those of gates; noisy channels can need far less.

None of this contradicts the theorem of Faist, Berta and Brandão. For channels that are not time-covariant they reach the capacity with Gibbs-preserving maps, which are free by assumption and need not be realizable by energy-conserving interactions with Gibbs reservoirs [8, 12, 21]; with thermal operations they treat time-covariant channels, for which $B(\Phi)$ equals the capacity. Nor does it contradict their thermal-operation implementation of any channel on a fixed, known i.i.d. input [2, Prop. 8.3], which consumes a coherence source of sublinear size. They treat that source as disposable, but it in fact returns approximately, with energy and entropy changes $o(n)$, so it meets our return condition, and their conclusion that Gibbs-preserving maps and thermal operations become asymptotically equivalent on given i.i.d. inputs [8] stands. What the batch law prices is a device that must be correct on every input and return its controller's energy and entropy; for such devices the equivalence fails (Proposition 12).

7 A process and its time reverse

For a Gibbs-preserving channel Φ, the reverse process of quantum fluctuation theorems is its Petz map [26, 27]

$$\Phi^\sharp(\sigma) = \gamma^{1/2}\Phi^*(\gamma^{-1/2}\sigma\gamma^{-1/2})\gamma^{1/2},$$

which also preserves γ. Thermal operations are closed under this reversal: the reverse of $\mathrm{Tr}_B[U(\cdot \otimes \gamma_B)U^*]$ is $\mathrm{Tr}_B[U^*(\cdot \otimes \gamma_B)U]$ [26, 39], and Aw, Zaw, Balanzó-Juandó and Scarani showed that the Petz map of any channel is obtained by reversing a dilation once system–bath correlations are accounted for [28]. Once a clock must be returned, they need not be.

Proposition 14 (A process and its reverse can cost differently)**.**
On a qutrit with $\gamma = \mathrm{diag}(1/2, 1/3, 1/6)$ *and* $H = k_\mathrm{B}T\,\mathrm{diag}(0, \ln\frac{3}{2}, \ln 3)$*, let* $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$*,* $P_\pm = |\pm\rangle\langle\pm|$*,* $P_2 = |2\rangle\langle 2|$*, and*

$$\mathcal{N}(\rho) = \sum_{i=0}^{2} \mathrm{Tr}(M_i\rho)\,|i\rangle\langle i|, \qquad M_0 = P_+ + \tfrac{1}{2}P_2, \quad M_1 = \tfrac{4}{5}P_-, \quad M_2 = \tfrac{1}{5}P_- + \tfrac{1}{2}P_2.$$

Both $\mathcal{N}$ *and its Petz reverse* $\mathcal{N}^\sharp$ *preserve* γ *and have capacity zero. With a returned clock,* $\mathcal{N}^\sharp$ *costs no work per use in batch, while* $\mathcal{N}$ *costs exactly* $c\,k_\mathrm{B}T$ *per use, where* $c = 0.15471030022\ldots$ *is the value of a strictly concave three-variable maximization, given explicitly in terms of the unique positive solution of two equations, in Appendix A.12.*

The reverse reads only populations, so its energy change is diagonal and lies in the span of its Kraus products. The forward channel reads the populations of $|\pm\rangle$; its Kraus products have no matrix element between $|+\rangle$ and $|-\rangle$, while H does, so inputs that differ only in their $|+\rangle\langle-|$ coherence are indistinguishable to the environment but differ in energy. The value is $B(\mathcal{N})$, reached by Theorem 7.

In coherence the same pair is already asymmetric. The inputs $|+\rangle$ and $|-\rangle$ of $\mathcal{N}$ have orthogonal outputs and $\langle+|H - \mathcal{N}^*(H)|-\rangle = -\frac{1}{2}k_\mathrm{B}T\ln\frac{3}{2} \neq 0$, so by the theorem of Tajima and Takagi an implementation of $\mathcal{N}$ by thermal operations needs coherence that diverges as the error vanishes, while $\mathcal{N}^\sharp$, which reads energy populations and prepares coherent states, is of the coherence-creating

kind that they found to have finite coherence cost [21]. Proposition 14 shows that this asymmetry between detecting and creating coherence survives in work when the coherence is supplied by a clock that must return, and that the two channels are each other's thermodynamic time reverse.

Which direction is free is decided by linear algebra. Write $\Phi(X) = \sum_i A_i X A_i^*$. The Petz reverse has Kraus operators $\gamma^{1/2} A_i^* \gamma^{-1/2}$, so the criterion of Theorem 6 for $\Phi^\sharp$ can be written in terms of Φ alone:

$$\text{forward free} \iff H - \Phi^*(H) \in \operatorname{span}\{A_i^* A_j\}, \qquad \text{reverse free} \iff \gamma H - \Phi(\gamma H) \in \operatorname{span}\{A_i \gamma A_j^*\}.$$

Both, one or neither of the two tests can hold.

Proposition 15 (The asymmetry is not generic)**.** *If a Gibbs-preserving channel Φ has a Choi matrix of full rank, then Φ and $\Phi^\sharp$ both cost nothing in batch with a returned clock. In particular, for every Gibbs-preserving Φ and every $0 < \delta \le 1$, the channel $(1-\delta)\Phi + \delta\, \gamma \operatorname{Tr}$ and its reverse both cost nothing. Separated pairs lie on the boundary of the Gibbs-preserving channels, in continuous families: replacing $|\pm\rangle$ in Proposition 14 by $\cos\theta|0\rangle + \sin\theta|1\rangle$ and $-\sin\theta|0\rangle + \cos\theta|1\rangle$, with $M_0 = P_{u_\theta} + zP_2$, $M_1 = \frac{2}{2+z} P_{v_\theta}$, $M_2 = \frac{z}{2+z} P_{v_\theta} + (1-z)P_2$ and $z = \sin^2\theta$, gives a separated pair for every $0 < \theta < \pi/2$.*

So the cost of the forward channel jumps from $c\, k_{\mathrm{B}} T$ to zero under any fixed amount of full-rank thermal noise. Tajima and Takagi found the same geometry for coherence: every Gibbs-preserving map that needs unbounded coherence lies on the boundary of the Gibbs-preserving channels, and every map in their relative interior has finite coherence cost [21]. For a finite device, what matters is the expected number of noisy uses.

Proposition 16 (Noise horizon)**.** *Let Φ be Gibbs-preserving with $B(\Phi) > 0$, and let $\Phi_\delta = (1-\delta)\Phi + \delta\, \gamma \operatorname{Tr}$. Call a family of batch implementations of $\Phi_{\delta_N}^{\otimes N}$ admissible if its error tends to zero, its controller returns approximately with actual energy and entropy changes $o(N)$, and its battery starts in an energy eigenstate and ends with entropy and mean deviation from its nominal endpoint $o(N)$; write W_N for its nominal work. Although Φ_δ costs nothing for every fixed $\delta > 0$, the noise helps only through the expected number $N\delta_N$ of replacements:*

(a) if $N\delta_N \to 0$, every admissible family has $W_N \ge NB(\Phi) - o(N)$, and some has $W_N \le NB(\Phi) + o(N)$;

(b) some admissible family has $W_N = o(N)$ if and only if $N\delta_N \to \infty$, and then one has $W_N = O\big(N(N\delta_N)^{-1/4} + \log N\big)$;

(c) for $\lambda \in (0, \infty)$, let $L(\lambda)$ and $U(\lambda)$ be the infima of $\liminf_N W_N/N$ and of $\limsup_N W_N/N$ over admissible families with $N\delta_N \to \lambda$; they do not depend on the sequence. Then $0 < L(\lambda) \le U(\lambda) \le B(\Phi)$; both are nonincreasing, $\lambda L(\lambda)$ and $\lambda U(\lambda)$ are nondecreasing, and both tend to $B(\Phi)$ as $\lambda \to 0$.

So the asymmetry of Proposition 14 is invisible to fixed noise but not to a finite device whose noise per use is small compared with the inverse of the number of uses. Tajima and Takagi found a related sensitivity, in the coherence some Gibbs-preserving maps need as the error vanishes [21]; here it is work, and the noise rather than the error sets the threshold. The positivity of $L(\lambda)$ is qualitative: its constant can be very small. The proof is in Appendix B. The asymmetry is an implementation asymmetry created by the obligation to return the clock, not a violation of any fluctuation theorem: those compare a protocol with its reversed protocol, not the cheapest devices for a channel and for its reverse. Nor is it carried by the clock, whose entropy changes by $o(n)$. On the equilibrium input $\gamma^{\otimes n}$, which both channels fix, the work of a deterministic implementation

equals the correlations it leaves plus the departure of its reservoirs from equilibrium, up to $o(n)$; with ideal reservoirs this is ordinary entropy production in the baths. For qubits the asymmetry cannot occur, since every Gibbs-preserving qubit channel costs nothing; likewise, the construction of Tajima and Takagi of cost-diverging maps on a single system needs at least three distinct energies [21].

8 A controller that returns only in state

The batch law requires the controller to come back jointly with everything else and to refund its actual energy and entropy up to $o(n)$. The refund is not a technicality: a broad reference could hand over energy of order n while staying close to its initial state (Section 2). We say that a controller *comes back only in state* if its final state is close to its initial one in trace distance, with no condition on its energy or entropy, and we ask what such a controller buys. Write

$$F_C(\varrho) = \operatorname{Tr} H_C\, \varrho - k_{\mathrm{B}} T\, S(\varrho)$$

for the free energy of a state ϱ of the controller, with the entropy in nats, and R_n for the largest drop $F_C(\alpha_C) - F_C(\alpha'_C)$ over inputs, where α'_C is the controller's final state on that input. As in Section 5, the size of the controller is $K = \beta \operatorname{span}(H_C) + \ln d_C$, where $\operatorname{span}(H_C)$ is its energy range and d_C its dimension.

Corollary 17 (A controller returned only in state is a second battery)**.** *Consider batch implementations of n uses of a channel Φ with error tending to zero, independently supplied Gibbs reservoirs, finite energy-conserving gates, a controller in any declared initial state, and a battery that starts in an energy eigenstate and ends with entropy $o(n)$ and with mean energy within $o(n)$ of its nominal endpoint. Write W_n for the nominal work.*

(a) Whatever the controller does, $W_n + R_n \geq nB(\Phi) - o(n)$.

(b) If, on every input, the final state of the controller is within $\eta_n > 0$ of α_C in trace distance, then

$$R_n \leq k_{\mathrm{B}} T\, [\eta_n K + r_0(\eta_n)], \qquad r_0(x) = (1+x)\ln(1+x) - x \ln x.$$

So if $W_n \leq n[B(\Phi) - \Delta]$ with a fixed $\Delta > 0$, the size of the controller is at least $K \geq [\beta n \Delta - r_0(\eta_n) - o(n)]/\eta_n$.

(c) For every real $w \leq B(\Phi)$, every energy $\omega > 0$ and every sequence $\eta_n \in (0, 1]$, some family meets the conditions of Theorem 7 on the error and the battery and has nominal work $nw + o(n)$; its whole controller returns jointly with everything else within $\eta_n + o(1)$ in trace distance, with entropy change $o(n)$, but its free energy drops by $n[B(\Phi) - w] + o(n)$ on every input. The part that pays is a classical ladder of spacing ω and energy range at most $(1+1/\eta_n)\,(n[B(\Phi)-w]+\omega)$.

Part (a) is Proposition 19 with the controller's actual change of free energy kept rather than bounded. Part (b) is the usual control of an approximately returned catalyst by the product of its return error and its size [22, 23], applied to that change. Part (c) adds to the device of Theorem 7 a ladder in a uniform mixture of N levels, which pays the battery in advance by one energy-conserving swap and moves by k/N in trace distance when it is lowered by k levels. Applied to the ladder alone, the bound of (b) is then met up to a factor $1 + \eta_n$ and terms of lower order, so it cannot be improved by more than that factor, and it is sharp as $\eta_n \to 0$. The payment needs no coherence; a pure ladder in a sine-shaped superposition also works, at a range larger by a factor that approaches π as $\eta_n \to 0$. The proof is in Appendix A.14.

This answers, under both ways of counting, the case left aside in Section 6. A controller that comes back only in state, whether it holds a coherent reference, an embezzling state or a plain store of energy, lets the battery's work reach the capacity, or any smaller value, negative ones included, so the battery's work alone has no meaningful optimum. What it cannot do is lower the total. Counting its free energy, no batch implementation, with any time reference, returned or not, draws less than $nB(\Phi) - o(n)$ from battery and controller together on its most demanding input, and the returned clock of Theorem 7 attains this. So $B(\Phi)$ is the batch cost of Φ with every time reference; it exceeds the capacity for X-dephasing (Theorem 10) and for the Gibbs-preserving qutrit channel of Proposition 12. A controller that comes back only in state is a second battery, and hiding its payment behind a small return error takes size: to pay a fixed amount of work per use within return error η, it needs a size K at least of order n/η.

9 Discussion

In real time, a lack of time-translation symmetry adds nothing to the operating rate but is expensive to own; in batch, a returned clock can add to the operating rate (Section 6). A clock that returns approximately runs every channel at the covariant rate, but for a gate its capital, the range of its baths, the effort of interacting with them and the spread of its own energy all grow faster than the number of uses as the error vanishes. In batch, a returned clock makes every channel cost $B(\Phi)$: two inputs can be charged separately for energy and for entropy exactly when the environment cannot tell them apart. No other time reference does better: a controller that comes back only in state can pay part of the work, but only from its own free energy, and to pay a fixed amount per use within return error η it needs a size K at least of order n/η (Section 8). The capacity of Faist, Berta and Brandão [8, 2] is reached in the balanced case, where the environment can carry the energy change; otherwise waiting can save less, or nothing, as for X-dephasing. And on the boundary of the Gibbs-preserving channels, a process and its time reverse can cost different amounts of work.

Open problems. Our constructions use large capital and astronomically large reservoirs. What capital and what reservoirs does the batch law need at error ε, and how much does finite size cost? How large can the gap between a channel and its Petz reverse be in dimension d? Our constructions give a gap growing like $k_{\mathrm{B}}T \ln d$, without a sharp constant. Proposition 3 concerns the energy the controller exchanges with everything else; can the controller's own energy-change operator nonetheless concentrate? At errors n^{-a} the capital bound and our construction leave a gap of $n^{a/2}$; in every construction we know, the return error is of order nE/Q for a clock of energy capital Q, and we conjecture that this holds for every architecture. The small-spread strong-converse threshold is sharp up to constants: part (c) of Proposition 4 covers every reusable controller whose spread is $o(n/\sqrt{\eta_n})$, while the flat bank reaches that scale and runs resets at half the vanishing-error rate with error $1/2$. This does not make spread alone a sufficient condition for reliable implementation. By Proposition 8 the same threshold holds at the batch cost, and reliable implementations of unbalanced channels need spread at least of order $n/\sqrt{\eta_n(\varepsilon_n + \eta_n)}$; the optimal trade-off between spread, return error and service error is open. In the noise horizon, whether the two envelopes $L(\lambda)$ and $U(\lambda)$ coincide, so that a single crossover curve exists, is open. Finally, our constructions use nonlocal gates and external scheduling; autonomous clocks with these properties are open.

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/price-of-keeping-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 constant chains of the converses and of the clock-size proposition, the sinc and no-programming inequalities, the energy-conserving permutations of the constructions, the swap and shift identities of the paying ladder, and the exact identities of the worked examples, including the dual certificate of the qutrit pair and the flat bank. 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 Proofs

A.1 The range of the baths

Suppose an implementation of n uses of a unitary U has error ε. For orthogonal inputs a, b, let σ_a, σ_b be the final states of everything except the outputs. Feeding the superposition $(|0\rangle a + |1\rangle b)/\sqrt{2}$ with a kept qubit, and using that the ideal outputs are pure, the device's state must nearly factor off, which gives $d_{\mathrm{tr}}(\sigma_a, \sigma_b) \le 2\sqrt{2\varepsilon}$ [40]. Take a and b to be n-fold products of eigenvectors of D for its largest and smallest eigenvalues. The delivered energies differ by at least $n(\lambda_{\max} - \lambda_{\min})(D) - 2n\varepsilon \operatorname{span}(H)$. The work is input-independent and the controller's mean energy changes by $o(n)$ in both runs, so by energy conservation the mean energies of the heat baths differ by the same amount up to $o(n)$. For any bath Hamiltonian of total span R, $|\operatorname{Tr} H_{\mathrm{baths}}(\sigma_a - \sigma_b)| \le R\, d_{\mathrm{tr}}(\sigma_a, \sigma_b)$. Hence

$$R \;\ge\; \frac{n(\lambda_{\max} - \lambda_{\min})(D) - 2n\varepsilon \operatorname{span}(H) - o(n)}{2\sqrt{2\varepsilon}},$$

which grows faster than n as $\varepsilon \to 0$ whenever U is not time-covariant. Rare, large heat transfers are therefore unavoidable, as in the construction.

A.2 What the clock costs

Capital. Let Λ be a time-covariant channel from the controller C to a system Y, L a contraction on Y with $[H_Y, L] = -qL$ for some $q > 0$, and $A = \Lambda^*(L)$. Then $A^*A \le I$, $[H_C, A] = -qA$, and $\operatorname{Tr}\gamma_C A^*A = e^{-\beta q}\operatorname{Tr}\gamma_C AA^* \le e^{-\beta q}$. Measuring the effect A^*A and using $\operatorname{Tr}\alpha_C A^*A \ge |\operatorname{Tr}\alpha_C A|^2$,

$$D_{\mathrm{nat}}(\alpha_C\|\gamma_C) \;\ge\; \beta q\,|\operatorname{Tr}\alpha_C A|^2 - \ln 2,$$

with no dependence on any dimension. For the Hadamard gate, let an observer feed K uses the state $(|0\rangle^{\otimes K}|q_0\rangle + |1\rangle^{\otimes K}|q_1\rangle)/\sqrt{2}$, where a kept qubit has energies 0 and $-KE$, so that the whole tester is stationary. The ideal output is $(|+\rangle^{\otimes K}|q_0\rangle + |-\rangle^{\otimes K}|q_1\rangle)/\sqrt{2}$, and $L = Z^{\otimes K}\otimes|q_1\rangle\langle q_0|$ lowers energy by KE and has expectation $1/2$ there. If the actual output is within δ, then $|\operatorname{Tr}\rho L| \ge 1/2 - 2\delta$. Everything except the controller starts stationary, so the whole experiment is a covariant channel from C, and the lemma gives $\beta^{-1}D(\alpha_C\|\gamma_C) \ge KE(1/2-2\delta)^2 - \ln 2/\beta$.

A controller that returns approximately can serve m consecutive blocks of n uses, with fresh reservoirs for each block, at total error at most $m\varepsilon + (m-1)\eta$. To see this, compare with a mathematical hybrid in which the controller is replaced by a fresh copy of α_C at block boundaries. Each replacement costs at most η by approximate product return, and the hybrid with all replacements runs m independent blocks, each within ε of ideal. Take $m = \lfloor 1/(8(\varepsilon+\eta))\rfloor$ and $K = mn$; then $\delta \le 1/8$ and (a) follows for the Hadamard gate. *General unitaries.* The lemma above needs only that every matrix element of L between energy eigenspaces lowers the energy by at least q. Then $A = \Lambda^*(L)$ has the same property, $A^*A \le I$, and detailed balance still gives $\operatorname{Tr}\gamma_C A^*A \le e^{-\beta q}\operatorname{Tr}\gamma_C AA^* \le e^{-\beta q}$. For a general U, pick energy eigenvectors $|a\rangle, |b\rangle$, of energies h_a, h_b, whose diagonal values of D are extreme, $d_b - d_a = g$. The observer feeds K uses the state $(|a\rangle^{\otimes K}|q_a\rangle + |b\rangle^{\otimes K}|q_b\rangle)/\sqrt{2}$, where the kept states $|q_a\rangle, |q_b\rangle$ are orthogonal, with energies $-Kh_a$ and $-Kh_b$, so that the tester is stationary. The ideal output is $(\psi_a + \psi_b)/\sqrt{2}$ with $\psi_x = (U|x\rangle)^{\otimes K}|q_x\rangle$. On ψ_x the total energy is the kept energy plus a sum of K independent payload energies, each of range $\operatorname{span}(H)$, and its mean is Kd_x. Let P_x project, within the label q_x, onto total energies in $[Kd_x - Kg/4,\, Kd_x + Kg/4]$. Hoeffding's inequality gives $p_x = \|P_x\psi_x\|^2 \ge 1 - h_K$ with $h_K = 2e^{-Kg^2/(8\operatorname{span}(H)^2)}$. With $u_x = P_x\psi_x/\sqrt{p_x}$, the contraction $L = |u_a\rangle\langle u_b|$ lowers energy by at least $Kg/2$, because the bottom of the window of b lies $Kg/2$ above the top of the window of a, and its expectation on the ideal output is $\sqrt{p_a p_b}/2 \ge (1-h_K)/2$. The lemma gives

$$\beta^{-1}D(\alpha_C\|\gamma_C) \;\ge\; \frac{Kg}{2}\Big[\frac{1-h_K}{2} - 2\delta\Big]^2 - \frac{\ln 2}{\beta}.$$

With the same reuse, $K = mn$ and $\delta \le 1/8$; once $h_n \le 1/4$, which holds for all large n since $h_K \le h_n$, the bracket is at least $1/8$, and $K \ge n/(16(\varepsilon+\eta))$ gives the constant $g/2048$.

Baths. Record the baths' initial energies in an inert purification and let Q be the heat they absorb. For the two product inputs of (b), the decoupling argument above makes the two laws of Q within $\delta = 2\sqrt{2\varepsilon}$, while energy conservation, the deterministic work and the controller's mean return make their means differ by L_n. If the laws are supported in an interval of length R, then $L_n \le R\delta$, which gives the range bound; and a circuit of m interactions, each changing the baths' energy by at most q_0, has $R \le 2q_0 m$. For the integrated strength, an interaction built from terms of span at most q_0 and total strength A gives the exponential moments $\mathbb{E}\,e^{\pm Q/(2q_0)} \le e^{2(\sqrt{e}-1)A}$, by a Dyson expansion in the style of [41]. Write $M = e^{2(\sqrt{e}-1)A_n}$, and write the difference of the two laws P, P' of Q as $\mu_+ - \mu_-$ with $\mu_+ \le P$, $\mu_- \le P'$ and common mass $\delta' \le \delta$. Jensen's inequality for μ_+/δ' gives $\int Q\,d\mu_+ \le 2q_0\delta'\ln(\mathbb{E}_P e^{Q/(2q_0)}/\delta') \le 2q_0\delta'\ln(M/\delta')$, and in the same way

$-\int Q\,d\mu_- \le 2q_0\delta'\ln(M/\delta')$. Since $x \mapsto x\ln(M/x)$ increases for $x \le M/e$, and $\delta \le 1/e \le M/e$ when $\varepsilon \le 1/(8e^2)$, the means differ by at most $4q_0\delta\ln(M/\delta)$, so $L_n \le 4q_0\delta\big[2(\sqrt{e}-1)A_n + \ln(1/\delta)\big]$, which is the bound on A_n. All three bounds use only the two product inputs and energy conservation, so they hold in batch as well.

A cheaper clock. In the construction of [9] for the Hadamard gate, the reference's widths were chosen against the worst case, in which every one of the n contacts emits rare heat. Given the past, a contact emits it with probability at most b/M, so the number J of emissions has $\mathbb{E}J^2 \le nb/M + (nb/M)^2$, and J exceeds a cutoff K_c only with probability at most $(nb/M)^{K_c+1}/(K_c+1)!$. Factoring the phase of each emission into an orthogonal record of the bath, the reference's phase sensitivity becomes $O(n+\sqrt{nm})$ instead of $O(nm)$, and the reference's box needs a buffer of only $n + m(K_c+2) + 2$, with the rare histories beyond the cutoff handled by a coherent comparison with the ideal circuit and no guard. With errors $O(1/\log n)$ the capital is then $\Theta(E\,n\log^2 n/\log\log n)$; with errors $O(n^{-a})$ it is $\Theta(E\,n^{1+3a/2})$. The work is still exactly $nE/\sqrt{2}$, and the controller returns its energy and entropy up to $o(1)$ on average. Finally, a qubit of large gap M can be replaced by a block of $k = \lceil M/(E\tanh(\beta E/2))\rceil$ Gibbs qubits of gap E flipped together by $iX^{\otimes k}$: a flip always moves a block to an orthogonal state, so the heat effect and the disturbance of the payload are as before.

A.3 The clock's energy

Proof of Proposition 3. First, $g > 0$: in an energy eigenbasis with energies h_i, $\sum_{i,j}|U_{ji}|^2(h_j-h_i)^2 = -2\sum_i h_i d_i$, and if every d_i vanished this would force $[U,H]=0$.

Let R be everything except the controller and the payload outputs, including the heat baths and the battery. Since R and the input $|i\rangle^{\otimes n}$ start in states that commute with their Hamiltonians, measuring their total energy at the start does not disturb anything, and comparing it with a final measurement defines a positive instrument on the controller whose outcomes shift its energy by definite amounts q. By energy conservation, q is minus the change of energy of R and the payload. Its first two moments are those of $V^*H_CV - H_C$ on these inputs; its higher moments in general are not.

Accurate service makes the final states of R on the two feeds within $\delta = 2\sqrt{2\varepsilon}$ in trace distance (the range of the baths, above). On the feed $|i\rangle^{\otimes n}$ the payload gains energy nd_i up to $n\varepsilon\,\mathrm{span}(H)$, and the sum of the n independent output energies concentrates around its mean, with Hoeffding tails $e^{-n(g-2a)^2/(2\,\mathrm{span}(H)^2)}$. If $|q| \le an$ with high probability on both feeds, the energy change of R would concentrate on two windows centred about ng apart, each of width less than ng. Two distributions within δ in total variation cannot do this, which gives the sum bound. For the standard deviation, the comparison is made in Hellinger distance; this is where the factor $\varepsilon^{-1/2}$ enters. Feed the two product inputs in superposition, marked by a kept qubit. By Uhlmann's theorem there is a pure state e of everything except the outputs such that the actual final state has fidelity at least $1-\varepsilon$ with the ideal output tensored with e, so on each feed its branch has overlap at least $1-2\varepsilon$ with that product. Let ℓ_x be the energy change of R on feed x, read off from the same energy measurements, and p_x its law. Both laws have Bhattacharyya coefficient at least $1-2\varepsilon$ with the law of the same measurement on e, so $\|\sqrt{p_a} - \sqrt{p_b}\|_2 \le 4\sqrt{\varepsilon}$. For every function f, the Cauchy–Schwarz inequality gives the classical comparison

$$|\mathbb{E}_a f - \mathbb{E}_b f| \;\le\; \|\sqrt{p_a} - \sqrt{p_b}\|_2 \,\big(2\,\mathbb{E}_a f^2 + 2\,\mathbb{E}_b f^2\big)^{1/2}.$$

Take $f = \ell$ minus the midpoint of its two means, and let Δ be the distance between the means.

Then $\mathbb{E}_x f^2 = \operatorname{Var}_x \ell + \Delta^2/4$, so $\Delta^2 \le 16\varepsilon(\Delta^2 + 2\operatorname{Var}_a \ell + 2\operatorname{Var}_b \ell)$, and

$$\max_x \operatorname{Var}_x \ell \;\ge\; \frac{\Delta^2(1-16\varepsilon)}{64\varepsilon}.$$

By energy conservation $\ell_x = -q_x - \Delta S_x$, where ΔS_x is the energy gained by the payload. Its means on the two feeds differ by at least $ng - 2n\varepsilon \operatorname{span}(H)$, while those of q differ by at most $2u$, so $\Delta \ge ng - 2n\varepsilon \operatorname{span}(H) - 2u$. Finally $\sigma(q_x) \ge \sigma(\ell_x) - \sigma(\Delta S_x)$, and $\sigma(\Delta S_x) \le (\sqrt{n}/2 + \sqrt{\varepsilon}\, n) \operatorname{span}(H)$: the ideal output energies are n independent variables, each with standard deviation at most $\operatorname{span}(H)/2$, and the actual law differs from the ideal one with probability at most ε, on an interval of length $n \operatorname{span}(H)$. This gives the stated bound. The variance bound on a coherent input follows from the uncertainty relation between $V^* H_C V - H_C$ and an output observable that detects the relative phase of the ideal output. □

A.4 A strong converse with a reusable controller

Proof of Proposition 4. Write $\sigma = \Phi(\rho)$, $D_{\text{in}} = D(\rho\|\gamma)$ and $D_{\text{out}} = D(\sigma\|\gamma)$, so that $\beta\mathcal{T}(\Phi) = D_{\text{out}} - D_{\text{in}}$, and let $W = E_i - E_f$.

(a) Choose a smoothed input $\tau_n \le e^{n(D_{\text{in}}+\kappa)}\gamma^{\otimes n}$ with $\|\tau_n - \rho^{\otimes n}\|_1 \le f_n$, and an output test P_n with $\operatorname{Tr} P_n \sigma^{\otimes n} \ge 1 - f_n$ and $\operatorname{Tr} P_n \gamma^{\otimes n} \le e^{-n(D_{\text{out}}-\kappa)}$. Let $\mathcal{M}$ be the success instrument on the controller: run the device on τ_n, apply $P_n \otimes |f\rangle\langle f|$, and trace out everything except C. It is completely positive. Because the dilation conserves energy, it fixes the global Gibbs state, so $\mathcal{M}(\gamma_C) \le r\,\gamma_C$ with $r = e^{\beta W - n(\beta\mathcal{T}(\Phi) - 2\kappa)} \le e^{-bn}$ for $\kappa = \beta\Delta/8$. Since $\alpha_C \le e^{D_{\max}(\alpha_C\|\gamma_C)}\gamma_C$, $\operatorname{Tr}\mathcal{M}(\alpha_C) \le e^{D_{\max}(\alpha_C\|\gamma_C) - bn}$, while $\operatorname{Tr}\mathcal{M}(\alpha_C) \ge 1 - \varepsilon_n - q_n - 2f_n$.

(b) For a density $X = \sum_a x_a |a\rangle\langle a|$ and a positive definite $\Gamma = \sum_b g_b |b\rangle\langle b|$, the weights $x_a |\langle a|b\rangle|^2$ form a probability law for $L = \ln(x_a/g_b)$, whose characteristic function is $\Xi(X;\Gamma,t) = \operatorname{Tr} X^{1+it}\Gamma^{-it}$ [42]. Its mean is $D(X\|\Gamma)$ when Γ is normalized. Two properties are used.

Acceptance. If a projection A commutes with Γ, then for every real x,

$$\operatorname{Tr} AX \le e^x \operatorname{Tr} A\Gamma + \Pr(L \ge x). \tag{1}$$

Indeed, in a joint eigenbasis of A and Γ, $\operatorname{Tr} AX = \sum_{b\in A}\sum_a x_a |\langle a|b\rangle|^2$, and the terms with $L < x$ have $x_a < e^x g_b$.

Continuity. For densities A, B and every positive definite Γ,

$$|\Xi(A;\Gamma,t) - \Xi(B;\Gamma,t)| \le 2\sqrt{1+t^2}\,\|A - B\|_1^{1/2}.$$

Put $z = (1+it)/2$. The map $u \mapsto u^{1+it}$ is $\sqrt{1+t^2}$-Lipschitz on $u \ge 0$, so expanding in eigenbases and using the Powers–Størmer inequality [43],

$$\|A^z - B^z\|_2^2 \le (1+t^2)\|\sqrt{A} - \sqrt{B}\|_2^2 \le (1+t^2)\|A - B\|_1.$$

Then $A^{1+it} - B^{1+it} = A^z(A^z - B^z) + (A^z - B^z)B^z$, Hölder's inequality with $\|A^z\|_2 = \|B^z\|_2 = 1$, and unitarity of Γ^{-it} give the claim.

Conservation. Let $\Gamma = \gamma_C \otimes \gamma^{\otimes n} \otimes \gamma_G \otimes e^{-\beta H_W}$ and let Γ_R be the same without γ_C. The dilation V conserves energy, so it commutes with Γ and leaves Ξ unchanged. The baths contribute $\Xi(\gamma_G;\gamma_G,t) = 1$, so the final state Ω has $\Xi(\Omega;\Gamma,t) = \Xi(\alpha_C;\gamma_C,t)\,\Xi(\rho;\gamma,t)^n\, e^{it\beta E_i}$, while $\Xi(\alpha_C \otimes \Omega_R;\Gamma,t) = \Xi(\alpha_C;\gamma_C,t)\,\Xi(\Omega_R;\Gamma_R,t)$. Centering the controller's law at a minimizing c gives $|\Xi(\alpha_C;\gamma_C,t)| \ge 1 - |t|s_n \ge \tfrac{1}{2}$ for $|t| \le 1/(2B)$, $B = \max\{1, s_n\}$. With v the variance of the one-copy law of

(ρ,γ), $|\Xi(\rho;\gamma,t)^n - e^{itnD_{\mathrm{in}}}| \le nvt^2/2$. Dividing by the controller's factor and using continuity with $\|\Omega - \alpha_C \otimes \Omega_R\|_1 \le 2\eta_n$, for $|t| \le 1/(2B)$,

$$\bigl|\Xi(\Omega_R;\Gamma_R,t) - e^{itc_R}\bigr| \le 4\sqrt{2\eta_n}\sqrt{1+t^2} + nvt^2/2, \qquad c_R = \beta E_i + nD_{\mathrm{in}}. \tag{2}$$

Test. Let $\kappa = \beta\Delta/2$. Pinching $\sigma^{\otimes n}$ in the at most $(n+1)^d$ eigenspaces of $\gamma^{\otimes n}$ gives a projection P_n that commutes with $\gamma^{\otimes n}$, with $\operatorname{Tr} P_n\sigma^{\otimes n} \ge 1 - f_n$ and $\operatorname{Tr} P_n\gamma^{\otimes n} \le e^{-n(D_{\mathrm{out}}-\kappa)}$ [44]. Take $A = P_n \otimes I_G \otimes |f\rangle\langle f|$. It commutes with Γ_R, $\operatorname{Tr} A\Gamma_R \le e^{-\beta E_f - n(D_{\mathrm{out}}-\kappa)}$, and $\operatorname{Tr} A\Omega_R \ge 1 - \varepsilon_n - q_n - f_n$.

Combination. Put $g = n(D_{\mathrm{out}} - D_{\mathrm{in}} - \kappa) - \beta W \ge n\beta\Delta/2$ and $x = c_R + g/2$ in (1); the first term is at most $e^{-g/2}$. For the second, let $h = \min\{1/(2B), 2/g\}$ and average (2) over $t \in [-h,h]$: $\mathbb{E}[1 - \operatorname{sinc}(h(L_R - c_R))] \le 4\sqrt{2\eta_n}\sqrt{1+h^2} + nvh^2/6$. Since $1 - \operatorname{sinc} y \ge \min\{y^2,1\}/8$,

$$\Pr(L_R \ge x) \le \Pr(|L_R - c_R| \ge g/2) \le 51\sqrt{\eta_n}\max\{1, 16B^2/g^2\} + \frac{16nv}{3g^2}.$$

The last term is $O(1/n)$, and $\max\{1, 16B^2/g^2\} \le \max\{1, (8s_n/(\beta\Delta n))^2\}$ whenever $s_n \ge 1$ or $n \ge 8/(\beta\Delta)$; for the finitely many remaining n we put $f_n = 1$. Collecting $e^{-g/2}$, the test error and $16nv/(3g^2)$ into f_n proves (b). Finally, $s_n \le \mathbb{E}|L_C| = D(\alpha_C\|\gamma_C) + 2\,\mathbb{E}(L_C)_-$, and $\mathbb{E}(L_C)_- \le 1/e$ because $x\ln(1/x) \le 1/e$.

(c), general case. As in the reuse argument of Chiribella, Yang and Renner [5], run the same device m times on the input $\rho^{\otimes n}$, reusing only the controller: each run gets a fresh input block, fresh Gibbs stocks and a fresh battery in $|i\rangle$, and earlier blocks are never touched again. This is a composition of energy-conserving unitaries, so it commutes with $\Gamma_m = \gamma_C \otimes \Gamma_R^{\otimes m}$. After k runs, compare the actual state with $\alpha_C \otimes \Omega_R^{\otimes k}$, tensored with the untouched fresh blocks. The next run contracts their distance, and it maps the comparison state, whose controller is exactly α_C and independent of everything else, to $\Omega_R^{\otimes k} \otimes \Omega_{C,\mathrm{rest}}$, which the return condition puts within η_n of $\alpha_C \otimes \Omega_R^{\otimes(k+1)}$. So the final state $\Omega^{(m)}$ satisfies $\frac12\|\Omega^{(m)} - \alpha_C \otimes \Omega_R^{\otimes m}\|_1 \le m\eta_n$. Conservation and factorization give $\Xi(\Omega^{(m)};\Gamma_m,t) = \Xi(\alpha_C;\gamma_C,t)\,\Xi(\rho;\gamma,t)^{nm}e^{itm\beta E_i}$ and $\Xi(\alpha_C \otimes \Omega_R^{\otimes m};\Gamma_m,t) = \Xi(\alpha_C;\gamma_C,t)\,\Xi(\Omega_R;\Gamma_R,t)^m$. With $\phi(t) = e^{-itc_R}\Xi(\Omega_R;\Gamma_R,t)$ and $\psi(t) = e^{-itnD_{\mathrm{in}}}\Xi(\rho;\gamma,t)^n$, continuity and division by the controller's factor give $|\phi^m - \psi^m| \le 4\sqrt{2m\eta_n}\sqrt{1+t^2}$ for $|t| \le 1/(2B)$.

Let $0 < \eta_n \le 2^{-14}$, $m = \lfloor 1/(8192\eta_n)\rfloor$, $a = n\beta\Delta/4 \ge 1$ and

$$h = \min\{1/2,\ 1/(2B),\ (8mnv)^{-1/2},\ 1/(2a)\}.$$

On $[-h,h]$, $|\psi^m - 1| \le mnvt^2/2 \le 1/16$, so $r = \phi/\psi$ is continuous, $r(0) = 1$, and, since $\sqrt{1+t^2} \le \sqrt{5/4}$ for $|t| \le h$, $|r^m - 1| \le \tau = (16/3)\sqrt{2m\eta_n} \le 1/12$; in particular r has no zero. The continuous logarithm of r that vanishes at 0 satisfies $m\operatorname{Log} r = \operatorname{Log}(r^m)$ on the whole interval, because the difference is continuous, takes values in $2\pi i\mathbb{Z}$ and vanishes at 0. Hence $|\operatorname{Log} r| \le 12\tau/(11m)$ and $|\phi - \psi| \le 2\tau/m$: continuity for one block that is linear in η_n. Averaging over $[-h,h]$ gives $\mathbb{E}[1 - \operatorname{sinc}(h(L_R - c_R))] \le nvh^2/6 + 2\tau/m$, and since $ha \le 1/2$,

$$\Pr(L_R - c_R \ge a) \le \frac{12nv}{a^2} + \frac{16}{3m}\max\Bigl\{1, \frac{B^2}{a^2}\Bigr\} \le \frac{192v}{n\beta^2\Delta^2} + 2^{17}\eta_n \max\Bigl\{1, \Bigl(\frac{4B}{n\beta\Delta}\Bigr)^2\Bigr\}.$$

With the test of (b), $g \ge n\beta\Delta/2$ and $x = c_R + g/2$, the first term of (1) is at most e^{-a} and the second is bounded by this tail. For $\eta_n > 2^{-14}$ the claim is trivial, and for $\eta_n = 0$ any m gives $\tau = 0$. Since $\max\{1, (4B/(n\beta\Delta))^2\} \le \max\{1, (8s_n/(n\beta\Delta))^2\}$ once $n\beta\Delta \ge 4$ (for the finitely many n with $n\beta\Delta < 4$ put $f_n = 1$), this proves (c) with $C = 2^{17}$.

(c), commuting case. The initial state $\alpha_C \otimes \rho^{\otimes n} \otimes \gamma_G \otimes |i\rangle\langle i|$ commutes with every local Hamiltonian, hence with Γ, and so does Ω. Since $\operatorname{Tr}_C[\Omega, H_C \otimes I] = 0$, the partial trace of

$[\Omega, H_C + H_R] = 0$ gives $[\Omega_R, H_R] = 0$, so $\alpha_C \otimes \Omega_R$ commutes with Γ too. Both are block diagonal in the eigenspaces of Γ, on which Γ^{-it} is a scalar. For positive semidefinite X, Y and $f_\delta(x) = (x+\delta)^{1+it}$, $\frac{d}{ds} \operatorname{Tr} f_\delta(Y + s(X-Y)) = \operatorname{Tr} f'_\delta(Y + s(X-Y))(X-Y)$ and $|f'_\delta| = \sqrt{1+t^2}$, so $|\operatorname{Tr} f_\delta(X) - \operatorname{Tr} f_\delta(Y)| \le \sqrt{1+t^2}\|X-Y\|_1$; letting $\delta \to 0$ and summing over blocks gives $|\Xi(\Omega;\Gamma,t) - \Xi(\alpha_C \otimes \Omega_R;\Gamma,t)| \le 2\sqrt{1+t^2}\,\eta_n$. In (2) the term $4\sqrt{2\eta_n}$ becomes $4\eta_n$, and the combination step of (b) gives $32\sqrt{5/4}\,\eta_n \le 36\,\eta_n$ in place of $51\sqrt{\eta_n}$, which is (c) with $C = 36$. A maximizer that commutes with H exists for time-covariant Φ: writing $D(\rho\|\gamma) = D(\mathcal{P}\rho\|\gamma) + A(\rho)$, with $\mathcal{P}$ the pinching in energy and A the relative entropy of asymmetry, which does not increase under covariant channels, the objective at ρ is at most its value at $\mathcal{P}\rho$.

The proof is a statement about one implementation; no controller is measured and nothing is postselected. □

The flat bank. A bank with levels jA, $A = \lceil n/2 \rceil \ln 2/\beta$, $j = 0, \dots, L+1$, in the flat state on a subspace of rank $2j(L+1-j) - 2$ at interior levels and L at the two ends, together with about n pure seed qubits that it returns and a fresh Gibbs bit, runs n resets of zero-Hamiltonian qubits at work $\lceil n/2 \rceil \ln 2/\beta$ with error exactly $1/2$. Its mean energy and entropy return exactly, its joint state with all physical systems (not the purifications of the Gibbs reservoirs) is within $\eta_n = 6/((L+1)(L+2))$ of the product, and its spread is $s_n = \lceil n/2 \rceil \ln 2\, \mathbb{E}|j - (L+1)/2|$ over the bank's level j, which is $\left(3\sqrt{6}\ln 2/32 + O(1/L)\right) 2\lceil n/2 \rceil\, \eta_n^{-1/2}$, about $0.159\, n\eta_n^{-1/2}$ for even n and within 1.5% of that for every $L \ge 10$. What the bank lends is energy variance, A^2, not energy or entropy. If L grows faster than n, for instance $L = n^2$, its return error weighted by its size, $\eta_n K$ with $K = \beta \operatorname{span}(H_C) + \ln \dim C$, tends to zero.

A.5 The single-letter bound

Proof of Proposition 5. (a) The lower bound is Proposition 19, with error and budgets $o(1)$ and $o(n)$, applied to product pairs. If $\Psi^{\otimes N}(\rho) = \Psi^{\otimes N}(\sigma)$, equality passes to every pair of marginals; the energy term is additive, and the entropy decrease is subadditive because local channels cannot increase the total correlation $D(\sigma\|\sigma_1 \otimes \cdots \otimes \sigma_N)$. So $B(\Phi^{\otimes N}) \le NB(\Phi)$, and products attain it. If the environment states differ by $\delta_N \to 0$ in trace distance, let B_δ be the same maximum over pairs whose environment states are within δ; compactness in finite dimension gives $B_\delta \to B(\Phi)$ as $\delta \to 0$, and marginalization, additivity of energy and subadditivity of the entropy decrease give the bound up to $o(N)$.

(b, c) Taking $\rho = \sigma$ gives $B \ge$ the capacity. If $D = \Psi^*(Y)$, then $\operatorname{Tr} \rho D = \operatorname{Tr} Y \Psi(\rho) = \operatorname{Tr} \sigma D$ on every fibre, and the objective is the capacity's. If $\Phi(\gamma) = \gamma$ but $D \notin \operatorname{range} \Psi^*$, there is a Hermitian $X \in \ker \Psi$ with $\operatorname{Tr} XD \ne 0$; then $\rho = \gamma \pm tX$ and $\sigma = \gamma$ are admissible for small t and give $B > 0$. For a qubit, suppose $\Psi(X) = 0$ for a nonzero traceless X. Uhlmann's theorem applied to the two eigenvectors of X gives an isometry relation that, with orthogonality, forces Φ to be unital. A unital Gibbs-preserving qubit channel with $E > 0$ and $\beta < \infty$ is phase-covariant.

□

Proof of Proposition 12. Completeness and Gibbs preservation follow from $x^2\gamma_{11} = y^2\gamma_{00}$. With $\gamma'_0 = (1 + e^{-\beta E_1})^{-1}$, $\gamma'_{01} = \operatorname{diag}(\gamma'_0, 1-\gamma'_0, 0)$ and $a = 1/(2\gamma'_0) \le 1$, the input $\rho = |0\rangle\langle 0|$ and the fixed point $\sigma = a\gamma'_{01} + (1-a)|2\rangle\langle 2|$ both leave the environment in $\operatorname{diag}(1-y^2, y^2)$. Since $\Phi(\sigma) = \sigma$ and $\operatorname{Tr} \rho D = y^2 E_1$, the necessary condition gives $ny^2E_1 - o(n)$. The channel is not covariant: on the input $(|0\rangle + e^{i\phi}|1\rangle)/\sqrt{2}$ the two Kraus operators carry the output coherence with opposite phases. □

A.6 A returned clock reaches the capacity

Proof of Theorem 6. The lower bound is Proposition 5(b). For the construction, let V be the dilation and K the environment Hamiltonian with $V^*(H \otimes I + I \otimes K)V = H$; add a level of energy zero to K if needed, and take the environment F from the reservoir. Stage 1 prepares F in a zero-energy eigenstate of K, at work $k_{\mathrm{B}}T \ln Z_F$ per use with $Z_F = \operatorname{Tr} e^{-\beta K}$, by the deterministic protocols of [45]. Stage 2 applies V to each input, lifted to an exactly energy-conserving unitary by one shared ladder clock in a flat superposition of $L = \lceil n^8 \rceil$ levels: each matrix element between energy eigenstates of SF moves the clock by the energy difference, which is bounded. By exact conservation and the balance identity, the clock's mean energy change equals minus that of $O^n F^n$, which vanishes for every joint input; its joint state with everything else is within $O(n/\sqrt{L})$ of its initial state tensored with the rest, and its entropy changes by $o(1)$. Stage 3 holds the outputs and erases F^n conditionally on O^n, by position-based decoding [2, Prop. 7.1] with the universal conditional and relative typical smoothing operator of [2, Prop. 6.1] for $\Gamma_{OF} = e^{-\beta(H+K)}$ and $\Gamma_O = e^{-\beta H}$. Since these commute, it can be chosen to be a projector commuting with both [2, Prop. 6.1], hence with the physical Hamiltonian, although Φ need not be covariant, and in [2, Thm. 7.1] covariance is used only to make the dilation energy-conserving, which here the clock does: on each input ρ the environment carries $D(\omega^\rho_{OF} \| \gamma_F \otimes \omega^\rho_O) = \ln Z_F + q(\rho)$ nats with $q(\rho) = \beta[F(\rho) - F(\Phi\rho)]$, so the erasure returns $nk_{\mathrm{B}}T[\ln Z_F + \inf_\rho q(\rho)] - o(n)$ uniformly over product inputs, and the postselection technique [46] extends this to arbitrary inputs with polynomial overhead. The outputs and their correlations with any reference are untouched. Converting purity back to mechanical work with [45] gives the net rate $-k_{\mathrm{B}}T \inf_\rho q(\rho)$, which is the capacity. For incommensurate gaps, a finite product of clock ladders handles a rational basis of the energy differences.

For a Gibbs-preserving channel the capacity is zero, so the rate is zero exactly when D lies in the range of Ψ^*, by Proposition 5(c). For a qubit, Proposition 5(c) shows that Gibbs preservation already implies this. □

A.7 The batch law

Proof of Theorem 7. The lower bound is Proposition 19; the construction, with its complete physical ledger, is Theorem 20. □

A.8 Unbalanced channels need a broad clock

Proof of Proposition 8. Write $U_t = e^{-itH}$, $\Phi_t = \mathrm{Ad}_{U_t} \circ \Phi \circ \mathrm{Ad}_{U_t^*}$, $\alpha_t = e^{-itH_C} \alpha_C e^{itH_C}$, $h = \operatorname{span}(H)$, and let Θ_n be the implemented channel on n uses.

Programming. The dilation conserves energy and the reservoirs and battery start stationary, so starting the controller in α_t implements $\mathrm{Ad}_{U_t^{\otimes n}} \circ \Theta_n \circ \mathrm{Ad}_{U_t^{*\otimes n}}$, which is ε_n-close to $\Phi_t^{\otimes n}$; by uniform return, the controller started in α_t also returns within η_n.

Rotating the controller. $\frac{1}{2}\|\alpha_t - \alpha_C\|_1 \le \sqrt{2|t|s_n/\beta}$. Purify $\alpha_C = \sum_a \lambda_a |a\rangle\langle a|$ as $\sum_a \sqrt{\lambda_a}|a\rangle|a\rangle$, let c minimize the spread, and compare e^{-itH_C} on the first factor with the phases e^{-ith_a}, $\beta h_a = c - \ln\lambda_a - \ln Z_C$, on the second. With $\gamma_C = \sum_b \mu_b |b\rangle\langle b|$, $\mu_b = e^{-\beta E_b}/Z_C$, the squared distance of the two purifications is $\sum_{a,b} \lambda_a |\langle a|b\rangle|^2 |e^{-itE_b} - e^{-ith_a}|^2 \le (2|t|/\beta) \sum_{a,b} \lambda_a |\langle a|b\rangle|^2 |\ln(\lambda_a/\mu_b) - c|$, since $|e^{ix} - 1|^2 \le 2|x|$.

A code from the fibre. Since D is not in the range of Ψ^*, there is a Hermitian X with $\Psi(X) = 0$ and $\operatorname{Tr} XD \neq 0$; its normalized positive and negative parts $\rho_\pm$ leave the same environment state, and after relabelling $a = \operatorname{Tr} D(\rho_+ - \rho_-) > 0$. For a fixed $M \ge 2$, let ρ_j, $j = 0, \dots, M-1$, interpolate linearly from ρ_- to ρ_+, and purify each with an orthogonal reference flag, giving orthonormal $|c_j\rangle$. Equality of the environment states gives the Knill–Laflamme conditions $\langle c_j | A_i^* A_k \otimes I | c_{j'} \rangle = c_{ik} \delta_{jj'}$,

so a recovery isometry undoes the Kraus operators on the code. Since $\sum_i A_i^* \dot{A}_i = -iD$, the encoded and recovered channel Λ_t has $\Lambda_0 = \mathrm{id}$ and $\dot{\Lambda}_0 = -i[D_L, \cdot\,]$ with $D_L = \mathrm{diag}(\mathrm{Tr}\,\rho_j D)$, and it is within $C_0 t^2$ of $\mathrm{Ad}_{e^{-itD_L}}$, with C_0 depending only on Φ and H.

Reading the clock. Encode the logical GHZ state $M^{-1/2}\sum_j |j\rangle^{\otimes n}$ into n uses. At the phases $t_k = 2\pi(M-1)k/(Mna)$, $k < M$, the ideal outputs after recovery are orthogonal up to $C_0 K_M^2/n$, where $K_M = 2\pi(M-1)^2/(Ma)$, so a fixed measurement identifies k from the implemented outputs $\sigma_n(t_k)$ with probability at least $1 - \varepsilon_n - C_0 K_M^2/n$. Replacing every $\sigma_n(t_k)$ by $\sigma_n(0)$ leaves success $1/M$ and changes it by at most $\max_k \frac{1}{2}\|\alpha_{t_k} - \alpha_C\|_1 \le \sqrt{2K_M s_n/(\beta n)}$, so $\varepsilon_n \ge 1 - 1/M - C_0 K_M^2/n - \sqrt{2K_M s_n/(\beta n)}$; if $s_n = o(n)$, letting $n \to \infty$ and then $M \to \infty$ gives $\varepsilon_n \to 1$, with no return condition. For the second statement, suppose $\varepsilon_n \le 1 - \delta$ along a subsequence, and take $M = \lceil 4/\delta \rceil$. For large n the success exceeds $1/M$ by $\delta/2$, so some $\sigma_n(t_k)$ is at trace distance at least $\delta/2$ from $\sigma_n(0)$, and since $t_k \le K_M/n$ there is a phase t_* with $\|\sigma_n'(t_*)\|_1 \ge \delta n/K_M$. The orbit acts on n ports only, so $\|\sigma_n''\|_1 \le 4n^2h^2$. For $t = 1/(2hn\sqrt{m})$ with $m \ge (4hK_M/\delta)^2$, a two-outcome measurement therefore separates $\sigma_n(t_*)$ from $\sigma_n(t_* + t)$ by at least $c_\delta/\sqrt{m}$, $c_\delta = \delta/(8hK_M)$, and by multiplicativity of fidelity their m-fold products are at trace distance at least $p_\delta = 1 - e^{-c_\delta^2/2}$.

Reuse. Run the same device m times, as in the proof of Proposition 4(c), once with the controller started in α_{t_*} and once in α_{t_*+t}, with identical fresh inputs, reservoirs and batteries. Each run is within $m\eta_n$ of m independent copies, and the two runs differ only in the controller's initial state, so

$$p_\delta \;\le\; 2m\eta_n + \sqrt{\frac{s_n}{\beta h n\sqrt{m}}}.$$

Put $A_n = \max\{1, (s_n/n)^2\}$ and $R_n = \eta_n A_n \to 0$. If $\eta_n > 0$ take $m = \lceil A_n/\sqrt{R_n} \rceil$, and if $\eta_n = 0$ take $m = \lceil nA_n \rceil$; then $m \to \infty$, $m\eta_n \to 0$ and $s_n/(n\sqrt{m}) \to 0$, a contradiction. For the reliable bound, take $M = 2$ in the code from the fibre and encode one logical qubit into the repetition code $|j\rangle \mapsto |c_j\rangle^{\otimes n}$; the ideal logical block is within nC_0t^2 of $\mathrm{Ad}_{e^{-itnD_L}}$, and the implemented block F_t is within ε_n of it. Pass the same logical qubit through k consecutive blocks, re-encoding it each time and keeping every syndrome, so that $\Gamma_t = F_t^k$ is within $k\varepsilon_n + knC_0t^2$ of $\mathrm{Ad}_{e^{-itknD_L}}$. With logical input $|+\rangle$ and $t_0 = \pi/(kna)$, the outputs $\sigma(0)$ and $\sigma(t_0)$ are at trace distance at least $1/2$ once $k\varepsilon_n \le 1/64$ and n is large, so some $\theta_* \le t_0$ has $\|\sigma'(\theta_*)\|_1 \ge kna/\pi$, while $\|\sigma''\|_1 \le 4k^2n^2h^2$. For $\sqrt{g} \ge \max\{16, 4\pi h/a\}$ and $t = 1/(2knh\sqrt{g})$ a two-outcome measurement separates $\sigma(\theta_*)$ from $\sigma(\theta_* + t)$ by at least $a/(8\pi h\sqrt{g})$, so their g-fold products are at trace distance at least a constant $p > 0$. Running the device for kg blocks, with every logical system and syndrome kept, independent return costs $2kg\eta_n$, and as above $p \le 2kg\eta_n + \sqrt{s_n/(\beta hkn\sqrt{g})}$. Taking $k \asymp 1/(\varepsilon_n + \eta_n)$ and $g \asymp 1/(k\eta_n)$, with constants depending only on Φ and H, gives $s_n \ge c\,\beta n/\sqrt{\eta_n(\varepsilon_n + \eta_n)}$. For X-dephasing the code is exact, the logical channel being $\mathrm{Ad}_{e^{-itD_L}}$ followed by dephasing by $\cos(Et)$, and one may take $N = 64$, $b = p/8192$ and $c = 9p^3E/(4096\sqrt{2})$ with $p = 1 - e^{-1/(128\pi^2)}$. For the third statement, if D lies in the span, equal environment states force $\mathrm{Tr}\,\rho D = \mathrm{Tr}\,\sigma D$, so $B(\Phi)$ is the capacity and Proposition 4(c) applies. □

A.9 How large the clock must be

Proof of Proposition 9. We use the notation and the first two steps of the proof of Proposition 8: starting the controller in α_t implements the channel rotated along the orbit Φ_t, and $\frac{1}{2}\|\alpha_t - \alpha_C\|_1 \le \sqrt{2|t|s_n/\beta}$. The first step uses that the battery and the reservoirs start in states that commute with their Hamiltonians; neither step uses return. Since the complete strategy error bounds the batch error, the lower bounds hold in real time as well.

(a) Since $\Phi_{t+u} = \mathrm{Ad}_{U_t} \circ \Phi_u \circ \mathrm{Ad}_{U_t^*}$, a vanishing derivative $\dot\Phi_0$ would make Φ_t constant, so $\dot\Phi_0 \neq 0$. Choose an input ρ with a reference such that $B = \frac{d}{dt}(\Phi_t \otimes \mathrm{id})(\rho)|_{t=0} \neq 0$, let P project onto the positive part of B, and measure $M = (I+P)/3$. Its probability $p(t)$ on $(\Phi_t \otimes \mathrm{id})(\rho)$ lies in $[1/3, 2/3]$, and $p'(0) = \mathrm{Tr}\, PB/3 = \|B\|_1/6 =: \lambda > 0$. Since $\|\ddot\Phi_t\|_\diamond \leq 4h^2$, $p'(t) \geq \lambda/2$ for $|t| \leq \tau = \lambda/(8h^2)$. Feed ρ, each with its own reference, to the n uses and measure M on every output. For an integer $J \geq 2$ and $v > 0$ put $t_j = jv/\sqrt{n}$, $j < J$, and let $n \geq ((J-1)v/\tau)^2$. Under $\Phi_{t_j}^{\otimes n}$ the outcomes are independent, adjacent means differ by at least $\lambda v/(2\sqrt{n})$, and the sample mean has variance at most $1/(4n)$, so by Chebyshev's inequality the nearest mean identifies j with error at most $\delta = 4/(\lambda v)^2$. The device started in α_{t_j} is within ε_n of $\Phi_{t_j}^{\otimes n}$, so it identifies j with probability at least $1 - \delta - \varepsilon_n$. With every α_{t_j} replaced by α_C the outputs no longer depend on j and the average success is $1/J$, while the replacement changes it by at most $\max_j \frac{1}{2}\|\alpha_{t_j} - \alpha_C\|_1$. Hence

$$\varepsilon_n \;\geq\; 1 - \frac{4}{(\lambda v)^2} - \frac{1}{J} - \sqrt{\frac{2(J-1)v\, s_n}{\beta\sqrt{n}}}.$$

If $s_n = o(\sqrt{n})$, let $n \to \infty$ at fixed J and v, and then let J and v grow.

(b) Put $\zeta = \varepsilon_n + \eta_n \leq 1/32$ and $k = \lfloor 1/(16\zeta) \rfloor \geq 1/(32\zeta)$. As in the proof of Proposition 4(c), run the same device on k blocks of n uses, with fresh inputs, reservoirs and batteries for every block, and keep all old outputs, reservoirs, batteries and references aside. By uniform return the run is within $(k-1)\eta_n$ of k runs with independent controllers, so it implements $\Phi^{\otimes kn}$ within $(k-1)\eta_n + k\varepsilon_n \leq 1/16$, and started in α_t it implements $\Phi_t^{\otimes kn}$ within the same error, since uniform return holds at α_t too. Apply the test of (a) with $J = 2$ and $v = 8/\lambda$ to these kn uses, taking $n \geq N = \lceil (v/\tau)^2 \rceil$. The ideal outputs at the phases 0 and $v/\sqrt{kn}$ are then at distance at least $7/8$, so

$$\tfrac{7}{8} - \tfrac{1}{8} \;\leq\; \tfrac{1}{2}\|\alpha_{v/\sqrt{kn}} - \alpha_C\|_1 \;\leq\; \sqrt{\frac{2vs_n}{\beta\sqrt{kn}}},$$

and $s_n \geq 9\beta\sqrt{kn}/(32v) \geq \frac{9\beta\lambda}{1024\sqrt{2}}\sqrt{n/(\varepsilon_n + \eta_n)}$.

(c), errors that cancel between uses. Balance gives a Stinespring isometry $V : A \to A \otimes E$ and a Hermitian K on E with $V^*(H \otimes I + I \otimes K)V = H$. Add to E a seed level $|0\rangle$, from which V starts, and put $H_E = K + e_0 I$, giving the seed the energy e_0, with e_0 large enough that H_E is positive; with $W = H \otimes I + I \otimes H_E$ this reads $V^*WV = H + e_0 I$. Put $R = WV - V(H + e_0 I)$ and $\kappa = \|R\|$. Then $V^*R = 0$, and $\kappa > 0$, since $\kappa = 0$ would make Φ time-covariant. For n copies, $R_n = W_n V^{\otimes n} - V^{\otimes n}(H_n + ne_0 I)$ is the sum of the terms with R in one factor and V in the others. Every cross term of $R_n^* R_n$ contains V^*R or R^*V, so $R_n^* R_n = \sum_i (R^*R)_i$ and $\|R_n\| \leq \sqrt{n}\,\kappa$ on all inputs, entangled ones included. The energy orbit $V_t = e^{-itW_n} V^{\otimes n} e^{it(H_n + ne_0)}$ therefore satisfies $V_t^* W_n V_t = H_n + ne_0 I$ for every t, and, since its derivative is a unitary conjugate of R_n,

$$\|V_t - V^{\otimes n}\| \;\leq\; |t|\sqrt{n}\,\kappa.$$

(c), the clock. Suppose first that all energy differences lie in $\omega\mathbb{Z}$. Lift $V^{\otimes n}$ to a ladder clock of spacing ω by $\hat V = \sum_{x,y} \Pi_x V^{\otimes n} \Pi_y \otimes T_{(e_y - e_x)/\omega}$, with Π_x the energy projections and T the shifts of the ladder. $\hat V$ is an isometry that conserves energy exactly; complete it, in each total-energy eigenspace, to a full square energy-conserving unitary. Start the clock in a sine packet on ℓ consecutive levels, with nu empty levels on each side, where u bounds the shift of one use. In the Fourier representation of the clock the lift acts as $V_{\theta/\omega}$ at phase θ, and the phase law of a sine packet has $\mathbb{E}\theta^2 \leq \pi^4/(4(\ell+1)^2)$ [9]. So the vector error from $V^{\otimes n}$ tensored with the packet is at most $\pi^2\kappa\sqrt{n}/(2\omega(\ell+1))$, which is at most $a/4$ for $\ell = \lceil 2\pi^2\kappa\sqrt{n}/(\omega a) \rceil$, and the spread of the pure

clock is at most $\beta\omega(\ell+1)/2 \le C\sqrt{n}/a$. Since $V_t^* W_n V_t = H_n + ne_0 I$ for every t, the clock's mean energy does not change, and its final state lies in the span of the $2nu+1$ translates of the packet, so its entropy is at most $\log(2nu+1)$.

If the energy differences generate a group of rank $m \ge 2$, with basis $\omega_1, \dots, \omega_m$, narrow clocks on each coordinate do not suffice: balance of the energy does not give balance of each coordinate. Use instead one clock on $\mathbb{Z}^m$ with Hamiltonian $\omega \cdot q$ and the packet $\phi_\Lambda(q) \propto f_L(\omega \cdot q)\, G(q/\Lambda)$, where $f_L(x) = L^{-1/2}\cos(\pi x/(2L))$ on $|x| \le L$ and G is a smooth, even, nonnegative bump on $(-1,1)^m$. Its energy lies in an interval of length $2L$, so $s_n \le \beta L$ however large Λ is. For each fixed $z \in \mathbb{Z}^m$, $\langle \phi_\Lambda | T_z \phi_\Lambda \rangle \to \int f_L(x) f_L(x - \omega \cdot z)\, dx$ as $\Lambda \to \infty$, by equidistribution of $(\omega_2 q_2 + \cdots + \omega_m q_m)/\omega_1$ modulo one: approximate the periodized profile by trigonometric polynomials; each nonzero harmonic carries an irrational ratio ω_j/ω_1 and its sums stay bounded, while the zero harmonic is a Riemann sum. With the continuous profile f_L, for which $\|f_L'\|_2 = \pi/(2L)$, the same estimate gives squared vector error at most $n\kappa^2\pi^2/(4L^2)$ and exact mean balance. Take $L = 2\pi\kappa\sqrt{n}/a$, $M_n = d^n (d\, d_E)^n$, and Λ so large that the finitely many Gram coefficients that enter are within $\zeta_n = \min\{a^2/(32M_n),\, E_0/(2n^2\|W\|M_n^2)\}$ of their limits, for a fixed energy E_0. The squared vector error is then at most $a^2/8$, and the finite clock's mean energy changes by at most $n\|W\|M_n^2\zeta_n \le E_0/n$. Its entropy is at most $m\log(2nu+1)$. No rate in Λ is claimed.

(c), returning the seeds. The n seeds belong to the controller and must return. After the lift, erase each seed with its own fresh Gibbs record cell R and its own sharp two-level battery W of gap $w > \operatorname{span}(H_E)$. Let R have a ground level of energy 0 and levels r_j of energy $w + e_j - e_0$, start W in $|1\rangle$, and swap $|j\rangle_E|1\rangle_W|g\rangle_R$ with $|0\rangle_E|0\rangle_W|r_j\rangle_R$ for every j, acting as the identity elsewhere. Each swap conserves energy, so this is a full energy-conserving unitary; it fails only if R starts excited, with probability $q \le d_E e^{-\beta(w - \operatorname{span} H_E)}$ independent of the input. Choose w so that $q \le \min\{a^2/(4n), n^{-2}\}$; then $w = O(\log(n/a))$. With probability at least $1 - nq$ every seed returns to $|0\rangle$; this costs at most $\sqrt{nq} \le a/2$ in trace distance and never touches the outputs. So the error is at most $a/2$ and the return error at most $2a$. The seeds' entropy and energy change by at most $n[h(q) + q\log(d_E - 1)]$ and $nq\operatorname{span}(H_E)$, both $o(1)$, so the controller's mean energy changes by $o(1)$ and its entropy by $O(\log n)$. The battery starts sharp and ends at a fixed level except with probability $nq \le 1/n$.

(d), lower bound. Encode a logical qubit as $|1\rangle_S|0\rangle_A$ and $|0\rangle_S|1\rangle_A$, with A held by the observer, and recover with the permutation that writes the logical bit and the syndrome. On this code $\Phi_{X,t}$ acts after recovery as $\mathrm{Ad}_{e^{-itD_L}} \circ \Delta_{\cos Et}$, where $D_L = \operatorname{diag}(-E/2, E/2)$ and Δ_r multiplies the off-diagonal entries by r. With the n-fold repetition code, the decoded implemented channel F_t is within $\varepsilon_n + nE^2t^2/4$ of the unitary $U_t = e^{-itnD_L}$, because $\frac{1}{2}\|\Delta_r - \mathrm{id}\|_\diamond = (1-r)/2$ and $1 - \cos^n x \le nx^2/2$. Take $t = 64\sqrt{\varepsilon_n}/(nE)$. For $n \ge 4096$, F_0 is within ε_n of the identity and F_t within $2\varepsilon_n$ of U_t, whose eigenvalues are $e^{\pm i\varphi/2}$ with $\varphi = 64\sqrt{\varepsilon_n} \le 1$.

Let $|p\rangle$ and $|q\rangle$ purify α_C and α_t with overlap f, their root fidelity, and let W_p and W_q be the resulting Stinespring isometries of the logical channel, keeping everything else as environment. They differ only in the program, so $W_p^* W_q = fI$. If a qubit channel is within δ of a unitary U, each of its Stinespring isometries W has $\|W - U \otimes e\| \le 2\sqrt{\delta}$ for some unit vector e: the Bell state of the channel has overlap at least $1 - \delta$ with that of U, Uhlmann's theorem aligns the environment within $\sqrt{2\delta}$, and the operator norm is at most $\sqrt{2}$ times that distance. With $z = \langle e_p | e_q \rangle$ this gives $\|fI - zU_t\| \le 2\sqrt{\varepsilon_n} + 2\sqrt{2\varepsilon_n} < 8\sqrt{\varepsilon_n}$. On the two eigenvectors of U_t one of $|f - ze^{\pm i\varphi/2}|$ is at least $f\sin(\varphi/2)$, so $f\sin(32\sqrt{\varepsilon_n}) \le 8\sqrt{\varepsilon_n}$ and $f \le \pi/8$. Hence $\frac{1}{2}\|\alpha_t - \alpha_C\|_1 \ge 1 - f > 1/2$, the rotation bound gives $1/4 < 2ts_n/\beta$, and $s_n \ge \beta/(8t) = \beta En/(512\sqrt{\varepsilon_n})$. No return, work bound or variance enters.

(d), construction. Let

$$\tau = \min\{\eta, \sqrt{\varepsilon}, \tfrac{1}{4}\}, \qquad L = \Big\lceil \frac{100\pi^2 n}{\tau} \Big\rceil, \qquad M = mE, \qquad z = \frac{1}{1+e^{\beta M}},$$
$$m = \Big\lceil \max\Big\{2, \frac{10^4 n}{\tau^2}, \frac{1}{\beta E} \ln \frac{10^4 n}{\tau^2}\Big\} \Big\rceil.$$

Each use takes a fresh program bit B in the state $I/2$ and a fresh Gibbs qubit G of gap M, and applies $(U \otimes I_G)J$, where $U = |0\rangle\langle 0|_B \otimes I + |1\rangle\langle 1|_B \otimes X$ and J is the rotation between the system and G with blocks $A = \sqrt{I - H/M}$ and $F = \sqrt{H/M}$, as in [9]. Tracing B and G gives Φ_X composed with a map within $1/(2m)$ of the identity, and the initialized energy gain of the system, B and G is the operator $(E/2)I - 2zH$. Lift each contact to a ladder clock P of spacing $E/2$, prepared in a sine packet on L levels in the middle of a ladder of $L + 2b$ levels, $b = n(2m+3)+1$, a margin that no history, rare hot ones included, can exhaust; raise P by one level per use while a sharp battery loses $E/2$; complete each stage shell by shell to a full energy-conserving unitary; schedule the stages with a zero-Hamiltonian counter that returns exactly; and swap each used program bit out with a fresh flat bit. At a fixed phase of P, tracing the stationary G and B leaves Φ_X conjugated by an energy phase, one phase for the whole strategy; the phase law is even with second moment at most $\pi^4/(4(L+1)^2)$, and the second derivative of the observer's state is at most $16n^2$. So the error is at most $\pi^4 n^2/(L+1)^2 + n/(2m) < \varepsilon$ for every adaptive observer. The energy of P drifts by the operator $2zL_S(H)$, so $0 \le \Delta E_C \le 2nzE$. Compared with the same circuit without the contact J, which moves the clock by at most $3n$ levels in each Fourier fibre, the full state is within $v \le 2\sqrt{nz} + \sqrt{n/m} + n/m + 3\pi^2 n/(2(L+1)) < \tau/10$ of the packet tensored with the ideal rest. So the controller returns jointly with all physical systems within $2v < \eta$, and its entropy changes by at most $h(v^2) + v^2 \ln \dim P = o(n)$. The battery moves from $nE/2$ to 0 on every path, and the spread of the pure packet is at most $\beta EL/4 = O(\beta En/\tau)$. □

A.10 The two families

(a) The Kraus operators of a channel that fixes every state of a basis are diagonal in that basis, so its environment output depends only on the diagonal of the input. Apply the necessary condition of Section 6 with $\rho = \psi$, a top eigenvector of $\Phi^*(H) - H$, and σ its diagonal, so that $S(\sigma) = S(\Phi(\sigma))$. For a mixture of commuting unitaries, the upper bound applies the rare-excitation lift of [9] to the controlled unitary $\sum_j V_j \otimes |j\rangle\langle j|$ with a Gibbs program register; this dilation leaves the device's entropy unchanged on every input. The criterion is the case of equality in the strict concavity of entropy.

(b) The lower bound takes $\sigma = \tau_A \otimes \rho_B$. The upper bound swaps B with a flat cell and uses the returning-clock theorem of [9]. The criterion is the optimality condition of the maximization defining L: a positive operator with zero partial trace must vanish. For the Bell example the finite construction gives the exact gain $\frac{3}{4}EI - 2zH$ with deterministic work, independently of that theorem.

A.11 A channel that reaches its capacity

Proof of Proposition 13. Both σ_0 and σ_1 are states ($\det \sigma_0 = 1/6$, $\det \sigma_1 = 0$), and $\frac{3}{4}\sigma_0 + \frac{1}{4}\sigma_1 = \gamma$, so Φ is a Gibbs-preserving measure-and-prepare channel; it maps the stationary state $|0\rangle\langle 0|$ to σ_0, which does not commute with H. Moreover $\Phi \circ \Delta = \Phi$ and $\Delta \circ \Phi(\rho) = \gamma$ for every ρ, where Δ removes coherences in the energy basis. Gibbs preservation and data processing give $F(\Phi\rho) \le F(\rho)$ with equality at γ, so the capacity is zero; by the accounting identity of [9], no implementation spends

less than $-o(n)$. Directly: on the input $\gamma^{\otimes n}$ the output is $o(1)$-close to $\gamma^{\otimes n}$, so its energy and, by Fannes' inequality, its entropy change by $o(n)$; with the controller's energy and entropy refunds and the battery's entropy refund, conservation of entropy and subadditivity give $\Delta S_R \geq -o(n)$ for the reservoir, and $\Delta E_R \geq \beta^{-1}\Delta S_R$ because it starts in its Gibbs state, so $W_n = \Delta E_R + o(n) \geq -o(n)$.

For the construction, take from the reservoir a qubit D with zero Hamiltonian, in its Gibbs state $I/2$, and let $\Psi(\rho_{AD}) = \gamma_A \otimes \Phi(\rho_A)_D$. Since Φ ignores input coherences and its output sits on a zero-Hamiltonian system, Ψ is time-covariant. Using $S(\rho_{AD}) \leq S(\rho_A) + \ln 2$ and $F(\Phi\rho_A) \leq F(\rho_A)$, its capacity is $a = \beta^{-1}[\ln 2 - S(\gamma)] = E(\frac{3}{4} - \ln 2/\ln 3) \approx 0.119E$, attained at $\gamma \otimes I/2$. Stage 1 runs $\Psi^{\otimes n}$ with the thermal-operation construction for time-covariant channels of [2, Thm. 7.1], at work $n(a+\delta) + o(n)$ for any $\delta > 0$ and error tending to zero. That construction uses only energy-conserving unitaries and Gibbs ancillas (the dilation environments and the copies used for position-based decoding, all reservoir cells), and counts work in zero-Hamiltonian information batteries. We draw those registers from the reservoir in their Gibbs state and convert the purity they consume and return into mechanical work with the protocols for degenerate registers of [45], which give deterministic work at a loss $o(n)$ and vanishing error. So the controller holds only the clock. Stage 2 swaps A^n and D^n in the computational basis, mapping $|a\rangle_A|z\rangle_D|j\rangle_C \mapsto |z\rangle_A|a\rangle_D|j+|a|-|z|\rangle_C$ on a clock with levels of energy jE, which conserves energy. Stage 3 extracts the free energy $na - o(n)$ of the known state $\gamma^{\otimes n}$, now on the zero-Hamiltonian register D^n, which needs no coherence. The net work is $o(n)$. Each stage has deterministic work, and δ is arbitrary.

Since $\Delta \circ \Phi = \gamma$, the computational labels z of $\Phi^{\otimes n}(\rho)$ are independent with law γ for every input ρ, including inputs entangled across uses and with a reference, and so are the labels a of $\gamma^{\otimes n}$. The clock shift $K = |a| - |z|$ is therefore a sum of n independent increments taking the values ± 1 with probability $3/16$ each and 0 with probability $5/8$, so $\mathbb{E}K = 0$ and $\mathbb{E}K^2 = 3n/8$. Take the clock on the ladder $j = 0, \ldots, 2n + L - 1$, starting in the uniform superposition of the levels $n, \ldots, n + L - 1$; on basis states that would leave the ladder, which carry no weight, the gate acts as the identity, which keeps it an energy-conserving unitary. The overlap of the clock with its translate by k is $\max(0, 1 - |k|/L)$ [3], so, purifying everything, the joint state is within trace distance $\sqrt{2\,\mathbb{E}|K|/L} \leq \sqrt{2\sqrt{3n/8}/L}$ of the exact output of Stage 1 tensored with the initial clock; the error of Stage 1 adds to this. The clock's marginal is a mixture of at most $2n+1$ translates of its initial state, so its entropy changes by $O(\log n)$. Its mean energy is unchanged on the exact output of $\Psi^{\otimes n}$, and by energy conservation of the gate it changes by $O(\varepsilon_1 nE) = o(n)$ on the implemented one, where ε_1 is the error of Stage 1. □

The proof uses only $\Phi \circ \Delta = \Phi$, $\Delta \circ \Phi = \gamma\,\mathrm{Tr}$ and $\Phi(\gamma) = \gamma$: every channel with these three properties, on a system with commensurate gaps, reaches capacity zero in the same way, with increments of mean zero and variance $2\,\mathrm{Var}_\gamma(H)$. On a qubit these channels form the family whose output coherences satisfy $c_1 = -(\gamma_0/\gamma_1)c_0$ and $|c_1|^2 \leq \gamma_0\gamma_1$; the channel of the proposition is its extreme point, where σ_1 is pure.

A.12 A process and its reverse

Proof of Proposition 14. The M_i sum to the identity, and $\mathrm{Tr}(P_\pm\gamma) = \frac{5}{12}$, $\mathrm{Tr}(P_2\gamma) = \frac{1}{6}$ give output populations $(\frac{1}{2}, \frac{1}{3}, \frac{1}{6})$, so $\mathcal{N}(\gamma) = \gamma$. The reverse is $\mathcal{N}^\sharp(\rho) = \sum_i \rho_{ii}\,\gamma^{1/2}M_i\gamma^{1/2}/\gamma_i$, which ignores input coherences; its Kraus products contain every $|i\rangle\langle i|$ and $\mathcal{N}^{\sharp *}(H)$ is diagonal, so Theorem 6 gives zero cost.

The Kraus products of $\mathcal{N}$ span $\mathcal{B}(\mathrm{span}\{|+\rangle, |2\rangle\}) + \mathcal{B}(\mathrm{span}\{|-\rangle, |2\rangle\})$, so two inputs leave the environment in the same state exactly when they differ only in their $|+\rangle\langle -|$ coherence. Dephasing σ in the basis $|+\rangle, |-\rangle, |2\rangle$ raises its entropy without changing $\mathcal{N}(\sigma)$, and the coherence of ρ between

$|\pm\rangle$ and $|2\rangle$ can be removed without changing feasibility or energy. So the maximum in $B(\mathcal{N})$ is attained at $\sigma = pP_+ + qP_- + rP_2$ and $\rho = (\sqrt{p}|+\rangle + \sqrt{q}|-\rangle)(\sqrt{p}\langle+| + \sqrt{q}\langle-|) + rP_2$, and with $a = \ln\frac{3}{2}$, $b = \ln 3$,

$$\frac{B(\mathcal{N})}{k_{\mathrm{B}}T} = \max_{p+q+r=1}\Big[-\tfrac{a}{2}p + (\tfrac{3a}{10} + \tfrac{b}{5})q - \tfrac{b}{2}r + a\sqrt{pq} + h(p,q,r) - h\big(p + \tfrac{r}{2}, \tfrac{4q}{5}, \tfrac{q}{5} + \tfrac{r}{2}\big)\Big],$$

where h is the Shannon entropy in nats. The objective is strictly concave, its maximum is interior, and equating the three partial derivatives gives, for $t = \sqrt{q/p}$ and $z = r/(2p)$,

$$12z^2(1+z)\big(\tfrac{3}{2}\big)^{t-1} = z + \tfrac{t^2}{5}, \qquad (1+z)^5 = 3\big(\tfrac{4}{5}\big)^4\Big(\tfrac{1}{5} + \tfrac{z}{t^2}\Big)\big(\tfrac{3}{2}\big)^{4+\frac{5}{2}(1/t-t)},$$

whose positive solution is unique, $t \approx 1.02099$ and $z \approx 0.16236$. Euler's relation for the homogeneous extension gives the value $c = \ln(1+z) + \frac{1}{2}\ln\frac{3}{2}\,(t-1)$. Interval arithmetic at a rational point near the optimum certifies $0.1547103002210 < c < 0.1547103002225$. For the upper bound, put $X = \mathrm{diag}(\frac{a}{2}(t-1), \frac{3a+2b}{10} + \frac{a}{2t}, -\frac{b}{2})$ in the basis $|+\rangle, |-\rangle, |2\rangle$ and $Y_{(i,s)} = X_s/\beta$ on the five Kraus labels. Then $\beta(D - \Psi^*(Y)) = -\frac{a}{2}\begin{pmatrix} t & -1 & 0 \\ -1 & 1/t & 0 \\ 0 & 0 & 0 \end{pmatrix} \le 0$, so $t_Y = 0$, and $\beta g_Y = c$, attained at the faithful optimal σ_*. The construction of Theorem 7 with this Y therefore attains $B(\mathcal{N}) = c\,k_{\mathrm{B}}T$ with deterministic work; since $\sigma_* > 0$, it needs no regularization of $\mathcal{N}$. □

Proof of Proposition 15. If the Choi matrix of Φ has full rank, its Kraus operators span all matrices, and so do their products $A_i^* A_j$; both tests then hold trivially. The map $A_i \mapsto \gamma^{1/2} A_i^* \gamma^{-1/2}$ is invertible, so $\Phi^\sharp$ also has full Kraus rank. The channel $\gamma\,\mathrm{Tr}$ has Choi matrix $I \otimes \gamma > 0$, so every mixture $(1-\delta)\Phi + \delta\,\gamma\,\mathrm{Tr}$ with $\delta > 0$ has full rank. For the family, the M_i form a POVM because $\frac{2}{2+z} + \frac{z}{2+z} = 1$, and $\langle u_\theta|\gamma|u_\theta\rangle = \frac{1}{2} - \frac{z}{6}$ and $\langle v_\theta|\gamma|v_\theta\rangle = \frac{2+z}{6}$ give Gibbs preservation; the Kraus products have no matrix element between u_θ and v_θ, while $\langle u_\theta|H|v_\theta\rangle = k_{\mathrm{B}}T\ln\frac{3}{2}\sin\theta\cos\theta \neq 0$, and the reverse reads only populations. □

A.13 The noise horizon

Proof of Proposition 16. Appendix B gives a lower bound at fixed error with explicit budgets, the uniform noise construction, and the envelopes. □

A.14 A controller that returns only in state

Proof of Corollary 17. (a) Follow the proof of Proposition 19, with $m = n$, but keep the controller's actual changes ΔE_C and ΔS_C on the input ρ of an exact-fibre pair instead of bounding them. Energy conservation, the battery budget u_W, the bath inequality $\beta\Delta E_{\mathrm{bath}} \ge \Delta S_{\mathrm{bath}}$ and $\Delta S_{\mathrm{bath}} \ge A_\rho - \Delta S_C - S(W_f)$ give

$$W_{\mathrm{nom}} \ge \Delta E_{\mathrm{payload}} + \beta^{-1}A_\rho - (\beta^{-1}\Delta S_C - \Delta E_C) - u_W - \beta^{-1}b_W,$$

and $\beta^{-1}\Delta S_C - \Delta E_C = F_C(\alpha_C) - F_C(\alpha'_C) \le R_n$, where α'_C is the controller's final state on ρ. The payload term and A_ρ refer to the same run on ρ, and the lower bound on A_ρ in that proof uses no property of the controller beyond the independence of its initial state from the input. Hence, with $t = \min\{1, 4\varepsilon^{1/4}\}$,

$$W_{\mathrm{nom}} + R_n \ge nB(\Phi) - \beta^{-1}\big\{n[(2\beta h + \ln d)\varepsilon + 4t\ln d] + r_0(\varepsilon) + r_0(t)\big\} - u_W - \beta^{-1}b_W.$$

Under the budgets of Proposition 19, $R_n \le e_C + \beta^{-1} s_C$, so this contains (4). With $\varepsilon \to 0$ and $u_W, b_W = o(n)$ it gives (a).

(b) On a given input, let $x \le \eta_n$ be the trace distance between α'_C and α_C. Writing $\alpha_C - \alpha'_C = x(\tau_+ - \tau_-)$ with states $\tau_\pm$ gives $|\operatorname{Tr} H_C(\alpha_C - \alpha'_C)| \le x \operatorname{span}(H_C)$, and Lemma 18 gives $|S(\alpha_C) - S(\alpha'_C)| \le x \ln d_C + r_0(x)$. Both bounds increase with x, since $r_0'(x) = \ln(1 + 1/x) > 0$. So $F_C(\alpha_C) - F_C(\alpha'_C) \le k_B T[\eta_n K + r_0(\eta_n)]$ on every input, which is the bound on R_n; together with (a) it gives the bound on K.

(c) Take the family of Theorem 20, with controller C_0, battery W and nominal work $nB(\Phi)+o(n)$, whose controller returns jointly within $\eta_n^0 \to 0$ with actual energy and entropy changes $o(n)$. Let $k = \lceil n[B(\Phi)-w]/\omega \rceil$; if $k = 0$ the family itself serves. Otherwise let $N = \lceil k/\eta_n \rceil \ge k$, add to the controller a ladder L with Hamiltonian $\omega \sum_{\ell=1}^{k+N} \ell\, |\ell\rangle\langle\ell|$ in the uniform mixture $\mu_L = N^{-1} \sum_{j=1}^{N} |k+j\rangle\langle k+j|$, and add to the battery a register W_2 with Hamiltonian $k\omega|1\rangle\langle 1|$, starting in $|0\rangle$. Before anything else, apply the permutation that exchanges $|k+j\rangle_L|0\rangle_{W_2}$ with $|j\rangle_L|1\rangle_{W_2}$ for $j = 1, \dots, N$ and fixes every other basis vector. It is an involution and conserves energy, since both vectors of each pair have energy $(k+j)\omega$, and it maps $\mu_L \otimes |0\rangle\langle 0|$ to $\mu'_L \otimes |1\rangle\langle 1|$, with $\mu'_L = N^{-1} \sum_{j=1}^{N} |j\rangle\langle j|$. Then run the device of Theorem 20, which does not act on L or W_2.

The error is that of the original family. The battery $W \otimes W_2$ starts in an energy eigenstate, and W_2 ends in $|1\rangle$ with certainty, so the battery reaches its new nominal endpoint except with the original vanishing probability, with the original entropy and mean deviation; its nominal work is $nB(\Phi) + o(n) - k\omega = nw + o(n)$, since $0 \le k\omega - n[B(\Phi) - w] < \omega$. The ladder ends in μ'_L, in a product with everything else, so by the triangle inequality the joint return error of $C_0 \otimes L$ is at most $\eta_n^0 + \mathsf{D}(\mu'_L, \mu_L)$, with $\mathsf{D} = \frac{1}{2}\|\cdot\|_1$. The two uniform distributions share $N - k$ levels and differ on k levels each, so $\mathsf{D}(\mu'_L, \mu_L) = k/N \le \eta_n$. The ladder's entropy is $\ln N$ before and after, and it loses exactly the energy $k\omega$. So the controller's entropy changes by $o(n)$ and its free energy drops by $k\omega + o(n) = n[B(\Phi) - w] + o(n)$ on every input. The ladder's energy range is $(k + N - 1)\omega \le (1 + 1/\eta_n)k\omega \le (1 + 1/\eta_n)(n[B(\Phi) - w] + \omega)$.

Applied to the ladder alone, the bound of (b) requires $\beta k\omega \le \eta_n[\beta \operatorname{span}(H_L) + \ln(k+N)] + r_0(\eta_n)$, that is, a range of at least $k\omega/\eta_n$ up to terms of order $k_B T[\ln(k + N) + r_0(\eta_n)/\eta_n]$; the range $(k + N - 1)\omega \le (1 + \eta_n)k\omega/\eta_n$ exceeds this by at most the factor $1 + \eta_n$ and those terms.

If the ladder must be pure, take instead $\psi = \sum_{j=1}^{N} a_j |k + j\rangle$ with $a_j = \sqrt{2/(N+1)} \sin(j\pi/(N + 1))$, so that $\sum_j a_j^2 = 1$, and $N = \lceil k\pi/\eta_n \rceil$. The same permutation maps $\psi \otimes |0\rangle$ to $\varphi \otimes |1\rangle$ with $\varphi = \sum_j a_j |j\rangle$, and for pure states $\mathsf{D}(\varphi, \psi) \le \|\varphi - \psi\|$. Let J lower the ladder by one level, $J|\ell\rangle = |\ell - 1\rangle$ for $\ell \ge 2$ and $J|1\rangle = 0$. It is isometric on vectors supported on levels ≥ 2, which every vector it acts on below is, and $\varphi = J^k \psi$. Telescoping, $\|\psi - J^k\psi\| \le \sum_{i=0}^{k-1} \|J^i(\psi - J\psi)\| = k\|\psi - J\psi\|$. With $\theta = \pi/(N + 1)$ and $a_0 = a_{N+1} = 0$, the coefficients satisfy $a_{j-1} + a_{j+1} = 2\cos\theta\, a_j$, so $2\langle\psi|J\psi\rangle = 2\sum_{j=1}^{N-1} a_j a_{j+1} = \sum_{j=1}^{N} a_j(a_{j-1} + a_{j+1}) = 2\cos\theta$ and $\|\psi - J\psi\|^2 = 2 - 2\cos\theta \le \theta^2$. Hence $\mathsf{D}(\varphi, \psi) \le k\pi/(N + 1) \le \eta_n$. The pure ladder also loses exactly $k\omega$, stays pure and stays in a product with everything else, and the range bound holds with π/η_n in place of $1/\eta_n$. □

B Physical construction for the batch law and the noise horizon

This appendix proves Theorem 7 and Proposition 16 with every physical resource counted. Section B.1 states the contract and Section B.2 proves the lower bound with explicit finite budgets. Section B.3 builds the device for a fixed channel, and Sections B.4 and B.5 treat noisy channels, first at a uniform rate and then over finite horizons.

B.1 The operational contract

All logarithms and entropies are natural. Input and output have dimension d and the same Hamiltonian H, shifted to ground zero. Write $h = \operatorname{osc} H = \|H\|$ and fix $\beta > 0$. Every implemented gate is a full finite-dimensional unitary conserving the fixed additive Hamiltonian. Every reservoir supply starts in its Gibbs state, independently of the controller, inputs, other supplies and battery. Used, failed, idle and retired supplies remain part of the physical rest. Mathematical purifications are inert: they are neither manipulated nor required to be returned.

The whole controller C includes the phase bank, initially pure environment seeds, flat routing label, dilation flag and scheduling counter. Its initial state α_C is independent of the inputs. The mechanical battery W starts in a sharp energy eigenstate. Debit is initial minus nominal final battery energy. Uniformly over every joint input and reference, require half diamond error $\varepsilon_n \to 0$ and

$$\mathsf{D}(\omega_{CR_{\mathrm{phys}}}, \alpha_C \otimes \omega_{R_{\mathrm{phys}}}) \to 0, \qquad |\Delta E_C| = o(n), \quad |\Delta S_C| = o(n).$$

Here $\mathsf{D} = \frac{1}{2}\|\cdot\|_1$, R_{phys} is all physical systems other than C, and the endpoint changes are actual and unconditional. The battery has a common nominal sharp endpoint with vanishing failure, $S(W_f) = o(n)$ and $\operatorname{Tr}|H_W - E_{W,\mathrm{nom}}|\omega_W = o(n)$. Startup capital $\beta^{-1}D(\alpha_C\|\gamma_C)$ is declared separately, unrestricted, and not charged as operating work.

This is a batch contract: all outputs are held until cleanup, and one device covers inputs entangled across all uses and with a reference. Real-time release and efficient resources are outside its scope.

B.2 The lower bound with explicit budgets

Let $V : S \to OE$ be a minimal Stinespring isometry, with channel Φ, complement Ψ, and $\dim E = r$. Put

$$D = \Phi^* H - H, \quad h_\Phi(\sigma) = S(\sigma) - S(\Phi\sigma), \quad B(\Phi) = \max_{\Psi\rho=\Psi\sigma} [\operatorname{Tr}\rho D + \beta^{-1} h_\Phi(\sigma)].$$

$h_\Phi(\sigma) = S(E|O)_{V\sigma V^*}$ is concave. The faithful feasible pair $\rho = \sigma = I/d$ gives Slater strong duality after removal of redundant affine constraints:

$$B(\Phi) = \inf_{Y=Y^*} (t_Y + g_Y), \quad t_Y = \lambda_{\max}(D - \Psi^* Y), \quad g_Y = \max_\sigma [\beta^{-1} h_\Phi(\sigma) + \operatorname{Tr}\sigma\Psi^* Y]. \tag{3}$$

Indeed, add $\operatorname{Tr} Y(\Psi\sigma - \Psi\rho)$ to the objective and maximize independently over the two states. The entropy and energy in g_Y use the same σ. The infimum is attained: the same relative-interior feasible point gives a Kuhn–Tucker vector [47, Thm. 28.2].

Exact complement equality for a block implies marginal equalities. Relative-entropy data processing for total correlation gives

$$S(\sigma_n) - S(\Phi^{\otimes n}\sigma_n) \le \sum_i [S(\sigma_i) - S(\Phi\sigma_i)].$$

The energy is additive; hence $B(\Phi^{\otimes n}) = nB(\Phi)$, with products attaining the reverse bound. With complement tolerance ζ, marginal contraction gives $nB_1(0) \le B_n(\zeta) \le nB_1(\zeta)$. Compactness of the single-copy state-pair set gives $B_1(\zeta) \downarrow B_1(0)$ as $\zeta \downarrow 0$.

Lemma 18 (Active-dimension continuity)**.** *Set $r_0(x) = (1+x)\ln(1+x) - x\ln x$, $r_0(0) = 0$. For states at half trace distance $x \le 1$,*

$$|S(A|F)_\rho - S(A|F)_\sigma| \le 2x\ln d_A + r_0(x), \qquad |S(A)_\rho - S(A)_\sigma| \le x\ln d_A + r_0(x).$$

No dimension of F enters.

Proof. The Jordan decomposition is $\rho - \sigma = x(\tau_+ - \tau_-)$. The common state $(\rho + x\tau_-)/(1+x) = (\sigma + x\tau_+)/(1+x)$ is a mixture of weight $p = x/(1+x)$. Conditional entropy of a mixture lies between its averaged value and that average plus $h_2(p)$: use concavity and then append the classical mixture label. Subtract the bounds and use the range $[-\ln d_A, \ln d_A]$. For ordinary entropy the range is $[0, \ln d_A]$. Finally $(1+x)h_2(x/(1+x)) = r_0(x)$. □

Proposition 19 (A finite-error lower bound with explicit budgets)**.** *A finite device implements $\Phi^{\otimes m}$ with half diamond error $\varepsilon \le 1$. Suppose uniformly $|\Delta E_C| \le e_C$, $|\Delta S_C| \le s_C$, $S(W_f) \le b_W$, and $|\langle H_W\rangle_f - E_{W,\mathrm{nom}}| \le u_W$. Define $t = \min\{1, 4\varepsilon^{1/4}\}$ and $b = e_C + u_W + \beta^{-1}(s_C + b_W)$. Then*

$$W_{\mathrm{nom}} \ge mB(\Phi) - \beta^{-1}\{m[(2\beta h + \ln d)\varepsilon + 4t\ln d] + r_0(\varepsilon) + r_0(t)\} - b. \tag{4}$$

Controller and reservoir dimensions, gaps and athermal capital may grow arbitrarily.

Proof. Let I be all final physical hardware except visible outputs. We may assume that the outputs occupy the input registers, appending a conserving swap otherwise, so that I consists of the same systems at the start and at the end; and F purify its initial product state. Purify input τ by Q, $\dim Q \le d^m$. The actual $QOIF$ state is pure; QF remains product. Thus $A_\tau = S(QO|F)_{\mathrm{act}} \ge S(\tau) - S(O_{\mathrm{act}}(\tau))$. Stinespring continuity [40, Thm. 1] gives a common environmental embedding of actual and ideal channels at operator distance at most $a = \sqrt{2\varepsilon}$. For an exact-fibre pair, that is, inputs with $\Psi(\rho) = \Psi(\sigma)$, the aligned ideal IF marginals coincide, so the actual IF marginals are at half trace distance at most $2a$. Uhlmann aligns their actual purifications by a mathematical unitary on QO at half trace distance at most $2\sqrt{a}$. Conditional entropy is invariant under that unitary. Lemma 18 on dimension at most d^{2m} gives $A_\rho \ge A_\sigma - 4mt\ln d - r_0(t)$; the cap $t = 1$ follows from the entropy range. These alignments are comparisons, not physical operations on F. Output entropy continuity then gives

$$A_\rho \ge h_{\Phi^{\otimes m}}(\sigma) - m\varepsilon\ln d - r_0(\varepsilon) - 4mt\ln d - r_0(t).$$

Initial $S(F) = S(I_{\mathrm{initial}})$ and final purity give $\Delta S_I = A_\rho$. Endpoint subadditivity implies $\Delta S_{\mathrm{bath}} \ge A_\rho - \Delta S_C - S(W_f)$. The initial product Gibbs bath obeys $\beta\Delta E_{\mathrm{bath}} \ge \Delta S_{\mathrm{bath}}$ by relative entropy positivity, with no bath dimension or span bound. Conservation and the nominal-to-mean battery budget give $W_{\mathrm{nom}} \ge \Delta E_{\mathrm{payload}} + \beta^{-1}A_\rho - b$. The actual payload gain differs from its ideal gain by at most $mh\varepsilon$; (4) uses the looser $2mh\varepsilon$. Maximize over exact-fibre pairs and apply additivity. □

B.3 The device for a fixed channel

Theorem 20 (The batch law)**.** *For every finite channel whose input and output carry the same Hamiltonian, one family of implementations meets the contract of Appendix B.1 for all inputs, including inputs entangled across uses and with a reference, with nominal mechanical work $nB(\Phi) + o(n)$. Every admissible family has $\liminf W_n/n \ge B(\Phi)$. For each fixed Y in (3), the construction below attains $W_n \le n(t_Y + g_Y) + o(n)$.*

We describe the design of the eraser before the contact that feeds it, so the order of the subsections is not the order in which the device runs. The physical device runs as follows: it starts with the invested environment seeds and the phase bank initialized; it pumps the sharp work budget into the bank; it runs the n lifted contacts, each preceded by its rare-excitation rotation; it decodes the environment's position using the held outputs; it converts the resulting label into mechanical work in the factory; it resets the routing label by a swap with a fresh Gibbs stock; and it closes the counter's cycle. The auxiliary Hamiltonians and the band filter below are tools for designing

one uniform, energy-conserving test; they are not physical resources supplied to the device. All nonstandard physical steps are proved below. The only standard mathematical imports are the finite Hayashi–Nagaoka inequality, Stinespring continuity and Uhlmann's theorem; no asymptotic eraser or general source compiler is assumed.

B.3.1 Faithful auxiliary balance

Fix Y, $\eta = (n+1)^{-6}$ and $\Phi_n = (1-\eta)\Phi + \eta \operatorname{Tr}(\cdot) I/d$. Use Kraus list $\sqrt{1-\eta}K_j$ and $\sqrt{\eta/d}|a\rangle\langle b|$, with fixed size $r' = r + d^2$. On its environment put

$$Y_n = (1-\eta)^{-1} Y \oplus 0, \quad X = \Psi_n^* Y_n = \Psi^* Y, \quad p = \lambda_{\max} Y_n, \quad H_E = pI - Y_n, \quad Z_E = \operatorname{Tr} e^{-\beta H_E}.$$

Write $D_n = \Phi_n^* H - H$, $t_n = \lambda_{\max}(D_n - X)$, $g_n = \max_\sigma [\beta^{-1} h_{\Phi_n}(\sigma) + \operatorname{Tr} X\sigma]$ and $a = p - g_n$. Physical H_E has bounded span for fixed Y. The maximizer σ_n is faithful: adding a small full-rank component to a state with a kernel increases input entropy by a positive $-u \ln u$ term, while $\Phi_n(\sigma) \geq \eta I/d$ makes the output entropy change only linear, as is the X term. Stationarity, with $\omega_n = \Phi_n(\sigma_n)$, therefore gives

$$-\ln \sigma_n + \Phi_n^* \ln \omega_n + \beta X = \beta g_n I.$$

For auxiliary, never supplied Hamiltonians $H'_{\text{in}} = -\beta^{-1} \ln \sigma_n$, $H'_O = -\beta^{-1} \ln \omega_n$,

$$V_n^* (H'_O + H_E) V_n = H'_{\text{in}} + aI. \tag{5}$$

The expectation in σ_n identifies the stationary scalar. The bounds $\|\ln \omega_n\| \leq \ln(d/\eta)$ and $\|\ln \sigma_n\| \leq \ln(d/\eta) + 2\beta\|X\| + \ln d$ follow directly. For $R = (H'_O + H_E)V_n - V_n(H'_{\text{in}} + a)$, $V_n^* R = 0$ and $\|R\| \leq C_Y \ln(n+2)$, where $C_Y = 20[\beta^{-1}(1 + \ln d) + \|Y\|]$ suffices. Tensor residual ranges are orthogonal; hence $\|R^{(n)}\| \leq \sqrt{n}\|R\|$.

B.3.2 Uniform test and two pinchings

Use the positive normalized kernel

$$k(s) = \frac{3}{8\pi} \left[\frac{\sin(s/4)}{s/4} \right]^4, \quad k_w(t) = w k(wt), \quad w = 100 C_Y (n+1)^{3/4} \ln(n+2).$$

Its Fourier support is $[-1, 1]$ before scaling and $\int |s| k(s)\, ds = 12 \ln 2/\pi \leq 12/\pi$. The finite matrix

$$A = \int k_w(t) e^{-it(H'_O + H_E)^{(n)}} V_n^{\otimes n} e^{it((H'_{\text{in}})^{(n)} + na)} dt$$

is a contraction, with exact energy-shift band $na \pm w$. Duhamel gives $\delta_A = \|A - V_n^{\otimes n}\| \leq 12\sqrt{n}\|R\|/(\pi w) = O(n^{-1/4})$. The integral defines a matrix, not a physical filter.

Pinch in auxiliary output energy, with at most $\nu' \leq (n+1)^{dr'}$ shells. Set $\alpha = \theta = (n+1)^{-(dr'+6)}$ and $Q = 1_{[\alpha,\infty)}(\mathcal{P}'(AA^*))$. The pinching inequality $AA^* \leq \nu' \mathcal{P}'(AA^*)$ yields $\|(I-Q)A\| \leq \sqrt{\nu'\alpha}$. In shell b, the exact band gives

$$\Pi_b A \sigma_n^{\otimes n} A^* \Pi_b \geq e^{-\beta(b - na + w)} \Pi_b A A^* \Pi_b.$$

Therefore

$$e^{-\beta(H'_O + H_E)^{(n)}} Q \leq \alpha^{-1} e^{-\beta(na - w)} \mathcal{P}'(A \sigma_n^{\otimes n} A^*).$$

Positive-kernel operator Jensen bounds $A\sigma_n^{\otimes n}A^*$ by $\int k_w(t)U_tV_n^{\otimes n}\sigma_n^{\otimes n}V_n^{*\otimes n}U_t^*dt$. Its environment partial trace is $\omega_n^{\otimes n}$, also after auxiliary pinching. Sandwiching by $\omega_n^{-1/2}$ proves

$$\mathrm{Tr}_{E^n}[e^{-\beta H_E^{(n)}}Q] \le \alpha^{-1}e^{-\beta(na-w)}I. \tag{6}$$

This uses total-energy commutation and cyclicity only on the traced environment.

Now let $Z = \mathcal{P}_{\mathrm{phys}}(Q)$ pinch in the physical energy $H_O^{(n)} + H_E^{(n)}$, and put $P = 1_{[\theta,\infty)}(Z)$. The number of shells is $\nu \le (n+1)^{dr'}$. Since $Q \le \nu Z$, $\|(I-P)Q\| \le \sqrt{\nu\theta}$. Since $P \le Z/\theta$, physical pinching preserves the weighted partial-trace bound in (6). Thus

$$\begin{aligned}
&[P, H_O^{(n)} + H_E^{(n)}] = 0,\\
&\mathrm{Tr}\, PV_n^{\otimes n}\rho V_n^{*\otimes n} \ge 1-\kappa, \quad \kappa = (\delta_A + \sqrt{\nu'\alpha} + \sqrt{\nu\theta})^2,\\
&\mathrm{Tr}_{E^n}[e^{-\beta H_E^{(n)}}P] \le (\alpha\theta)^{-1}e^{-\beta(na-w)}I.
\end{aligned} \tag{7}$$

All bounds act on the entire code and survive any reference extension.

B.3.3 Conserving position decoder

Set $\varepsilon' = (n+1)^{-4}$ and $L_{\mathrm{pos}} = \varepsilon'\alpha\theta Z_E^n e^{\beta(na-w)}$. When $L_{\mathrm{pos}} \ge 2$, set $q = \lfloor L_{\mathrm{pos}} \rfloor$. Controller J is a zero-energy q-level system initially in its true Gibbs state π_q. A separate zero-energy qubit K starts pure 0 and is invested controller capital. Supply q independent Gibbs blocks $A_j \simeq E^n$ in $\gamma_E^{\otimes n}$. First swap E^n with A_j conditioned on $J = j$, a conserving unitary W_1. Define $P_j = P_{O^nA_j}$, $\Sigma = \sum_j P_j$, and $\Omega_j = \Sigma^{-1/2}P_j\Sigma^{-1/2}$ on its support. All effects commute with total energy. The finite Hayashi–Nagaoka inequality is

$$I - (A+B)^{-1/2}A(A+B)^{-1/2} \le 2(I-A) + 4B \quad (0 \le A \le I,\ B \ge 0).$$

The correct slot obeys (7); each wrong slot is independent Gibbs. Its weighted rejection bound gives $\langle\Omega_j\rangle \ge 1-b$, $b = 2\kappa + 4\varepsilon' = O(n^{-1/2})$.

For cyclic shifts S_J, the operator $D_{\mathrm{dec}} = \sum_j \Omega_j \otimes S_J(-j)$ is a contraction: compress the block-diagonal shifts with the isometry having columns $\sqrt{\Omega_j}$ and the unused complement. For $T = D_{\mathrm{dec}}W_1$ the full Halmos square

$$U_T = \begin{pmatrix} T & \sqrt{I-TT^*} \\ \sqrt{I-T^*T} & -T^* \end{pmatrix}_K$$

is unitary; $T\sqrt{I-T^*T} = \sqrt{I-TT^*}T$ proves this, and functional calculus preserves energy commutation. An input branch $J = j, K = 0$ has overlap at least $1-b$ with its post-swap vector and $J = K = 0$. Purify π_q by an inert orthogonal j label and sum these amplitudes. The full purified overlap remains at least $1-b$, retaining all cross-history Gram matrices. Half trace error is at most $\sqrt{2b}$ and $\Pr[J \ne 0 \text{ or } K \ne 0] \le 2b$. No factorization of the post-decoder environment is needed next.

B.3.4 Exact mechanical factory

The factory works on every E^n state, including coherences and entanglement with all retained systems, when $J = 0$. Take positive frequencies $\omega_1 = \beta^{-1}$ and each nonzero eigenenergy of H_E as separate further coordinates; redundant or rationally dependent frequencies are allowed. Let $s \le r'+1$. An environment eigenvector l has a coordinate charge v_l in $\{0,\dots,n\}^s$ and energy $\omega \cdot v_l$. Put

$$C_{\mathrm{tile}} = (n+1)^2, \quad k = \lfloor \ln q - n\ln Z_E - \ln(1+C_{\mathrm{tile}}^{-1})\rfloor, \quad T_{\mathrm{tile}} = \lceil (n+1)^8 + 4(n+|k|+1)\rceil.$$

Stock F has tiles $t \in \{0, \dots, T_{\text{tile}} - 1\}^s$ of energy $\omega \cdot t$ and integer degeneracy $d_t = \lceil C_{\text{tile}} e^{\beta\omega\cdot t} \rceil$. Its true Gibbs partition satisfies $C_{\text{tile}} T_{\text{tile}}^s \le Z_F \le (C_{\text{tile}} + 1) T_{\text{tile}}^s$. A coordinate-interior tile with margin $n + |k| + 1$ has bad probability

$$f \le 2s(1 + C_{\text{tile}}^{-1})(n + |k| + 1)/T_{\text{tile}} = O(n^{-7}).$$

The extraction battery has distinct orthogonal labels W_i, W_f even at $k = 0$, with energies

$$E_i = \max\{0, -k/\beta\}, \qquad E_f = \max\{0, k/\beta\}.$$

Its extraction is k/β, of either sign.

At coordinate-sum u, count all source columns $E_l, J0, W_i, F_{u-v_l}$ simultaneously:

$$\sum_l d_{u-v_l} \le (1 + C_{\text{tile}}^{-1}) C_{\text{tile}} e^{\beta\omega\cdot u} Z_E^n.$$

The target $E_{\text{ground}}, J_{\text{any}}, W_f, F_{u-ke_1}$ has dimension at least $qC_{\text{tile}} e^{\beta\omega\cdot u - k}$, dominating the entire source dimension by the choice of k. Restrict initialized source columns to good tiles. Their target tiles exist. Inject these columns into target basis columns, pair-swap each source with its target, include reverse swaps, and leave unmatched columns unchanged. Distinct battery labels make source and target orthogonal. Different u have disjoint target tiles even when scalar energies coincide, so redundant frequencies cause no collision. Every pair has equal total scalar energy. This is a full conserving square, not postselection. Linearity proves reset and extraction on arbitrary coherent entangled environment inputs. All initialized bad-tile and $J \neq 0$ columns are unchanged. Thus failure is at most $2b + f$. The factory never touches O^n or the bank.

Always swap J with a fresh independent zero-energy Gibbs π_q stock. Then $J = \pi_q$ exactly independently of physical rest; old correlations are declared reservoir residue. K is not discarded. If $L_{\text{pos}} < 2$, use $q = 1$, omit decoding, unconditionally swap E^n with a fresh Gibbs environment block, and run this factory, extracting $-n \ln Z_E/\beta - O(1)$. At the optimizing state $a = \operatorname{Tr} H_E \Psi_n(\sigma_n) - \beta^{-1} S(E|O)$. Conditional entropy and the Gibbs variational inequality show $c = \ln Z_E + \beta a \ge 0$. The fallback condition gives $nc < \beta w - \ln(\varepsilon' \alpha\theta) + \ln 2 = o(n)$; hence it has the same asymptotic rate bound.

B.3.5 Rare excitation and a finite conserving bank

Complete the pure-ground column V_n to a full U_{SE}. Let $B_{\text{hot}} = t_n I - (D_n - X) \ge 0$ and $w_0 = p + t_n$. Its initialized gain effect is $pI + (D_n - X)$. Testing a payload ground state in the nonnegative output energies proves $w_0 \ge 0$. Supply a true Gibbs qubit of gap $M = 100[\beta^{-1} + \|H\| + \|Y\|](n+1)^3$ and excited probability $z = (1 + e^{\beta M})^{-1}$. Before U_{SE} apply

$$J_{\text{hot}} = \begin{pmatrix} C_0 & -F_0 \\ F_0 & C_0 \end{pmatrix}, \quad C_0 = \sqrt{I - B_{\text{hot}}/M}, \quad F_0 = \sqrt{B_{\text{hot}}/M}.$$

For either Gibbs eigenstate the preprocessing is $\Lambda(\rho) = C_0 \rho C_0 + F_0 \rho F_0$, with half diamond disturbance at most $3\|B_{\text{hot}}\|/(2M)$. For $Q_0 = \Phi_n^* H + \Psi_n^* H_E$, exact external gain is

$$\begin{aligned} \text{gain} &= w_0 I + E_n, \\ E_n &= C_0 Q_0 C_0 + F_0 Q_0 F_0 - Q_0 - 2z B_{\text{hot}}, \\ \|E_n\| &\le \chi_n = 3\,\text{osc}(Q_0)\|B_{\text{hot}}\|/M + 2z\|B_{\text{hot}}\| = O(n^{-3}). \end{aligned}$$

The actual energy is controlled by this direct calculation of the effect; a bound by the channel error times the hot gap would not suffice. Before cleanup, the bare state is still in the exact $V_n^{\otimes n}$ code, after preprocessing, with all other systems treated as a reference.

Assign separate bank coordinates to distinct nonzero payload energies, environment energies, M and w_0; dependencies are harmless and $\ell \le d + r' + 2$ is fixed. On a bilateral mathematical bank,

$$\mathcal{L}(U) = \sum_{x,y} U_{xy} |x\rangle\langle y| \otimes T_{a_y - a_x}$$

is a unital norm-preserving $*$-homomorphism, by matrix multiplication and translation addition. It makes unitaries exactly conserving. Use an actual finite box of side $D_P = L + 8(n+1)$, $L = (n+1)^8$, with product sine packet

$$|\xi\rangle = \bigotimes_{j=1}^{\ell} \sqrt{\frac{2}{L+1}} \sum_{u=0}^{L-1} \sin \frac{\pi(u+1)}{L+1} |4(n+1) + u\rangle_j.$$

Pump nw_0 from a two-level sharp battery into the bank, then execute n lifted contacts. External basis charges use at most one unit in each relevant coordinate. All initialized paths, including excited Gibbs paths, remain buffered. Prescribe the lift on all reachable columns and complete orthogonal complements in each exact finite energy shell. Initialized boundary error is zero.

At torus phase ϑ, each contact differs from its phase-zero version by at most $6\|\vartheta\|_2$ and the pump contributes $n\|\vartheta\|_2$. The sine packet has $\mathbb{E}\vartheta_j^2 \le \pi^4/[4(L+1)^2]$. Telescoping the full purified history gives $k_{\text{front}} \le 10\pi^2\sqrt{\ell}(n+1)/(L+1) = O(n^{-7})$. The probability outside the initial bank vector is at most k_{front}^2; no phase measurement or history dephasing occurs.

For stationary fresh seeds and Gibbs qubits, expansion in equal source energy indices proves $\mathrm{Tr}_{EG}[\tau_{EG}\mathcal{L}(A)] = \mathcal{L}_S(\mathrm{Tr}_{EG}[\tau_{EG}A])$. This remains true for arbitrary payload–bank correlations. Conservation, the nw_0 pump and the initialized effect imply

$$|\Delta E_P| \le n\chi_n. \tag{8}$$

Cleanup never acts on the bank, so its growing span never multiplies a trace distance.

B.3.6 Whole-controller and battery refunds

A pure zero-energy counter of dimension $n + O(1)$ executes the pump, contacts, decoder, factory, unconditional label reset and closing identity through a cyclic controlled-unitary. Its full cycle returns the counter exactly. It is counted in C and in startup capital. Set conservatively

$$p_E = 2b + k_{\text{front}} + f, \quad p_K = 2b + k_{\text{front}}, \quad p_C = k_{\text{front}}^2 + p_E + p_K = O(n^{-1/2}).$$

The pure PE^nK block has probability at least $1 - p_C$ in its initial pure target; J and counter are returned exactly and independently of physical rest. Projection implies whole-controller product error at most $2\sqrt{p_C}$. Dephasing in that initial pure vector gives

$$0 \le \Delta S_C \le h_2(p_C) + p_C \ln(D_P^{\ell}(r')^n 2 - 1) = O(\sqrt{n}).$$

Together with (8) and bounded seed energy,

$$|\Delta E_C| \le n\chi_n + n\|H_E\|p_E = O(\sqrt{n}).$$

This bounds the entire block entropy, including cross-correlations. The pump battery is exactly sharp. The extraction battery fails with probability at most p_E, has entropy at most $h_2(p_E)$ and

absolute nominal energy tail at most $|k|p_E/\beta = O(\sqrt{n})$. The composite battery dimension is four. Complete service error is

$$n\eta + \frac{3n\|B_{\mathrm{hot}}\|}{2M} + k_{\mathrm{front}} + \sqrt{2b} = O(n^{-1/4}).$$

The factory and reset act on the complement only, so cannot change the output marginal together with its reference.

Rounding gives

$$W_{\mathrm{nom}} = nw_0 - k/\beta \le n(t_n + g_n) + w + \beta^{-1}[-\ln(\varepsilon'\alpha\theta) + \ln 2 + \ln(1 + C_{\mathrm{tile}}^{-1}) + 1].$$

Finite-dimensional continuity gives $t_n + g_n \to t_Y + g_Y$. Choose Y_j with $t_{Y_j} + g_{Y_j} \le B + 1/j$ and sufficiently large successive horizon thresholds that all errors and normalized budgets for device j are at most $1/j$. This slow diagonal supplies one family at rate B, without an attained infimum or bounded witnesses. Proposition 19 gives the converse.

Controller dimension is $D_P^{\ell}(r')^n q\, 2(n+O(1))$. Capital is the actual packet mean plus $\beta^{-1}[\ln Z_P + n\ln Z_E + \ln 2 + \ln(n+O(1))]$; the Gibbs label J adds no athermality. Bank span is $O(n^{11})$ for fixed Y. Supplies are n hot Gibbs qubits, q Gibbs environment blocks, the explicit factory stock, one flat reset stock and, if needed, a fallback stock. All are finite and retained.

Remark 1 (The eraser). Equation (7) verifies, on the whole code space, the hypotheses of the finite eraser of Faist, Berta and Brandão [2, Prop. 7.1]: an energy-invariant effect that accepts every true state uniformly and rejects Gibbs decoys. We add an energy-conserving decoder and a mechanical factory that turns the erased label into work, with the whole controller counted.

B.4 Uniform noise construction

Assume $\Phi(\gamma) = \gamma$, $B = B(\Phi) > 0$, and $\Phi_\delta = (1-\delta)\Phi + \delta\gamma\,\mathrm{Tr}$, $\gamma = e^{-\beta H}/\mathrm{Tr}\, e^{-\beta H}$. Write $D_\delta = \Phi_\delta^* H - H$, $p_\delta = \lambda_{\max} D_\delta \in [0, h]$ and $B_\delta = p_\delta I - D_\delta \ge 0$, $\|B_\delta\| \le 2h$. Use old Kraus labels $\sqrt{1-\delta}K_j$ and replacement labels $\sqrt{\delta\gamma_a}|a\rangle\langle b|$ in an energy basis. Environment dimension $R = r + d^2$ is fixed. Give old labels energy zero and replacement labels Hamiltonian $I_a \otimes B_\delta/\delta$. Expanding Kraus products gives exactly

$$\Psi_\delta^* H_{E,\delta} = B_\delta, \quad V_\delta^*(H + H_{E,\delta})V_\delta = H + p_\delta I, \quad V_\delta^* H_{E,\delta}^2 V_\delta = B_\delta^2/\delta.$$

The residual norm is at most $5h/\sqrt{\delta}$ and the n-copy norm at most $5h\sqrt{n/\delta}$, by orthogonal ranges. Although $\|H_{E,\delta}\| = O(\delta^{-1})$, its Gibbs mean e_δ satisfies $\beta e_\delta \le \ln R$. Also $r \ge 2$: a Gibbs-preserving unitary commutes with H and has $B = 0$. Therefore $r \le Z_{E,\delta} \le R$.

Theorem 21 (Diverging noise horizon)**.** *For every $t_n = n\delta_n \to \infty$, an admissible family has*

$$W_n = O(nt_n^{-1/4} + \ln(n+1)), \quad \varepsilon_n, \eta_{C,n} = O(t_n^{-1/4}) + o(1),$$

and actual controller energy, entropy and absolute battery tail $O(nt_n^{-1/2}) + o(n)$. Constants depend on fixed Φ, H, β, not on changing environment gaps.

Proof. Use the physical bandfilter with $w = 100\beta^{-1}(1+\beta h)n/t_n^{1/4}$, $\alpha = (n+1)^{-(dR+6)}$, $\varepsilon' = (n+1)^{-4}$. The residual estimate gives $\delta_A = O(t_n^{-1/4})$. Pinch AA^* in physical energy and threshold: then $\kappa = O(t_n^{-1/2} + n^{-6})$ and

$$\mathrm{Tr}_{E^n}(e^{-\beta H_E^{(n)}} Q) \le \alpha^{-1} e^{-\beta(np_\delta - w)} I.$$

Positive-kernel partial trace is the physical Gibbs operator because Φ_δ is Gibbs-preserving. Set $q = \lfloor \varepsilon' \alpha Z_E^n e^{\beta(np_\delta - w)} \rfloor$, eventually at least two since $Z_E \geq r \geq 2$ and $w = o(n)$. Decoder failure uses $b = 2\kappa + 4\varepsilon'$. There is no hot qubit or extra regularization. The finite sine bank has fixed coordinate rank, width n^8, and $k_{\text{front}} = O(n^{-7})$ independently of frequency magnitudes. Initialized mean gain is exactly $p_\delta I$, so the np_δ pump gives $\Delta E_P = 0$ exactly.

After decoding and before the factory, always swap E^n with a fresh independent $\gamma_{E,\delta}^{\otimes n}$ stock. Old code, correlations and failures become physical reservoir rest. The new environment is Gibbs independently of physical rest, not independently of its inert purifier. The factory acts identically on all initialized $J \neq 0$ and bad-tile columns. Good branches reset to ground; failed branches retain the fresh Gibbs environment. Those failure events are independent of it. Distinct battery labels remove cross-branch energy terms. Thus with $p_J \leq 2b + k_{\text{front}}$ and $f = O(n^{-7})$,

$$\Delta E_{E^n} \leq n e_\delta (p_J + f) \leq \beta^{-1} n \ln R (p_J + f) = O(n t_n^{-1/2}) + o(n).$$

This uses the actual Gibbs mean; a bound by $n\|H_E\|$ times the error could diverge. Set $p_E = p_J + f$, $p_K = 2b + k_{\text{front}}$, $p_C = k_{\text{front}}^2 + p_E + p_K$. The same pure-block projection and unconditional label reset give joint return at most $2\sqrt{p_C}$ and $\Delta S_C \leq h_2(p_C) + p_C \ln(D_P^\ell R^n 2 - 1) = o(n)$. Factory rounding gives $np_\delta - k/\beta \leq w + O(\ln n)$ and $|k| = O(n)$ uniformly since $Z_E \leq R$, $p_\delta \leq h$. Battery entropy and tail are bounded by $h_2(p_E)$ and $|k|p_E/\beta$. Capital may grow as $O(n^8(1 + h/\delta_n))$, and is declared rather than charged as operating work. □

B.5 Finite horizons

A genuine Gibbs flag with probabilities $(1 - \delta, \delta)$, controlling a swap with a fresh Gibbs payload, implements $(1 - \delta)\,\mathrm{id} + \delta\gamma\,\mathrm{Tr}$ at zero mechanical work. The flag label does not change, so the square conserves energy. Retain flag and swap stocks as physical rest. Append this module to the fixed-channel engine to obtain $W_N \leq NB + o(N)$ for every noise sequence.

B.5.1 Full-horizon padding and strict positivity

Feed $N - m$ ports of a horizon-N device with independent genuine Gibbs payload supplies, retain their outputs as reservoir rest, and run the entire N-horizon device and cleanup. This implements m visible ports with the original nominal work and the original full actual controller and battery budgets. There is no prefix controller refund. Error against $\Phi^{\otimes m}$ is at most

$$x_{N,m} = \min\{1, \varepsilon_N + 1 - (1 - \delta_N)^m\} \leq \min\{1, \varepsilon_N + m\delta_N\}.$$

Indeed, the product-channel mixture has an all-unreplaced branch of weight $(1 - \delta_N)^m$. This proves the diamond bound.

Define $t(x) = \min\{1, 4x^{1/4}\}$ and $f_0(x) = (2\beta h + \ln d)x + 4t(x) \ln d$. The explicit converse gives

$$\frac{W_N}{N} \geq \frac{m}{N}[B - \beta^{-1} f_0(x_{N,m})] - \frac{r_0(x_{N,m}) + r_0(t(x_{N,m}))}{\beta N} - \frac{b_N}{N}, \tag{9}$$

where b_N is the full-horizon actual budget in Proposition 19. For $m/N \to \theta > 0$, $b_N = o(N)$ is also enough per visible port. With $c = \beta B$, choose explicitly

$$x_* = \min\left\{\frac{1}{2}, \left(\frac{c}{64 \ln d}\right)^4, \frac{c}{4(2\beta h + \ln d)}\right\} > 0.$$

$B > 0$ implies $d \geq 2$, and $f_0(x_*) \leq c/2$. On a bounded subsequence of $N\delta_N$, choose fixed $\theta > 0$ so that $\theta \sup N\delta_N < x_*/2$ and $m = \lfloor \theta N \rfloor$. Vanishing service error and full $o(N)$ budgets then give a strictly positive rate on that subsequence. Thus $o(N)$ work requires $N\delta_N \to \infty$; the preceding construction proves sufficiency. If $N\delta_N \to 0$, take $m = N$ to get $W_N \geq NB - o(N)$, matched by $NB + o(N)$.

B.5.2 Envelopes

There is no finite-horizon optimum at vanishing error without an error and refund schedule. Let $\mathcal{F}_\lambda$ be all implementation families for $\Phi_{\lambda/N}^{\otimes N}$ (ignore finitely many $N < \lambda$) satisfying the contract of Appendix B.1. Define

$$L(\lambda) = \inf_{\mathcal{F}_\lambda} \liminf W_N/N, \qquad U(\lambda) = \inf_{\mathcal{F}_\lambda} \limsup W_N/N, \qquad L(0) = U(0) = B.$$

These optimize over families with vanishing error and actual budgets; they do not identify targets at a fixed prescribed finite error.

Theorem 22 (The finite-noise statement)**.** *For finite $\lambda > 0$, $0 < L(\lambda) \leq U(\lambda) \leq B$. Both envelopes are nonincreasing, and both $\lambda L(\lambda)$ and $\lambda U(\lambda)$ are nondecreasing. They are unchanged on replacing λ/N by any sequence with $N\delta_N \to \lambda$, and tend to B as $\lambda \downarrow 0$. This does not prove $L = U$ or a single limiting rate $w(\lambda)$.*

Proof. Full-horizon padding gives for every $\theta \in (0, 1]$ the common bound

$$L(\lambda), U(\lambda) \geq \theta[B - \beta^{-1} f_0(1 - e^{-\theta\lambda})].$$

Choose θ so that $1 - e^{-\theta\lambda} \leq x_*$; the bound is at least $\theta B/2$. The free replacement module gives upper bound B. Ordinary Gibbs-input accounting also bounds every admissible noisy-channel rate below by zero up to $o(1)$.

Two noise sequences with $N|\delta_N - \delta'_N| \to 0$ have product target half diamond distance at most $N|\delta_N - \delta'_N| \to 0$. The same family is admissible for either, with identical actual endpoint and battery budgets. For $\lambda_2 \geq \lambda_1$ append free replacement probability $(\delta_2 - \delta_1)/(1 - \delta_1)$ to each output. It changes Φ_{δ_1} exactly to Φ_{δ_2} and preserves the controller marginal, actual endpoints and mechanical work. Contractivity preserves joint return. Thus both envelopes are nonincreasing.

Let $0 < \mu = \theta\lambda \leq \lambda$. Padding gives $m = \lfloor \theta N \rfloor$ visible ports, $m\delta_N \to \mu$, vanishing error, and original budgets $o(N) = o(m)$. For an upper envelope, select one N_m for every large m (every m occurs as $\lfloor \theta N \rfloor$), obtaining $U(\mu) \leq U(\lambda)/\theta$. For a lower envelope, select a subsequence attaining a family's lower rate; after removing repeated m values it is increasing. Fill absent m with the bounded-rate fixed-channel engine and free replacement. The filled family's lower rate is no greater than the selected subsequence rate, so $L(\mu) \leq L(\lambda)/\theta$. Apply these arguments to families approaching the respective infima; no minimizing family is assumed. This proves the scaling monotonicity. The lower bound with $\theta = 1$ tends to B as $\lambda \downarrow 0$, and both envelopes are bounded above by B. □

C Results from the companion papers

Table 3 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	[9, Sec. 2]	Section 2
exact return: the error bound; approximate return: every channel at κ_β, a unitary at $\lambda_{\max}(D)$	[9, Prop. 10, Thm. 11]	Theorem 1, Proposition 11(b)
accounting identity	[9, Lemma 4]	Proposition 13
rare-excitation contact, energy-conserving lift, sine-packet clock and the Hadamard construction	[9, App. D]	Propositions 11(a) and 9(c), (d); Appendix A.2

Table 3: Results of the companion papers and of our earlier papers on which an argument here depends.

References

[1] Matteo Lostaglio, Kamil Korzekwa, David Jennings, and Terry Rudolph. Quantum coherence, time-translation symmetry, and thermodynamics. *Physical Review X*, 5:021001, 2015. arXiv:1410.4572.

[2] 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.

[3] Johan Åberg. Catalytic coherence. *Physical Review Letters*, 113:150402, 2014. arXiv:1304.1060.

[4] Kamil Korzekwa, Matteo Lostaglio, Jonathan Oppenheim, and David Jennings. The extraction of work from quantum coherence. *New Journal of Physics*, 18:023045, 2016. arXiv:1506.07875.

[5] Giulio Chiribella, Yuxiang Yang, and Renato Renner. Fundamental energy requirement of reversible quantum operations. *Physical Review X*, 11:021014, 2021. arXiv:1908.10884.

[6] Hiroyasu Tajima, Naoto Shiraishi, and Keiji Saito. Coherence cost for violating conservation laws. *Physical Review Research*, 2:043374, 2020. arXiv:1906.04076.

[7] Hiroyasu Tajima, Koji Yamaguchi, Ryuji Takagi, and Yui Kuramochi. Universal tradeoff relations between resource cost and irreversibility of channels: General-resource Wigner–Araki–Yanase theorems and beyond. arXiv:2507.23760, 2025. arXiv:2507.23760.

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

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

[10] Iman Marvian and Robert W. Spekkens. A no-broadcasting theorem for quantum asymmetry and coherence and a trade-off relation for approximate broadcasting. *Physical Review Letters*, 123:020404, 2019. arXiv:1812.08766.

[11] Matteo Lostaglio and Markus P. Müller. Coherence and asymmetry cannot be broadcast. *Physical Review Letters*, 123:020403, 2019. arXiv:1812.08214.

[12] 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.

[13] Artur S. L. Malabarba, Anthony J. Short, and Philipp Kammerlander. Clock-driven quantum thermal engines. *New Journal of Physics*, 17:045027, 2015.

[14] Mischa P. Woods, Ralph Silva, and Jonathan Oppenheim. Autonomous quantum machines and finite-sized clocks. *Annales Henri Poincaré*, 20:125–218, 2019.

[15] Paul Erker, Mark T. Mitchison, Ralph Silva, Mischa P. Woods, Nicolas Brunner, and Marcus Huber. Autonomous quantum clocks: does thermodynamics limit our ability to measure time? *Physical Review X*, 7:031022, 2017.

[16] Mischa P. Woods and Michał Horodecki. Autonomous quantum devices: when are they realizable without additional thermodynamic costs? *Physical Review X*, 13:011016, 2023.

[17] Ryuji Takagi and Naoto Shiraishi. Correlation in catalysts enables arbitrary manipulation of quantum coherence. *Physical Review Letters*, 128:240501, 2022. arXiv:2106.12592.

[18] Stephen D. Bartlett, Terry Rudolph, Robert W. Spekkens, and Peter S. Turner. Degradation of a quantum reference frame. *New Journal of Physics*, 8:58, 2006. arXiv:quant-ph/0602069.

[19] David Poulin and Jon Yard. Dynamics of a quantum reference frame. *New Journal of Physics*, 9:156, 2007. arXiv:quant-ph/0612126.

[20] Daniel Bedingham and Owen Maroney. The thermodynamic cost of quantum operations. *New Journal of Physics*, 18:113050, 2016. arXiv:1604.03749.

[21] Hiroyasu Tajima and Ryuji Takagi. Gibbs-preserving operations requiring infinite amount of quantum coherence. *Physical Review Letters*, 134:170201, 2025. arXiv:2404.03479.

[22] Fernando Brandão, Michał Horodecki, Nelly Ng, Jonathan Oppenheim, and Stephanie Wehner. The second laws of quantum thermodynamics. *Proceedings of the National Academy of Sciences*, 112:3275–3279, 2015.

[23] Nelly Huei Ying Ng, Laura Mančinska, Cristina Cîrstoiu, Jens Eisert, and Stephanie Wehner. Limits to catalysis in quantum thermodynamics. *New Journal of Physics*, 17:085004, 2015.

[24] Philippe Faist. Universal thermodynamic implementation of a process with a variable work cost. arXiv:2601.20155, 2026.

[25] Runyao Duan, Simone Severini, and Andreas Winter. Zero-error communication via quantum channels, noncommutative graphs, and a quantum Lovász number. *IEEE Transactions on Information Theory*, 59(2):1164–1174, 2013.

[26] Gavin E. Crooks. Quantum operation time reversal. *Physical Review A*, 77:034101, 2008.

[27] Hyukjoon Kwon and M. S. Kim. Fluctuation theorems for a quantum channel. *Physical Review X*, 9:031029, 2019.

[28] Clive Cenxin Aw, Lin Htoo Zaw, Maria Balanzó-Juandó, and Valerio Scarani. Role of dilations in reversing physical processes: Tabletop reversibility and generalized thermal operations. *PRX Quantum*, 5:010332, 2024.

[29] Minjeong Song, Hyukjoon Kwon, and Valerio Scarani. Exact and approximate conditions of tabletop reversibility: when is Petz recovery cost-free? arXiv:2510.26895, 2025.

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

[31] Anurag Anshu, Rahul Jain, and Naqueeb Ahmad Warsi. Building blocks for communication over noisy quantum networks. *IEEE Transactions on Information Theory*, 65(2):1287–1306, 2019.

[32] Rafał Demkowicz-Dobrzański, Jan Czajkowski, and Pavel Sekatski. Adaptive quantum metrology under general Markovian noise. *Physical Review X*, 7:041009, 2017.

[33] Sisi Zhou, Mengzhen Zhang, John Preskill, and Liang Jiang. Achieving the Heisenberg limit in quantum metrology using quantum error correction. *Nature Communications*, 9:78, 2018.

[34] Sisi Zhou and Liang Jiang. Asymptotic theory of quantum channel estimation. *PRX Quantum*, 2:010343, 2021.

[35] Aleksander Kubica and Rafał Demkowicz-Dobrzański. Using quantum metrological bounds in quantum error correction: a simple proof of the approximate Eastin–Knill theorem. *Physical Review Letters*, 126:150503, 2021.

[36] Sisi Zhou, Zi-Wen Liu, and Liang Jiang. New perspectives on covariant quantum error correction. *Quantum*, 5:521, 2021.

[37] Michael A. Nielsen and Isaac L. Chuang. Programmable quantum gate arrays. *Phys. Rev. Lett.*, 79:321, 1997.

[38] Aleksander M. Kubicki, Carlos Palazuelos, and David Pérez-García. Resource quantification for the no-programming theorem. *Phys. Rev. Lett.*, 122:080505, 2019.

[39] Álvaro M. Alhambra, Stephanie Wehner, Mark M. Wilde, and Mischa P. Woods. Work and reversibility in quantum thermodynamics. *Physical Review A*, 97:062114, 2018.

[40] Dennis Kretschmann, Dirk Schlingemann, and Reinhard F. Werner. The information-disturbance tradeoff and the continuity of Stinespring's representation. *IEEE Transactions on Information Theory*, 54:1708–1717, 2008. arXiv:quant-ph/0605009.

[41] Itai Arad, Tomotaka Kuwahara, and Zeph Landau. Connecting global and local energy distributions in quantum spin models on a lattice. *Journal of Statistical Mechanics: Theory and Experiment*, 2016:033301, 2016. arXiv:1406.3898.

[42] Michael Nussbaum and Arleta Szkoła. The Chernoff lower bound for symmetric quantum hypothesis testing. *Ann. Statist.*, 37:1040–1057, 2009.

[43] Robert T. Powers and Erling Størmer. Free states of the canonical anticommutation relations. *Commun. Math. Phys.*, 16:1–33, 1970.

[44] Masahito Hayashi. Optimal sequence of quantum measurements in the sense of Stein's lemma in quantum hypothesis testing. *J. Phys. A*, 35:10759–10773, 2002.

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

[46] Matthias Christandl, Robert König, and Renato Renner. Postselection technique for quantum channels with applications to quantum cryptography. *Physical Review Letters*, 102:020504, 2009.

[47] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, 1970.