Thermodynamically-Efficient Local Computation and the Inefficiency of Quantum Memory Compression

Samuel P. Loomis^{*} and James P. Crutchfield[†]

Complexity Sciences Center and Physics Department, University of California at Davis, One Shields Avenue, Davis, CA 95616

(Dated: February 1, 2020)

Modularity dissipation identifies how locally-implemented computation entails costs beyond those required by Landauer's bound on thermodynamic computing. We establish a general theorem for efficient local computation, giving the necessary and sufficient conditions for a local operation to have zero modularity cost. Applied to thermodynamically-generating stochastic processes it confirms a conjecture that classical generators are efficient if and only if they satisfy retrodiction, which places minimal memory requirements on the generator. This extends immediately to quantum computation: Any quantum simulator that employs quantum memory compression *cannot* be thermodynamically efficient.

Keywords: stochastic process, ϵ -machine, predictor, generator, causal states, modularity, Landauer principle. DOI: xx.xxxx/....

I. INTRODUCTION

Recently, Google AI announced a breakthrough in quantum supremacy, using a 54-qubit processor ("Sycamore") to complete a target computation in 200 seconds, claiming the world's fastest supercomputer would take more than 10,000 years to perform a similar computation [1]. Shortly afterward, IBM announced that they had proven the Sycamore circuit could be successfully simulated on the Summit supercomputer, leveraging its 250 PB storage and 200 petaFLOPS speed to complete the target computation in a matter of days [2]. This episode highlights two important aspects of quantum computing: first, the importance of memory and, second, the subtle relationship between computation and simulation.

Feynman [3] broached the notion that quantum computers would be singularly useful for the simulation of quantum processes, without supposing that this would also make them advantageous at simulating classical processes. Here, we explore issues raised by the recent developments in quantum computing, focusing on the problem of simulating classical stochastic processes via stochastic and quantum computers. We show that using quantum computers to simulate classical processes typically requires nonzero thermodynamic cost, while stochastic computers can theoretically achieve zero cost in simulating classical processes. This supports the viewpoint originally put forth by Feynman—that certain types of computers would each be advantageous at simulating certain physical processes—which challenges the current claims of quantum supremacy. Furthermore, we show that in

both classical and quantum simulations, thermodynamic efficiency places a lower bound on the required memory of the simulator.

To demonstrate both, we must prove a new theorem on the thermodynamic efficiency of local operations. Correlation is a resource: it has been investigated as such, in the formalism of *resource theories* [4] such as that of local operations with classical communication [5], with public communication [6], and many others, as well as the theory of local operations alone, under the umbrella term of *common information* [7–9]. Correlations have long been recognized as a thermal resource [10–13], enabling efficient computation to be performed when taken properly into account. Local operations that act only on part of a larger system are known to never increase the correlation between the part and the whole; most often, they are destructive to correlations and therefore resource-expensive.

Thermodynamic dissipation induced by a local operation say on system A of a bipartite system AB to make a new joint system CB—is classically proportional to the difference in mutual informations [14]:

$$\Delta S_{\text{loc}} = k_{\text{B}}T\left(I(A:B) - I(C:B)\right)$$

This can be asymptotically achieved for quantum systems [15]. By the data processing inequality [16, 17], it is always nonnegative: $\Delta S_{\text{loc}} \geq 0$. Optimal thermodynamic efficiency is achieved when $\Delta S_{\text{loc}} = 0$.

To identify the conditions, in both classical and quantum computation, when this is so, we draw from prior results on saturated information-theoretic inequalities [18–24]. Specifically, using a generalized notion of quantum sufficient statistic [24–27], we show that a local operation on part of a system is efficient if and only if it unitarily preserves the minimal sufficient statistic of the part for

^{*} sloomis@ucdavis.edu

[†] chaos@ucdavis.edu

the whole. Our geometric interpretation of this also draws on recent progress on fixed points of quantum channels [28–31].

Paralleling previous results on ΔS_{loc} [14], our particular interest in locality arises from applying it to thermal transformations that generate and manipulate stochastic processes. This is the study of *information engines* [12, 13, 32–35]. Rooted in computational mechanics [36– 39], which investigates the inherent computational properties of natural processes and the resources they consume, information engines embed stochastic processes and Markovian generators in the physical world, where Landauer's bound for the cost of erasure holds sway [10].

A key result for information engines is the *information*processing Second Law (IPSL): the cost of transforming one stochastic process into another by any computation is at least the difference in their Kolmogorov-Sinai entropy rates [33]. However, actual physical generators and transducers of processes, with their own internal memory dynamics, often exceed the cost required by the IPSL [14]. This arises from the temporal locality of a physical generator—only operating from timestep-to-timestep, rather than acting on the entire process at once. The additional dissipation ΔS_{loc} induced by this temporal locality gives the true thermodynamic cost of operating an information engine.

Previous work explored optimal conditions for a classical information engine to generate a process. Working from the hidden Markov model (HMM) [40] that determines an engine's memory dynamics, it was conjectured that the HMM must be *retrodictive* to be optimal. For this to hold, the current memory state must be a sufficient statistic of the *future* data for predicting the *past* data [14].

Employing a general result on conditions for reversible local computation, the following confirms this conjecture, in the form of an equivalent condition on the HMM's structure. We then extend this, showing that it holds for quantum generators of stochastic processes [15, 41– 49]. Notably, quantum generators are known to provide potentially unbounded advantage in memory storage when compared to classical generators of the same process [42, 43, 45–48]. Surprisingly, the advantage is contingent: optimally-efficient generators—those with $\Delta S_{\text{loc}} = 0$ must not benefit from any memory compression. We show this to be true not only for previously published quantum generators, but for a new family of quantum generators as well, derived from time reversal [47, 50–52]. While important on its own, this also provides a comple-

mentary view to our previous result on quantum generators which showed that a quantum-compressed generator is never less thermodynamically-efficient than the classical generator it compresses [15]. Combined with our current result, one concludes that a quantum-compressed generator is efficient with respect to the generator it compresses but, to the extent that it is compressed, it cannot be optimally efficient. In short, only classical retrodictive generators achieve the lower bound dictated by the IPSL. Practically, this highlights a pressing need to experimentally explore the thermodynamics of quantum computing.

II. THERMODYNAMICS OF QUANTUM INFORMATION RESERVOIRS

The physical setting of our work is in the realm of *in*formation reservoirs—systems all of whose states have the same energy level. Landauer's Principle for quantum systems says that to change an information reservoir Afrom state ρ_A to state ρ'_A requires a work cost satisfying the lower bound:

$$W \ge k_{\rm B}T\ln 2\left(\mathrm{H}[\rho_A] - \mathrm{H}[\rho'_A]\right)$$

where $H[\rho_A]$ is the von Neumann entropy [17]. Note that the lower bound $W_{\min} := k_B T \ln 2 (H[\rho_A] - H[\rho'_A])$ is simply the change in free energy for an information reservoir. Further, due to an information reservoir's trivial Hamiltonian, all of the work W becomes heat Q. Then the total entropy production—of system and environment—is:

$$\Delta S := Q + k_{\rm B} T \ln 2\Delta \operatorname{H}[A]$$
$$= W - W_{\rm min} \; .$$

Thus, not only does Landauer's Principle provide the lower bound, but reveals that any work exceeding W_{\min} represents dissipation.

Reference [14] showed that Landauer's bound may indeed be attained in the quasistatic limit for any channel acting on a *classical* information reservoir. This result generally does not extend to single-shot quantum channels [53]. However, when we consider asymptotically-many parallel applications of a quantum channel, we recover the tightness of Landauer's bound [15].

These statements are exceedingly general. To derive useful results, we must place further constraints on the system dynamics to see how Landauer's bound is affected. Reference [14] introduced the following perspective. Consider a bipartite information reservoir AB, on which we wish to apply the local channel $\mathcal{E} \otimes I_B$, where $\mathcal{E} : \mathcal{B}(\mathcal{H}_A) \to \mathcal{B}(\mathcal{H}_C)$ maps the states of system A into those of system C, transforming the initial joint state ρ_{AB} to the final state ρ_{CB} . The Landauer bound for $AB \to$ CB is given by $W_{\min} = k_B T \ln 2 (\mathrm{H}[\rho_{AB}] - \mathrm{H}[\rho_{CB}])$. However, since we constrained ourselves to use local manipulations, the lowest achievable bound is actually $W_{\text{loc}} := k_{\text{B}}T \ln 2 (\text{H}[\rho_A] - \text{H}[\rho_C])$. Thus, we must have dissipation of at least:

$$\begin{split} \Delta S &\geq W_{\rm loc} - W_{\rm min} \\ &= k_{\rm B} T \ln 2 \left({\rm H}[\rho_A] - {\rm H}[\rho_{AB}] - {\rm H}[\rho_C] + {\rm H}[\rho_{CB}] \right) \\ &= k_{\rm B} T \ln 2 \left({\rm I}[A:B] - {\rm I}[C:B] \right) \; . \end{split}$$

where $I[A : B] = H[\rho_A] + H[\rho_B] - H[\rho_{AB}]$ is the quantum mutual information. And so, we have a minimal *locality* dissipation:

$$\Delta S_{\text{loc}} := k_{\text{B}}T \ln 2 \left(\mathbf{I}[A:B] - \mathbf{I}[C:B] \right) ,$$

which arises because we did not use the correlations to facilitate our erasure. See Fig. 1 for a simple example of this phenomenon.

This local form of Landauer's Principle is still highly general, but the following shows how to examine it for specific classical and quantum computational architectures The key question we ask is: For which architectures can $\Delta S_{\rm loc}$ be made to exactly vanish? We first we consider this problem generally and then provide a solution.

III. REVERSIBLE LOCAL COMPUTATION

Suppose we are given a bipartite system AB with state ρ_{AB} . We wish to determine the conditions for a local channel $\mathcal{E}_A \otimes I_B$ that maps A to C:

$$\rho_{CB}' = \mathcal{E}_A \otimes I_B \left(\rho_{AB} \right)$$

to preserve the mutual information I(A : B) = I(C : B). Proofs of the following results are provided in the Supplementary Material.

Stating our result requires first defining the quantum notion of a sufficient statistic. Previously, quantum sufficient statistics of A for B were defined when AB is a classical-quantum state [27]. That is, when ρ_{AB} commutes with a local measurement on A. They were also introduced in the setting of sufficient statistics for a family of states [24, 25]. This corresponds to the case where AB is quantum-classical— ρ_{AB} commutes with a local measurement on B. Our definition generalizes these cases to fully-quantal correlations between A and B.

We start, as an example, by giving the following definition of a minimum sufficient statistic of a classical joint random variable $XY \sim \Pr(x, y)$ in terms of an equivalence relation. We define the predictive equivalence relation \sim for which $x \sim x'$ if and only if $\Pr(y|x) = \Pr(y|x')$ for all y. The minimum sufficient statistic (MSS) $[X]_Y$ is given by the



FIG. 1. Thermodynamics of locality: Suppose we have two bits XY in a correlated state where 1/2 probability is in XY = 00 and 1/2 probability is in XY = 11. (a) A thermodynamically irreversible operation can be performed to erase only X (that is, set X = 0 without changing Y) if we are not allowed to use knowledge about the state of Y. (b) A reversible operation can be performed to erase X if we are allowed to use knowledge about Y. Both operations have the same outcome given our initial condition, but the nonlocal operation (a) is more thermodynamically costly because it is irreversible. According to Thm. 1, operation (a) is costly since it erases information in X that is correlated with Y.

equivalence classes $[x]_Y := \{x' : x \sim x'\}$. Let us denote $\Sigma := [X]_Y$ and let $\Pr(y|\sigma) := \Pr(y|x)$ for any $x \in \sigma$.

This cannot be *directly* generalized to the quantum setting since correlations between A and B cannot always be described in the form of states conditioned on the outcome of a local measurement on A. If the latter were the case, the state would be classical-quantum, but general quantum correlations can be much more complicated than these. However, we can take the most informative local measurement that does not disturb ρ_{AB} and then consider the "atomic" quantum correlations it leaves behind.

Let ρ_{AB} be a bipartite quantum state. A maximal local

commuting measurement (MLCM) of A for B is any local measurement X with projectors $\{\Pi^{(x)}\}$ on system A such that:

$$\rho_{AB} = \bigoplus_{x} \Pr(X = x) \rho_{AB}^{(x)} ,$$

where:

$$\Pr(X = x) = \operatorname{Tr}\left((\Pi_X^{(x)} \otimes I_B)\rho_{AB}\right)$$

and:

$$\Pr(X = x)\rho_{AB}^{(x)} = (\Pi_X^{(x)} \otimes I_B)\rho_{AB}(\Pi_X^{(x)} \otimes I_B) ,$$

and any further local measurement Y on $\rho_{AB}^{(x)}$ disturbs the state:

$$\rho_{AB}^{(x)} \neq \sum_{y} (\Pi_{Y}^{(y)} \otimes I_B) \rho_{AB}^{(x)} (\Pi_{Y}^{(y)} \otimes I_B)$$

We call the states $\left\{\rho_{AB}^{(x)}\right\}$ quantum correlation atoms.

Proposition 1 (MLCM uniqueness). Given a state ρ_{AB} , there is a unique MLCM of A for B.

Now, as in the classical setting, we define an equivalence class over the values of the MLCM via the equivalence between their quantum correlation atoms. Classically, these atoms are simply the conditional probability distributions $Pr(\cdot|x)$; in the classical-quantum setting, they are the conditional quantum states $\rho_B^{(x)}$. Note that each is defined as a distribution on the variable Y or system B. In contrast, the general quantum correlation atoms $\rho_{AB}^{(x)}$ depend on both systems A and B.

The resulting challenge is resolved in the following way. Let ρ_{AB} be a bipartite quantum state and let X be the MLCM of A for B. We define the correlation equivalence relation $x \sim x'$ over values of X where $x \sim x'$ if and only if $\rho_{AB}^{(x)} = (U \otimes I_B)\rho_{AB}^{(x')}(U^{\dagger} \otimes I_B)$ for a local unitary U. Finally, we define the Minimal Local Sufficient Statistic (MLSS) $[X]_{\sim}$ as the equivalence class $[x]_{\sim} := \{x' : x' \sim x\}$ generated by the relation \sim between correlation atoms. Thus, our notion of sufficiency of A for B is to find the most informative local measurement and then coarse-grain its outcomes by unitary equivalence over their correlation atoms. The correlation atoms and the MLSS $[X]_{\sim}$ together describe the correlation structure of the system AB.

The machinery is now in place to state our result. The proof depends on previous results regarding the fixed points of stochastic channels [28–31] and saturated information-theoretic inequalities [18–24]. This background and the proof are described in the SM.

Theorem 1 (Reversible local operations). Let ρ_{AB} be a bipartite quantum state and let $\mathcal{E}_A \otimes I_B$ be a local operation with $\mathcal{E}_A : \mathcal{B}(\mathcal{H}_A) \to \mathcal{B}(\mathcal{H}_C)$. Suppose X is the MLCM of ρ_{AB} and Y, that of $\rho_{CB} = \mathcal{E}_A \otimes I_B(\rho_{AB})$. The decomposition into correlation atoms is:

$$\rho_{AB} = \bigoplus_{x} \Pr_{A}(x) \rho_{AB}^{(x)} and \qquad (1)$$

$$\rho_{CB} = \bigoplus_{y} \Pr_{C}(y) \rho_{CB}^{(y)} .$$
⁽²⁾

Then, I(A:B) = I(C:B) if and only if \mathcal{E}_A can be expressed by Kraus operators of the form:

$$K^{(\alpha)} = \bigoplus_{x,y} e^{i\phi_{xy\alpha}} \sqrt{\Pr(y,\alpha|x)} U^{(y|x)} , \qquad (3)$$

where $\phi_{xy\alpha}$ is any arbitrary phase and $\Pr(y, \alpha|x)$ is a stochastic channel that is nonzero only when $\rho_{AB}^{(x)}$ and $\rho_{CB}^{(y)}$ are equivalent up to a local unitary operation $U^{(y|x)}$ that maps $\mathcal{H}_{A}^{(x)}$ to $\mathcal{H}_{C}^{(y)}$.

The theorem's classical form follows as a corollary.

Corollary 1 (Reversible local operations, classical). Let XY be a joint random variable and let $\Pr(Z = z | X = x)$ be a channel from \mathcal{X} to some set \mathcal{Z} , resulting in the joint random variable ZY. Then I(X : Y) = I(Z : Y) if and only if $\Pr(Z = z | X = x) > 0$ only when $\Pr(Y = y | Z = z) = \Pr(Y = y | X = x)$ for all y.

In light of the previous section, there is a simple thermodynamic interpretation of Thm. 1 and Cor. 1: local channels that circumvent dissipation due to their locality (*i.e.*, those which have $\Delta S_{\text{loc}} = 0$) are precisely those channels that preserve the sufficiency structure of the joint state. They may create and destroy any information that is not stored in the sufficient statistic and the correlation atoms. However, the sufficient statistic itself must be conserved and the correlation atoms must be only unitarily transformed.

We now turn to apply this perspective to classical and quantum generators—systems that use thermodynamic mechanisms to produce stochastic processes. We compute the necessary and sufficient conditions for these generators to have zero locality dissipation: $\Delta S_{\text{loc}} = 0$. And so, in this way we determine precise criteria for when they are thermodynamically efficient.

IV. THERMODYNAMICS OF CLASSICAL GENERATORS

A classical generator is the physical implementation of a hidden Markov model (HMM) [40] $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\}),$



FIG. 2. Information ratchet sequentially generating a symbol string on an empty tape: At time step t, S_t is the random variable for the ratchet state. The generated symbols in the generated (output) process are denoted by $X_{t-1}, X_{t-2}, X_{t-3}, \ldots$ The most recently generated symbol X_t (green) is determined by the internal dynamics of the ratchet's memory, using heat Q from the thermal reservoir as well as work W from the work reservoir. (Inside ratchet): Ratchet memory dynamics and symbol emission are governed by the conditional probabilities $\Pr(s_{t+1}, x_t | s_t)$, where s_t is the current state at time t, x_t is the generated symbol, and s_{t+1} is the new state. Graphically, this is represented by a hidden Markov model, depicted here as a state-transition diagram in which nodes are states s and edges represent transitions $s \to s'$ labeled by the generated symbol and associated probability: $x : \Pr(s', x|s)$. (Reproduced with permission from Ref. [15].)

where (here) S is countable, \mathcal{X} is finite, and for each $x \in \mathcal{X}$, $\mathbf{T}^{(x)}$ is a matrix with values given by a stochastic channel from S to $S \times \mathcal{Y}$, $T_{s's}^{(x)} := \Pr_{\mathfrak{G}}(s', x|s)$. We define generators to use *recurrent* HMMs, which means the total transition matrix $T_{s's} := \sum_{x} T_{s's}^{(x)}$ is irreducible. In this case, there is a unique stationary distribution $\pi_{\mathfrak{G}}(s)$ over states S satisfying $\pi_{\mathfrak{G}}(s) > 0$, $\sum_{s} \pi_{\mathfrak{G}}(s) = 1$, and $\sum_{s} T_{s's} \pi_{\mathfrak{G}}(s) = \pi_{\mathfrak{G}}(s')$.

During its operation, a generator's function is to produce a *stochastic process*—for each ℓ , a probability distribution $\Pr_{\mathfrak{G}}(x_1 \dots x_{\ell})$ over words $x_1 \dots x_{\ell} \in \mathcal{X}^{\ell}$. The probabilities for words of length ℓ generated by \mathfrak{G} are defined by:

$$\Pr_{\mathfrak{G}}(x_1 \dots x_{\ell}) := \sum_{s_0 \dots s_{\ell} \in \mathcal{S}^{\ell+1}} T_{s_{\ell} s_{\ell-1}}^{(x_{\ell})} \dots T_{s_1 s_0}^{(x_1)} \pi(s_0) \; .$$

Typically, we view a generator as operating over discrete time, writing out a sequence of symbols from $x \in \mathcal{X}$ on a tape, while internally transforming its memory state; see Fig. 2. Starting with an initial state $S_0 \sim \pi(s)$ and empty tape at time t = 0, the entire system at time t is described by the joint random variable $X_1 \dots X_t S_t$, with distribution:

$$\Pr_{\mathfrak{G}}(x_1 \dots x_t, s_t) := \sum_{s_0 \dots s_{t-1} \in \mathcal{S}^t} T^{(x_t)}_{s_t s_{t-1}} \dots T^{(x_1)}_{s_1 s_0} \pi_{\mathfrak{G}}(s_0) \ .$$

Continuing this technique, one can compute the joint random variable $X_1 \dots X_t S_t X_{t+1} S_{t+1}$.

This picture of a generator as operating on a tape while continually erasing and rewriting its internal memory allows us to define the possible thermodynamics, also shown in Fig. 2. Erasure generally requires work, drawn from the work reservoir, while the creation of noise often allows the extraction of work, which is represented in our sign convention by drawing negative work from the reservoir. Producing a process $X_1 \dots X_t \sim$ $\Pr(x_1 \dots x_t)$ of length t has an associated work cost $W \geq -k_{\rm B}T \ln 2H(X_1 \dots X_t)$. The negative sign, as discussed, indicates work $k_{\rm B}T \ln 2 H(X_1 \dots X_t)$ may be transferred from the thermal reservoir to the work reservoir. For large t, this can be asymptotically expressed by the work rate $W/t \geq -k_{\rm B}T \ln 2 h_{\mu}$, where:

$$h_{\mu} := \lim_{t \to \infty} \frac{1}{t} \operatorname{H}[X_1 \dots X_t]$$

is the process' Kolmogorov-Sinai entropy rate [33]. This is a reasonable description of the average entropy rate of a process that is stationary—that is, $\Pr(X_t \dots X_{t+\ell} = x_1 \dots x_{\ell-1})$ is independent of t—and ergodic. Said differently, for large t a typical realization $x_1 \dots x_t$ contains the word $\hat{x}_1 \dots \hat{x}_{\ell}$ approximately $t \times \Pr(\hat{x}_1 \dots \hat{x}_{\ell})$ times. Recurrent generators produce exactly these sorts of processes.

Now, a given generator cannot necessarily be implemented as efficiently as the minimal work rate $W_{\min} := -k_{\rm B}T \ln 2 \ h_{\mu}$ indicates. This is because a generator acts temporally locally, only being able to use its current memory state to generate the next memory state and symbol. The true cost at time t must be bounded below by $W_{\rm loc} := W_{\rm min} + \Delta S_{\rm loc}$, where in this case the locality dissipation is [14]:

$$\Delta S_{\text{loc}} = k_{\text{B}} T \ln 2 \left(I \left(S_t : X_1 \dots X_t \right) \right.$$
$$\left. - I \left(S_{t+1} X_{t+1} : X_1 \dots X_t \right) \right) \,.$$

In this case, the dissipation does not represent work lost to heat but rather the increase in tape entropy that did not facilitate converting heat into work. To understand this in some detail, this section identifies the necessary and sufficient conditions for efficient generators—those with $\Delta S_{\rm loc} = 0$.

To state our result for classical generators, we must introduce two further notions regarding generators. As before, proofs of results are given in the SM. Consider a partition of $S: \mathcal{P} = \{\mathcal{P}_{\theta}\}, \mathcal{P}_{\theta} \cap \mathcal{P}_{\theta'}, \bigcup_{\theta} \mathcal{P}_{\theta} = S$, labeled by index θ . Let:

$$\Pr_{\mathfrak{GP}}(\theta', x|\theta) := \sum_{\substack{s' \in \mathcal{P}_{\theta'} \\ s \in \mathcal{P}_{\theta}}} \Pr_{\mathfrak{G}}(s', x|s) \pi(s|\theta)$$

with $\pi_{\mathfrak{G}}(s|\theta) = \pi_{\mathfrak{G}}(s)/\pi_{\mathfrak{GP}}(\theta)$ and $\pi_{\mathfrak{GP}}(\theta) = \sum_{s\in\mathcal{P}_{\theta}}\pi_{\mathfrak{G}}(s)$. We say a partition $\{\mathcal{P}_{\theta}\}$ is mergeable with respect to the generator $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ if the merged generator $\mathfrak{G}^{\mathcal{P}} = (\mathcal{P}, \mathcal{X}, \{\widetilde{T}_{\theta'\theta}^{(x)}\})$, with transitions $\widetilde{T}_{\theta'\theta}^{(x)} := \Pr(\theta', x|\theta)$, generates the same process as the original.

Pertinent to our goals here is the notion of retrodictive equivalence. Let $\Pr_{\mathfrak{G}}(x_1 \ldots x_t | s_t) := \Pr_{\mathfrak{G}}(x_1 \ldots x_t, s_t) / \pi_{\mathfrak{G}}(s_t)$. Given two states $s, s' \in \mathcal{S}$ of a generator $(\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$, we say that $s \sim s'$ if $\Pr_{\mathfrak{G}}(x_1 \ldots x_t | s) = \Pr_{\mathfrak{G}}(x_1 \ldots x_t | s')$ for all words $x_1 \ldots x_t$. The equivalence class $[S_t]_{\sim}$ is the sufficient statistic of S_t for predicting the past symbols $X_1 \ldots X_t$. The set $\mathcal{P}_{\sim} := \{[s]_{\sim} : s \in \mathcal{S}\}$ of equivalence classes is a partition on \mathcal{S} that we index by σ .

Proposition 2. Given a generator $(S, X, \{T_{s's}^{(x)}\})$, the partition $\mathcal{P} := \mathcal{P}_{\sim}$ induced by retrodictive equivalence is mergeable.

We now state our theorem for efficient classical generators:

Theorem 2. A generator $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ satisfies $I(S_t : X_1 \dots X_t) = I(S_{t+1}X_{t+1} : X_1 \dots X_t)$ for all t if and only if the retrodictively state-merged generator $\mathfrak{G}^{\mathcal{P}} = (\mathcal{P}_{\sim}, \mathcal{X}, \{\widetilde{T}_{\sigma'\sigma}^{(x)}\})$ satisfies $\widetilde{T}_{\sigma'\sigma}^{(x)} \propto \delta_{\sigma,f(\sigma',x)}$ for some function $f : \mathcal{S} \times \mathcal{X} \to \mathcal{S}$.

We say that a generator $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ satisfying $T_{s's}^{(x)} \propto \delta_{s,f(s',x)}$ for some f is co-unifilar. The dual property $T_{s's}^{(x)} \propto \delta_{s',f(s,x)}$ for some f is called unifilar [54]. For every process, there is a unique generator, called the reverse ϵ -machine, constructed by retrodictively statemerging any co-unifilar generator [48]. Similarly, using a different partition, called predictive equivalence on states, any unifilar generator for a process can be state-merged into a unique generator called the forward ϵ -machine of that process [48].

The reverse ϵ -machine has the following property. Let $\overrightarrow{X}_t := X_{t+1}X_{t+1}\dots$ represent all future generated symbols, the reverse ϵ -machine state Σ_t at time t is the minimum sufficient statistic of \overrightarrow{X}_t for predicting $X_1 \dots X_t$. Any generator whose state S_t is a sufficient statistic of \overrightarrow{X}_t for $X_1 \dots X_t$ is called a *retrodictor*. The reverse ϵ -machine can then be considered the minimal retrodictor.

Reference [14] conjectured that the necessary and sufficient condition for $\Delta S_{\text{loc}} = 0$ is that the generator in question is a retrodictor. In the SM we confirm this by establishing that the conditions of Thm. 8 imply the generator is a retrodictor.

A similar result, for classical generators, was presented in [34] where a lower bound on $\Delta S_{\rm loc}$ was derived for predictive generators (Eq. (A23) in [34]). A consequence of this bound is that $\Delta S_{\rm loc} = 0$ only when the predictor is also a retrodictor. However, this bound does not extend to nonpredictive generators. In contrast, Thm. 8 applies to all generators.

Our result is complemented by another recent result [35], which demonstrated how from a predictive generator one can construct a sequence of generators that asymptotically approach a retrodictor and whose dissipation ΔS_{loc} asymptotically approaches zero. Helpfully, this result points to possible perturbative extensions of Thm. 8.

These results bear on the trade-off between dissipation and *memory* for classical generators. The reverse (forward) ϵ -machine, being a state-merging of any co-unifilar (unifilar) generator, is minimal with respect to the counifilar (unifilar) generators via all quantifications of the memory, such as the number of memory states |S| and the entropy H[S] [48].

As a consequence, we now see that the above showed that any *thermodynamically efficient* generator can be statemerged into a co-unifilar generator. This means it can be further state-merged into the reverse ϵ -machine of the process it generates. In short, thermodynamic efficiency comes with a memory constraint. And, when the memory falls below this constraint, *dissipation must be present*.

V. THERMODYNAMICS OF QUANTUM MACHINES

A process' forward ϵ -machine, a key player in the previous section, may be concretely defined as the unique generator $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ for a given process satisfying [38]:

- 1. Recurrence: $T_{s's} := \sum_{x} T_{s's}^{(x)}$ is an irreducible matrix;
- 2. Unifilarity: $T_{s's}^{(x)} > 0$ only when s' = f(s, x) for some function $f : \mathcal{S} \times \mathcal{X} \to \mathcal{S};$
- 3. Predictively Distinct States: $\Pr(x_t x_{t+1} \dots x_{\ell} | s_t) = \Pr(x_t x_{t+1} \dots x_{\ell} | s'_t)$ for all ℓ and $x_t x_{t+1} \dots x_{\ell}$ implies $s_t = s_{t'}$.

 ϵ -Machines are a process' minimal unifilar generators, in the sense that they are smallest with respect to the number of memory states $|\mathcal{S}|$, the entropy H[S], and all other ways of measuring memory, such as the Rényi entropies $H_{\alpha}[S] := \frac{1}{1-\alpha} \log_2 (\sum_s \pi_{\mathfrak{G}}(s)^{\alpha})$. In this, they are unique. However, one can implement ϵ -machines with even lower memory costs, by encoding them in a quantum system and generating symbols by means of a noisy measurement. This encoding is called a *q*-machine. In terms of qubits, as a unit of size, these implementations can generate the same process at a much lower memory cost that the ϵ -machine's bit-based memory cost. It has also been shown that these quantum implementations have a lower locality cost W_{loc} than their corresponding ϵ -machine, and so they are more thermodynamically efficient [15].

This section identifies the constraints for quantum generators to have zero dissipation; that is, $\Delta S_{\rm loc} = 0$. We show that this results in a peculiar pair of constraints. First, the forward ϵ -machine memory must not be smaller than the memory of the reverse ϵ -machine. (This mirrors the results of Thm. 8 in SM.) Second, the quantum generator achieves no compression. That is, the memory of the quantum generator in qubits is precisely the memory of the forward ϵ -machine in bits. Thus, compression of memory and perfect thermodynamic efficiency are exclusive outcomes.

To state this precisely, we review q-machines and introduce several new definitions to capture their properties. (See the SM for the proofs.)

Given a forward ϵ -machine $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$, for any set of phases $\{\phi_{xs} : x \in \mathcal{X}, s \in \mathcal{S}\}$ there is an encoding $\{|\psi_s\rangle : s \in \mathcal{S}\}$ of the memory states \mathcal{S} into a Hilbert space \mathcal{H}_S and a set of Kraus operators $\{K^{(x)} : x \in \mathcal{X}\}$ on said Hilbert space such that:

$$K^{(x)} |\psi_s\rangle = e^{i\phi_{xs}} \sqrt{T^{(x)}_{f(s,x),s}} |\psi_{f(s,x)}\rangle$$

This expression implicitly defines the Kraus operators given the encoding $\{|\psi_s\rangle\}$. The encoding, in turn, is determined up to a unitary transformation by the following constraint on their overlaps:

$$\langle \psi_r | \psi_s \rangle = \sum_{x \in \mathcal{X}} e^{i(\phi_{xs} - \phi_{xr})} \sqrt{T_{r',r}^{(x)} T_{s',s}^{(x)}} \langle \psi_{r'} | \psi_{s'} \rangle \ ,$$

where r' = f(r, x) and s' = f(s, x). This equation has a unique solution for every choice of phases $\{\phi_{xs}\}$ [49].

We note that if $\pi_{\mathfrak{G}}(s)$ is the ϵ -machine's stationary distribution, then the stationary state of this quantum generator is given by the ensemble:

$$\rho_{\pi} = \sum_{s} \pi_{\mathfrak{G}}(s) \left| \psi_{s} \right\rangle \left\langle \psi_{s} \right|$$

and satisfies:

$$\rho_{\pi} = \sum_{x} K^{(x)} \rho_{\pi} K^{(x)\dagger}$$

When we say that a quantum generator uses less memory than its classical counterpart, we mean that dim $\mathcal{H}_S \leq |\mathcal{S}|$, $\mathrm{H}[\rho_{\pi}] \leq \mathrm{H}[S]$, and further that $\mathrm{H}_{\alpha}[\rho_{\pi}] \leq \mathrm{H}_{\alpha}[S]$, where $\mathrm{H}_{\alpha}[\rho_{\pi}] := \frac{1}{1-\alpha} \log_2 \mathrm{Tr}[\rho_{\pi}^{\alpha}]$ are the Rényi-von Neumann entropies [42, 47, 48].

To see this quantum generator as a physical system, as in Fig. 2, requires us interpreting the tape being written on as a series of copies of a single Hilbert space \mathcal{H}_A that represents one cell on the tape. On \mathcal{H}_A we define the computational basis $\{|x\rangle : x \in \mathcal{X}\}$ in which outputs will be written. The system at time t can be described using the joint Hilbert space $\mathcal{H}_{A_1} \otimes \mathcal{H}_{A_t} \otimes \mathcal{H}_S$, where each \mathcal{H}_{A_k} is unitarily equivalent to \mathcal{H}_A , and the state is:

$$\rho_{\mathfrak{G}}(t) := \sum_{x_1 \dots x_t} |x_1 \dots x_t\rangle \langle x_1 \dots x_t| \otimes K^{(x_t \dots x_1)} \rho_{\pi} K^{(x_t \dots x_1)\dagger}$$

where $K^{(x_t...x_1)} = K^{(x_t)} \dots K^{(x_1)}$ and $|x_1 \dots x_t\rangle = |x_1\rangle_{A_1} \otimes \dots \otimes |x_t\rangle_{A_t}$. From this we get the process generated by the ϵ -machine and quantum generator in terms of the Kraus operators as:

$$\Pr_{\mathfrak{G}}(x_1\dots x_t) := \operatorname{Tr}\left[K^{(x_t\dots x_1)}\rho_{\pi}K^{(x_t\dots x_1)\dagger}\right] .$$
(4)

Let us now briefly discuss the thermodynamic properties of quantum generators, homing in on our main result about conditions for their efficiency. The previous section discussed how a process, to be generated, requires the minimal work rate $W_{\rm min} = -k_{\rm B}T \ln 2 \ h_{\mu}$. However, this is not typically achievable for classical generators. The same principle holds for quantum generators: Since they act temporally locally, the true cost at time t is bounded below by $W_{\rm loc} = W_{\rm min} + \Delta S_{\rm loc}$ and the locality dissipation $\Delta S_{\rm loc}$ has the same form:

$$\Delta S_{\rm loc} = k_{\rm B} T \ln 2 \left(I \left(S_t : A_1 \dots A_t \right) - I \left(S_{t+1} A_{t+1} : A_1 \dots A_t \right) \right) \,.$$
(5)

There are two crucial differences, though. First, the mutual information I above is the quantum mutual information derived from the von Neumann entropy. Second, even the work rate $W_{\rm loc}$ is not necessarily achievable in the single-shot case [53]. However, it may be attained for asymptotically parallel generation [15]. We will not concern ourselves with this second problem here. Our intent is to focus, as in the previous section, on the necessary and sufficient conditions for $\Delta S_{\rm loc} = 0$.

The preceding material was, in fact, review. We now introduce a simple partition that may be constructed on the memory states of the ϵ -machine for a given quantum implementation. Specifically, we define the *maximal* commuting partition (MCP) on S to be the most refined partition { \mathcal{B}_{θ} } such that the overlap matrix $\langle \psi_r | \psi_s \rangle$ is



Thermal Efficiency

FIG. 3. Performance trade-offs for q-machines, whose variety and dependence on phases $\{\phi_{xs}\}$ is depicted by a torus: Under all ways of quantifying memory, the q-machines constructed from a predictor achieve nonnegative memory compression [48], and they also have a smaller dissipation ΔS_{loc} , rendering them more thermodynamically efficient [15]. However, to achieve positive compression, they must also have a nonzero ΔS_{loc} , rendering them less efficient than a classical retrodictor.

block-diagonal. That is, $\{\mathcal{B}_{\theta}\}$ is such that $\langle \psi_r | \psi_s \rangle = 0$ if $r \in \mathcal{B}_{\theta}$ and $s \in \mathcal{B}_{\theta'}$ for $\theta \neq \theta'$.

Our result on thermodynamically-efficient quantum generators is as follows.

Theorem 3 (Maximally-efficient quantum generator). Let $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ be a given process' ϵ -machine. Suppose we build from it a quantum generator with encoding $\{\psi_s\}$ and Kraus operators $\{K^{(x)}\}$. Let $\mathcal{B} := \{\mathcal{B}_{\theta}\}$ be the MCP of \mathcal{S} . Then the quantum generator has $\Delta S_{\text{loc}} = 0$ if and only if the partition \mathcal{B} is trivially maximal—in that $|\mathcal{B}_{\theta}| = 1$ for each θ —and the retrodictively state-merged generator $\mathfrak{G}^{\mathcal{B}}$ of \mathfrak{G} is co-unifilar.

We previously found that, in the limit of asymptotically parallel generation, a quantum generator is always more thermodynamically efficient than its corresponding ϵ -machine, in that it has a lower dissipation [15]. Yet this does not imply dissipation can be made to vanish for quantum generators of a process. In fact, only for processes whose forward ϵ -machine is also a retrodictor can dissipation be made to vanish. In these cases, the memory states will be orthogonally encoded, and so no memory compression is achieved, which is seen by the trivial maximality of \mathcal{B} . The situation is heuristically represented in Fig. 3.

VI. THERMODYNAMICS OF REVERSE q-MACHINES

We showed that forward ϵ -machines compressed via the q-machine cannot achieve the efficiency of a classical retrodictor. However, one may wonder what happens to a retrodictor's optimal efficiency if it is directly compressed. We now demonstrate a method for such compression, derived from the time-reversal of the q-machine, and prove that even here any nonzero compression of memory precludes optimal efficiency.

A process' reverse ϵ -machine may be defined similarly to the forward ϵ -machine as the unique generator $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ for a given process satisfying:

- 1. Recurrence: $T_{s's} := \sum_{x} T_{s's}^{(x)}$ is an irreducible matrix;
- 2. Co-unifilarity: $T_{s's}^{(x)} > 0$ only when s = f(s', x) for some function $f: \mathcal{S} \times \mathcal{X} \to \mathcal{S};$
- 3. Retrodictively Distinct States: $Pr(x_1 \dots x_t | s_t) = Pr(x_1 \dots x_t | s'_t)$ for all t and $x_1 \dots x_t$ implies $s_t = s_{t'}$.

Reverse ϵ -machines are a process' minimal co-unifilar generators, in the sense that they are smallest with respect to the number of memory states $|\mathcal{S}|$, the entropy H[S], and all other ways of measuring memory, such as the Rényi entropies $H_{\alpha}[S] := \frac{1}{1-\alpha} \log_2 \left(\sum_s \pi_{\mathfrak{G}}(s)^{\alpha}\right)$.

There is an intricate relationship between forward and reverse ϵ -machines that can only be appreciated in the language of time reversal. The time-reverse of a generator $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ is the generator $\widetilde{\mathfrak{G}} = (\mathcal{S}, \mathcal{X}, \{\widetilde{T}_{s's}^{(x)}\})$ where $\widetilde{T}_{s's}^{(x)} = \pi_s T_{s's}^{(x)} / \pi_{s'}$ [52]. The generator \widetilde{G} is associated with the reverse process, $\Pr_{\widetilde{G}}(x_1 \dots x_t) = \Pr_G(x_t \dots x_1)$. Note that time reversal preserves both the state space \mathcal{S} and the stationary distribution π_s .

Given a process' forward ϵ -machine \mathfrak{F} , its time reverse \mathfrak{F} is the *reverse* ϵ -machine of the *reverse* process. Conversely, given a process' reverse ϵ -machine \mathfrak{G} , its time reverse \mathfrak{G} is the *forward* ϵ -machine of the *reverse* process. Since the stationary distribution and state space are preserved under time reversal, \mathfrak{F} and \mathfrak{F} have the same memory costs, as do \mathfrak{G} and \mathfrak{G} . However, somewhat surprisingly, this does not mean that \mathfrak{F} and \mathfrak{G} have the same memory costs [51].

Previous work compared the results of compressing the forward ϵ -machine \mathfrak{F} of a process and the forward ϵ -machine \mathfrak{S} of the reverse process using the *q*-machine formalism. The result, for compressing \mathfrak{S} , is a *q*-machine that generates the reverse process—remarkably, with identical cost to the *q*-machine constructed from \mathfrak{F} [47].

The q-machine constructed from $\widetilde{\mathfrak{G}}$ is a quantum process and as such can itself undergo quantum time-reversal [50], resulting in a new process that we call the *reverse* q-machine. Just as the q-machine compresses $\widetilde{\mathfrak{G}}$, the reverse q-machine is a compression of \mathfrak{G} . Though the reverse q-machine is derived from the q-machine via timereversal, there is genuinely new physics present, as the dissipation ΔS_{loc} (Eq. (5)) is not invariant under timereversal. Thus, they must be approached as a separate case from the traditional q-machine when examining their thermodynamic efficiency.

The details of the time-reversal are handled in the SM. Here, we present the resulting technique for compressing the reverse ϵ -machine. Given a reverse ϵ -machine $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$, for any set of phases $\{\phi_{xs} : x \in \mathcal{X}, s \in \mathcal{S}\}$ there is an encoding $\{|\psi_s\rangle : s \in \mathcal{S}\}$ of orthogonal states into a Hilbert space \mathcal{H}_S and a set of Kraus operators $\{K^{(x)} : x \in \mathcal{X}\}$ on said Hilbert space such that:

$$K^{(x)} |\psi_s\rangle = \sum_{s' \in \mathcal{S}} e^{i\phi_{xs'}} \sqrt{T^{(x)}_{s's}} |\psi_{s'}\rangle \ .$$

The orthogonality of $\{|\psi_s\rangle\}$ allows us to turn this into an explicit definition of the Kraus operators:

$$K^{(x)} = \sum_{s' \in \mathcal{S}} e^{i\phi_{xs'}} \sqrt{T^{(x)}_{s'f(s',x)}} \left| \psi_{s'} \right\rangle \left\langle \psi_{f(s',x)} \right| \ .$$

The stationary state ρ_{π} of this machine is, unlike the *q*-machine, generically not expressible as an ensemble of the encoding states $\{|\psi_s\rangle\}$. If this were so, the orthogonality of $\{|\psi_s\rangle\}$ would make them a diagonalizing basis for ρ_{π} , and we would achieve no memory compression. Rather, compression is achieved for the reverse *q*-machine precisely because the stationary state ρ_{π} is generically not diagonal in the encoding states—in contrast to the *q*-machine, which derived compression from the nonorthogonality of the encoding states.

The reverse q-machine stochastic dynamics Eq. (4) and thermodynamics Eq. (5) are defined precisely as those for q-machines in the previous section. As before, to prove our result we must define a special partition of the generator states. Here, it is important to note a relationship between a process' forward ϵ -machine $\mathfrak{F} = \left(\mathcal{P}, \mathcal{X}, \left\{R_{p'p}^{(x)}\right\}\right)$ and its reverse ϵ -machine $\mathfrak{G} = \left(\mathcal{S}, \mathcal{X}, \left\{T_{s's}^{(x)}\right\}\right)$. Specifically, the state S_t of \mathfrak{G} after seeing the word $x_1 \dots x_t$ and the state P_t of \mathfrak{F} after the same are related by:

$$\Pr_{\widetilde{\mathfrak{G}}}(s_t|x_1\dots x_t) = \sum_{p_t} \Pr_{\mathcal{C}}(s_t|p_t) \Pr_{\widetilde{\mathfrak{F}}}(p_t|x_1\dots x_t)$$

for some channel $\Pr_{\mathcal{C}}(s|p)$. Let λ_p be the stationary distribution of \mathfrak{F} 's states and let $\Pr_{\mathcal{E}}(s'|s) = \sum_p \Pr_{\mathcal{C}}(s|p) \Pr_{\mathcal{C}}(s'|p) \lambda_p/\pi_s$. Let $\mathcal{B} = \{\mathcal{B}_{\theta}\}$ be the ergodic partition of $\Pr_{\mathcal{E}}(s'|s)$, such that $\Pr_{\mathcal{E}}(s'|s) > 0$ only when $\theta(s) = \theta(s')$. The SM shows that the ρ_{π} is diagonal in the blocks defined by \mathcal{B} .

Our result for reverse q-machines, proven in the SM, can now be stated:

Theorem 4 (Maximally-efficient reverse q-machine). Let $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ be a given process' reverse ϵ -machine.



Thermal Efficiency

FIG. 4. Performance trade-offs for reverse q-machines, whose variety and dependence on $\{\phi_{xs}\}$ is represented by a torus: Under all quantifications of memory, the reverse q-machines constructed from a retrodictor achieve nonnegative memory compression. However, to achieve positive compression, they must also have a nonzero dissipation ΔS_{loc} . The latter renders them less thermodynamically efficient.

Suppose we build from it a reverse q-machine with encoding $\{|\psi_s\rangle\}$ and Kraus operators $\{K^{(x)}\}$. Let $\mathcal{B} := \{\mathcal{B}_{\theta}\}$ be the MCP of \mathcal{S} . Then the reverse q-machine has $\Delta S_{\text{loc}} = 0$ if and only if the partition \mathcal{B} is trivially maximal—in that $|\mathcal{B}_{\theta}| = 1$ for each θ —and the predictively state-merged generator $\mathfrak{G}^{\mathcal{B}}$ of \mathfrak{G} is unifilar.

Notice that this is a similar statement to that made in the last section and is essentially its time reverse. It implies that the only reverse ϵ -machines which can be quantally compressed are those which are also predictive generators. Also, again the trivial maximality of the ergodic partition \mathcal{B} implies an inability to achieve nonzero compression. A heuristic diagram of the situation is shown in Fig. 4.

In conjunction with the previous section, this is a profound result on the efficiency of quantum memory compression. Distinct from the classical case, where Thm. 8 established that *every* process has certain generators that do achieve zero dissipation, Thms. 10 and 3 imply that only *certain* processes have zero-dissipation quantum generators and, moreover, those particular processes achieve no memory compression. The memory states, being orthogonally encoded, take no advantage of the quantum setting to reduce their memory cost.

VII. CONCLUDING REMARKS

We identified the conditions under which local operations circumvent the thermodynamic dissipation $\Delta S_{\rm loc}$ that arises from destroying correlation. We started by showing how a useful theorem can be derived using recent results on the fixed points of quantum channels. We applied it to the setting of local operations to determine the necessary and sufficient conditions for vanishing $\Delta S_{\rm loc}$ in classical and quantum settings, with the aid of a generalized notion of quantum sufficient statistic. We employed this fundamental result to review and extend previous results on the thermodynamic efficiency of generators of stochastic processes. We confirmed a recent conjecture regarding the conditions for vanishing ΔS_{loc} in a classical generator. And, then, we showed the exact same conditions hold for quantum generators, even to the point of requiring orthogonal encoding of memory states. This implies the profound result that quantum memory compression and perfect efficiency ($\Delta S_{\text{loc}} = 0$) are incompatible.

It is appropriate here to recall the lecture by Feynman in the early days of thinking about quantum computing, in which he observed that quantum systems can only be simulated on classical (even probabilistic) computers with great difficulty, but on a fundamentally-quantum computer they could be more realistically simulated [3]. Here, we considered the task of simulating a classical stochastic process by two means: one by using fundamentallyclassical but probabilistic machines and the other by using a fundamentally-quantum machine. Previous results generally indicated quantum machines are advantageous in memory for this task, in comparison to their classical counterparts. Historically, this led to a much stronger notion of "quantum supremacy" than Feynman proposed: quantum computers may be advantageous in all tasks [55].

However, the quantum implementation we examined, though advantageous in memory, requires nonzero dissipa-

- Frank Arute et al. Quantum supremacy using a programmable superconducting processor. Nature, 574(7779):505–510, 2019.
- [2] E. Pednault, J. A. Gunnels, G. Nannicini, L. Horesh, and R. Wisnieff. Leveraging Secondary Storage to Simulate Deep 54-qubit Sycamore Circuits. 2019. arXiv:1910.09534 [quant-ph]. 1
- [3] M. Junge, R. Renner, D. Sutter, M. M. Wilde, and A. Winter. Recoverability in quantum information theory. *Int. J. Theor. Phys.*, 21(6/7), 1982. 1, 10
- [4] B. Coecke, T. Fritz, and R. W. Spekkens. A mathematical theory of resources. *Info. Comp.*, 250:59, 2016. 1
- [5] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki. Quantum entanglement. *Rev. Mod. Phys.*, 81:865, 2009. 1
- [6] U. M. Maurer. Secret key agreement by public discussion from common information. *IEEE Trans. Info. Th.*, 39:3, 1993. 1
- [7] P. Gács and J. Körner. Common information is far less than mutual information. Problems of Control and Info. Theory, 2:119, 1973. 1

tion in order to cash in on that advantage. Furthermore, not every process necessarily has a quantum generator that achieves zero dissipation. This is in sharp contrast to the classical outcome. And so, this returns us to the spirit of Feynman's vision for simulating physics, in which it may sometimes be the case that the best machine to simulate a classical stochastic process is a classical stochastic computer—at least, thermodynamically speaking.

To further exercise these results, further extensions must be made to quantum generators, beyond the *q*-machine and its time reverse. We must determine if the exclusive relationship between compression and zero dissipation continues to hold in such extensions. We pursue this question in forthcoming work.

ACKNOWLEDGMENTS

The authors thank Fabio Anza, David Gier, Richard Feynman, Ryan James, Alexandra Jurgens, Gregory Wimsatt, and Ariadna Venegas-Li for helpful discussions and the Telluride Science Research Center for hospitality during visits. As a faculty member, JPC similarly thanks the Santa Fe Institute. This material is based upon work supported by, or in part by FQXi Grant FQXi-RFP-IPW-1902, the U.S. Army Research Laboratory and the U. S. Army Research Office under contract W911NF-13-1-0390 and grant W911NF-18-1-0028, and via Intel Corporation support of CSC as an Intel Parallel Computing Center.

- [8] A. D. Wyner. The common information of two dependent random variables. *IEEE Trans. Info. Theory.*, 21(2), 1975.
- [9] G. R. Kumar, C. T. Li, and A. El Gamal. Exact common information. 2014. arXiv:1402.0062 [cs.IT]. 1
- [10] R. Landauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Develop.*, 5(3):183–191, 1961. 1, 2
- [11] S. Lloyd. Use of mutual information to decrease entropy: Implications for the second law of thermodynamics. *Phys. Rev. A*, 39(10):5378, 1989.
- [12] A. B. Boyd, D. Mandal, and J. P. Crutchfield. Correlationpowered information engines and the thermodynamics of self-correction. *Phys. Rev. E*, 95:012152, 2017. 2
- [13] A. B. Boyd, D. Mandal, and J. P. Crutchfield. Leveraging environmental correlations: The thermodynamics of requisite variety. J. Stat. Phys., 167(6):1555, 2017. 1, 2
- [14] A. B. Boyd, D. Mandal, and J. P. Crutchfield. Thermodynamics of modularity: Structural costs beyond the Landauer bound. *Phys. Rev. X*, 8(3):031036, 2018. 1, 2, 5, 6

- [15] S. Loomis and J. P. Crutchfield. Thermal efficiency of quantum memory compression. *Phys. Rev. Lett.*, to appear, 2019. arXiv:1911.00998. 1, 2, 5, 7, 8
- [16] T. M. Cover and J. A. Thomas. Elements of Information Theory. Wiley-Interscience, New York, second edition, 2006. 1, 4
- [17] M. Nielsen and I. Chuang. Quantum Computation and Quantum Information. Cambridge University Press, New York, 2010. 1, 2
- [18] D. Petz. Sufficient subalgebras and the relative entropy of states of a von neumann algebra. *Comm. Math. Phys.*, 105(1):123, 1986. 1, 4
- [19] D. Petz. Sufficiency of channels over von neumann algebras. Quart. J. Math. Oxford, 2(39):97, 1988. 4
- [20] M. B. Ruskai. Inequalities for quantum entropy: A review with conditions for equality. J. Math. Phys., 43:4358, 2002.
- [21] P. Hayden, R. Josza, D. Petz, and A. Winter. Structure of states which satisfy strong subadditivity of quantum entropy with equality. *Comm. Math. Phys.*, 246(2):359, 2004.
- [22] M. Mosonyi and D. Petz. Structure of sufficient quantum coarse-grainings. Lett. Math. Phys., 68(1):19, 2004. 5
- [23] M. Mosonyi. Entropy, Information and Structure of Composite Quantum States. PhD thesis, KU Leuven, 2005.
- [24] D. Petz. Quantum Information Theory and Quantum Statistics. Springer, Berlin, Heidelberg, Springer-Verlag Berlin Heidelberg, 2008. 1, 3, 4
- [25] A. Jenčová and D. Petz. Structure of sufficient quantum coarse-grainings. Comm. Math. Phys., 263:259, 2006. 3
- [26] A. Łuczak. Quantum sufficiency in the operator algebra framework. Int. J. Theor. Phys., 53:3423, 2013.
- [27] M. S. Leifer and R. W. Spekkens. A Bayesian approach to compatibility, improvement, and pooling of quantum states. J. Phys. A: Math. Theor., 47:275301, 2014. 1, 3
- [28] B. Baumgartner and H. Narnhoffer. The structures of state space concerning quantum dynamical semigroups. *Rev. Math. Phys.*, 24(2):1250001, 2012. 2, 4, 5
- [29] R. Carbone and Y. Pautrat. Irreducible decompositions and stationary states of quantum channels. *Reports on Math. Phys.*, 77(3):293, 2016. 2, 5
- [30] J. Guan, Y. Feng, and M. Ying. The structure of decoherence-free subsystems. 2018. arXiv:1802.04904 [quant-ph]. 5
- [31] V. V. Albert. Asymptotics of quantum channels: conserved quantities, an adiabatic limit, and matrix product states. *Quantum*, 3:151, 2019. 2, 4, 5
- [32] D. Mandal and C. Jarzynski. Work and information processing in a solvable model of Maxwell's demon. *PNAS*, 109(29):11641, 2012. 2
- [33] A. B. Boyd, D. Mandal, and J. P. Crutchfield. Identifying functional thermodynamics in autonomous Maxwellian ratchets. *New J. Physics*, 18:023049, 2016. 2, 5
- [34] A. J. P. Garner, J. Thompson, V. Vedral, and M. Gu. Thermodynamics of complexity and pattern manipulation. *Phys. Rev. E*, 95:042140, 2017. 6

- [35] A. J. P. Garner. Oracular information and the second law of thermodynamics. 2019. arXiv:1912.03217 [quant-ph]. 2, 6
- [36] J. P. Crutchfield and K. Young. Inferring statistical complexity. *Phys. Rev. Let.*, 63:105–108, 1989. 2
- [37] C. R. Shalizi and J. P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. J. Stat. Phys., 104:817–879, 2001.
- [38] N. F. Travers and J. P. Crutchfield. Equivalence of history and generator ε-machines. arxiv.org:1111.4500 [math.PR]. 6, 13
- [39] J. P. Crutchfield. Between order and chaos. Nature Physics, 8:17–24, 2012. 2
- [40] D. R. Upper. Theory and Algorithms for Hidden Markov Models and Generalized Hidden Markov Models. PhD thesis, University of California, Berkeley, 1997. Published by University Microfilms Intl, Ann Arbor, Michigan. 2, 4, 9
- [41] M. Gu, K. Wiesner, E. Rieper, and V. Vedral. Quantum mechanics can reduce the complexity of classical models. *Nature Comm.*, 3(762), 2012. 2
- [42] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield. Occam's quantum strop: Synchronizing and compressing classical cryptic processes via a quantum channel. *Scientific Reports*, 6(20495), 2016. 2, 7
- [43] P. M. Riechers, J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield. A closed-form shave from Occam's quantum razor: Exact results for quantum compression. *Phys. Lett. A*, 93:052317, 2016. 2
- [44] F. C. Binder, J. Thompson, and M. Gu. A practical, unitary simulator for non-Markovian complex processes. *Phys. Rev. Lett.*, 120:240502, 2017.
- [45] C. Aghamohammadi, J. R. Mahoney, and J. P. Crutchfield. Extreme quantum advantage when simulating classical systems with long-range interaction. *Scientific Reports*, 7(6735), 2017. 2
- [46] C. Aghamohammadi, S. P. Loomis, J. R. Mahoney, and J. P. Crutchfield. Extreme quantum memory advantage for rare-event sampling. *Phys. Rev. X*, 8:011025, 2018.
- [47] J. Thompson, A. J. P. Garner, J. R. Mahoney, J. P. Crutchfield, V. Vedral, and M. Gu. Causal asymmetry in a quantum world. *Phys. Rev. X*, 8:031013, 2018. 2, 7, 8
- [48] S. Loomis and J. P. Crutchfield. Strong and weak optimizations in classical and quantum models of stochastic processes. J. Stat. Phys, 176:1317–1342, 2019. 2, 6, 7, 8
- [49] Q. Liu, T. J. Elliot, F. C. Binder, C. Di Franco, and M. Gu. Optimal stochastic modeling with unitary quantum dynamics. *Phys. Rev. A*, 99:062110, 2019. 2, 7, 12
- [50] G. Crooks. Quantum operation time reversal. *Phys. Rev.* A, 77:034101, 2008. 2, 8, 13
- [51] J. P. Crutchfield, C. J. Ellison, and J. R. Mahoney. Time's barbed arrow: Irreversibility, crypticity, and stored information. *Phys. Rev. Lett.*, 103(9):094101, 2009. 8
- [52] C. J. Ellison, J. R. Mahoney, R. G. James, J. P. Crutchfield, and J. Reichardt. Information symmetries in irreversible processes. *Chaos*, 21:037107, 2011. 2, 8, 15

- [53] O. C. O. Dahlsten, R. Renner, E. Rieper, and V. Vedral. Inadequacy of von Neumann entropy for characterizing extractable work. *New J. Phys.*, 13, 2011. 2, 7
- [54] R. B. Ash. Information Theory. John Wiley and Sons, New York, 1965. 6
- [55] J. Preskill. Quantum computing and the entanglement frontier. arxiv:1203.5813. 10
- [56] N. Travers and J. P. Crutchfield. Exact synchronization for finite-state sources. J. Stat. Phys., 145(5):1181–1201, 2011. 13
- [57] N. Travers and J. P. Crutchfield. Asymptotic synchronization for finite-state sources. J. Stat. Phys., 145(5):1202– 1223, 2011. 13

Supplementary Materials

Thermodynamically-Efficient Local Computation and the Inefficiency of Quantum Memory Compression

Samuel P. Loomis and James P. Crutchfield

The Supplementary Materials give a quick notational overview, review (i) fixed points of quantum channels, reversible computation, and sufficient statistics, (ii) quantum implementations of classical generators and q-machines and their thermodynamic costs, and (iii) provide details on example calculations. The intention is that the main development be accessible while, together with the Supplementary Materials, the full treatment becomes self-contained.

Appendix A: Notation and Basic Concepts

1. State, Measurement, and Channels

Quantum systems are denoted by letters A, B, and C and represented by Hilbert spaces \mathcal{H}_A , \mathcal{H}_B , \mathcal{H}_C , respectively. A system (say A) has state ρ_A that is a positive bounded operator on \mathcal{H}_A . The set of bounded operators on \mathcal{H}_A is denoted by $\mathcal{B}(\mathcal{H}_A)$ and the set of positive bounded operators by $\mathcal{B}_+(\mathcal{H}_A)$.

Measurements W, X, Y, and Z take values in the countable sets $W, \mathcal{X}, \mathcal{Y}$, and \mathcal{Z} , respectively. A measurement X on system A is defined by a set of Kraus operators $\{K^{(x)} : x \in \mathcal{X}\}$ over \mathcal{H}_A that satisfy the completeness relation: $\sum_x K^{(x)\dagger} K^{(x)} = I_A$ —the identity on \mathcal{H}_A . Given state ρ_A and measurement X, the probability of outcome $x \in \mathcal{X}$ is:

$$\Pr\left(X=x;\rho_A\right) := \operatorname{Tr}\left(K^{(x)}\rho_A K^{(x)\dagger}\right)$$

and the state is transformed into:

$$\rho_A|_{X=x} := \frac{K^{(x)}\rho_A K^{(x)\dagger}}{\Pr\left(x;\rho_A\right)} \ .$$

If the Kraus operators $\{K^{(x)} : x \in \mathcal{X}\}$ are orthogonal to one another and projective, then X is called a *projective* measurement. Otherwise, X is called a *positive-operator* valued measure (POVM).

Given a set of projectors $\{\Pi^{(x)}\}\$ on a Hilbert space \mathcal{H}_A , such that $\sum_x \Pi^{(x)} = I_A$ is the identity, we may decompose $\mathcal{H}_A = \bigoplus_x \mathcal{H}_A^{(x)}$ into a direct sum of subspaces $\mathcal{H}_A^{(x)}$, each one corresponding to the support of projector $\Pi^{(x)}$. We will use the direct sum symbol \bigoplus to indicate sums of orthogonal objects. For instance, if $\rho_A^{(x)}$ is a state on $\mathcal{H}_A^{(x)}$ for each x, then $\bigoplus_x \rho_A^{(x)}$ represents a block-diagonal density matrix over \mathcal{H}_A . Similarly, if $K_A^{(y|x)}$ represents a linear operator from $\mathcal{H}_A^{(x)}$ to $\mathcal{H}_A^{(y)}$, then $\bigoplus_{x,y} K_A^{(y|x)}$ represents a linear operator from \mathcal{H}_A to itself, defined by a block matrix with diagonal (x = y) and off-diagonal elements.

If $\mathcal{H}_A = \bigoplus_x \mathcal{H}_A^{(x)}$ and all the $\mathcal{H}_A^{(x)}$ are unitarily equivalent to one another, then we can equivalently write $\mathcal{H}_A = \mathcal{H}_{A_1} \otimes \mathcal{H}_{A_2}$ where \mathcal{H}_{A_1} is equivalent to each of the $\mathcal{H}_A^{(x)}$ and \mathcal{H}_{A_2} is equivalent to $\mathbb{C}^{|\mathcal{X}|}$.

Let ρ_A and σ_A be states in $\mathcal{B}_+(\mathcal{H}_A)$. If, for every measurement X and outcome $x \in \mathcal{X}$, $\Pr(x; \sigma_A) > 0$ implies that $\Pr(x; \rho_A) > 0$, then we write $\rho_A \ll \sigma_A$, indicating that ρ_A is absolutely continuous with respect to σ_A .

To discuss classical systems, we eschew states and instead focus directly on the measurements. In this setting, measurements are called *random variables*. Here, a random variable X is defined over a countable set \mathcal{X} , taking values $x \in \mathcal{X}$ with a specified probability distribution $\Pr(X = x)$. When the variable can be inferred from context, we will simply write the probabilities as $\Pr(x)$.

A random variable X can generate a quantum state ρ_A through an *ensemble* $\left\{\left(\Pr(X=x), \rho_A^{(x)}\right)\right\}$ of potentially nonorthogonal states such that:

$$\rho_A = \sum_{x \in \mathcal{X}} \Pr\left(X = x\right) \rho_A^{(x)} \,.$$

A quantum channel is a linear map $\mathcal{E} : \mathcal{B}(\mathcal{H}_A) \to \mathcal{B}(\mathcal{H}_C)$ that is trace-preserving and completely positive. These conditions are equivalent to requiring that there is a random variable X and a set of Kraus operators $\{K^{(x)} : x \in \mathcal{X}\}$ such that:

$$\mathcal{E}\left(\rho_{A}\right) = \sum_{x} K^{(x)} \rho_{A} K^{(x)\dagger} ,$$

for all ρ_A .¹ To every channel \mathcal{E} there is an *adjoint* \mathcal{E}^{\dagger} :

$$\operatorname{Tr}\left(\mathcal{E}(\rho_A)M\right) = \operatorname{Tr}\left(\rho_A \mathcal{E}^{\dagger}(M)\right)$$

for every state $\rho_A \in \mathcal{B}_+(\mathcal{H}_A)$ and operator $M \in \mathcal{B}(\mathcal{H}_A)$. The adjoint has the form:

$$\mathcal{E}^{\dagger}\left(M\right) = \sum_{x} K^{(x)\dagger} M K^{(x)}$$

Given a subspace $\mathcal{H}_B \subseteq \mathcal{H}_A$, we denote the restriction of \mathcal{E} to \mathcal{H}_B as $\mathcal{E}|_{\mathcal{H}_B} : \mathcal{B}(\mathcal{H}_B) \to \mathcal{B}(\mathcal{H}_C)$. A map $\mathcal{E} : \mathcal{B}(\mathcal{H}_A) \to \mathcal{B}(\mathcal{H}_A)$ is *ergodic* on a space \mathcal{H}_A if there is no proper subspace $\mathcal{H}_B \subseteq \mathcal{H}_A$ such that the range of $\mathcal{E}|_{\mathcal{H}_B}$ is limited to $\mathcal{B}(\mathcal{H}_B)$. \mathcal{E} is ergodic if and only if it has a unique stationary state $\mathcal{E}(\pi_A) = \pi_A$ [29, and references therein].

The classical equivalent of a quantum channel is, of course, the *classical channel* that maps a set \mathcal{X} to \mathcal{Y} according to the conditional probabilities $\Pr(Y = y | X = x)$. This may be alternately represented by the stochastic matrix $\mathbf{T} := (T_{yx})$ such that $T_{yx} = \Pr(Y = y | X = x)$. Stochastic matrices are also defined by the conditions that $T_{yx} \ge 0$ for all $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ and $\sum_{y} T_{yx} = 1$ for all $x \in \mathcal{X}$. A channel from \mathcal{X} to itself is *ergodic* if there is no proper subset $\mathcal{Y} \subset \mathcal{X}$ such that, for $x \in \mathcal{Y}$, $\Pr(X' = x' | X = x) > 0$ only when $x' \in \mathcal{Y}$ as well. This is equivalent to requiring that $T_{x'x} := \Pr(X' = x' | X = x)$ is an irreducible matrix.

A classical channel $\Pr(Y = y|X = x)$ can be induced from a joint distribution $\Pr(x, y)$ by the Bayes' rule $\Pr(y|x) := \Pr(x, y) / \sum_{y} \Pr(x, y)$, so that the joint distribution may be written as $\Pr(x, y) = \Pr(x) \Pr(y|x)$. Given three random variables XYZ, we write X - Y - Z to denote that X, Y, and Z form a *Markov chain*: $\Pr(x, y, z) = \Pr(x) \Pr(y|x) \Pr(y|x) \Pr(z|y)$. The definition is symmetric in that it also implies $\Pr(x, y, z) = \Pr(z) \Pr(y|z) \Pr(x|y)$.

2. Entropy, Information, and Divergence

The uncertainty of a random variable or measurement is considered to be a proxy for its information content. Often, uncertainty is measured by the *Shannon entropy*:

$$H(X) := -\sum_{x \in \mathcal{X}} \Pr(x) \log_2 \Pr(x) ,$$

when it does not diverge. Given a system A with state ρ_A , the quantum uncertainty is the von Neumann entropy:

$$H(A) := -\mathrm{Tr}\left(\rho_A \log_2 \rho_A\right) \ .$$

When the same system may have many states in the given context, this is written directly as $H(\rho_A)$. It is the smallest entropy that can be achieved by taking a projective measurement of system A. The minimum is attained by the measurement that diagonalizes ρ_A .

There are many ways of quantifying the information shared between two random variables X and Y. The most familiar

¹ It is sufficient for our purposes to consider the case where Kraus operators form a countable set. This holds, for instance, as long as we work with separable Hilbert spaces.

$$I(X:Y) := H(X) + H(Y) - H(XY)$$
$$= \sum_{x,y} \Pr(x,y) \log_2 \left(\frac{\Pr(x,y)}{\Pr(x)\Pr(y)}\right) .$$

Corresponding to the mutual information are the *conditional entropies*:

$$H(X|Y) := H(X) - I(X : Y)$$
 and
 $H(Y|X) := H(Y) - I(X : Y)$.

We will also have occasion to use the *conditional mutual information* for three variables:

$$I(X:Y|Z) := \sum_{x,y,z} \Pr(x,y,z) \log_2 \left(\frac{\Pr(x,y|z)}{\Pr(x|z) \Pr(y|z)} \right) \ .$$

Note that an equivalent condition for X - Y - Z is the vanishing of the conditional mutual information: I(X : Z|Y) = 0. For a bipartite quantum state ρ_{AB} the mutual information is usually taken to be the analogous quantity I(A : B) = H(A) + H(B) - H(AB) and the conditional entropies as H(A|B) = H(AB) - H(B) and so on.

We need a way to compare quantum systems and classical random variables. Consider an ensemble $\left\{\left(\Pr(X=x), \rho_A^{(x)}\right)\right\}$ generating average state ρ_A and define its *Holevo quantity*:

$$I(A:X) := H(\rho_A) - \sum_x \Pr(x) H\left(\rho_A^{(x)}\right) .$$

Consider the ensemble-induced state ρ_{AB} defined by:

$$\rho_{AB} = \sum_{x \in \mathcal{X}} \Pr\left(X = x\right) \rho_A^{(x)} \otimes \left|x\right\rangle \left\langle x\right| \ ,$$

where $\{|x\rangle\}$ is an orthogonal basis on \mathcal{H}_B . Such a state is called *quantum-classical*, or *classical-quantum* if the role of A and B are swapped. In the first case, I(A:B) = I(A:X), where I(A:B) is the mutual information of ρ_{AB} and I(A:X) is the Holevo quantity of the ensemble.

Consider random variables X_1 and X_2 over the same set \mathcal{X} with two possible distributions $\Pr(X_1 = x)$ and $\Pr(X_2 = x)$, respectively, and suppose that whenever $\Pr(X_2 = x) > 0$ we have $\Pr(X_1 = x) > 0$ as well. Then, we can quantify the difference between the two distributions with the *relative entropy*:

$$D\left(\Pr(X_1)\|\Pr(X_2)\right) := \sum_{x \in \mathcal{X}} \Pr(X_1 = x) \log_2 \frac{\Pr(X_1 = x)}{\Pr(X_2 = x)} .$$

Similarly, two states ρ_A and σ_A with $\rho_A \ll \sigma_A$ may be compared via the quantum relative entropy:

$$D(\rho_A \| \sigma_A) := \operatorname{Tr} \left(\rho_A \log_2 \rho_A - \rho_A \log_2 \sigma_A \right)$$

The relative entropy leads to allied information-theoretic quantities. For instance:

$$I(X:Y) = D\left(\Pr\left(X,Y\right) \|\Pr\left(X\right)\Pr\left(Y\right)\right) ,$$

and:

$$I(A:B) = D\left(\rho_{AB} \| \rho_A \otimes \rho_B\right) \;.$$

One of the most fundamental information-theoretic inequalities is the monotonicity of the relative entropy under



FIG. S1. Quantum channel decompositions: Conserved measurement X divides the Hilbert space "vertically" via an orthogonal decomposition, $\mathcal{H}_A^{(x=0)} \oplus \mathcal{H}_A^{(x=1)}$, represented above by labels $A^{(x=0)}$ and $A^{(x=1)}$. For each value of x, there is a "horizontal" decomposition into the tensor product of an ergodic subspace and a decoherence-free subspace: $\mathcal{H}_A^{(x)} = \mathcal{H}_{A_E}^{(x)} \otimes \mathcal{H}_{A_{Df}}^{(x)}$, represented respectively by the labels A_E and A_{Df} . According to Theorem 6, information-theoretic reversibility requires storing data in the conserved measurement and the decoherence-free subspace. Any information stored coherently with respect to the conserved measurement (stored in the ergodic subspace) will be irreversibly modified under the channel's action.

transformations. This is the *data processing inequality* [16]. For a quantum channel \mathcal{E} it says:

$$D\left(\mathcal{E}\left(\rho_{A}\right)\|\mathcal{E}\left(\rho_{A}\right)\right) \leq D\left(\rho_{A}\|\sigma_{A}\right) \tag{S1}$$

The condition for equality requires constructing the *Petz recovery channel*:

$$\mathcal{R}_{\sigma}(\cdot) = \sigma^{1/2} \mathcal{E}^{\dagger} \left(\mathcal{E}(\sigma_A)^{-1/2} \cdot \mathcal{E}(\sigma_A)^{-1/2} \right) \sigma^{1/2} .$$
(S2)

It is easy to check that $\mathcal{R}_{\sigma} \circ \mathcal{E}(\sigma_A) = \sigma_A$. A markedly useful result is that $D(\mathcal{E}(\rho_A) || \mathcal{E}(\rho_A)) = D(\rho_A || \sigma_A)$ if and only if $\mathcal{R}_{\sigma} \circ \mathcal{E}(\rho_A) = \rho_A$ as well [18, 19].

Two other forms of the data processing inequality are useful to note here. The first uses another quantity for measuring distance between states called the *fidelity*:

$$F(\rho,\sigma) := \left(\operatorname{Tr} \left[\sqrt{\sqrt{\rho}\sigma\sqrt{\rho}} \right] \right)^2 \,. \tag{S3}$$

It takes value F = 0 when the states ρ and σ are completely orthogonal and value F = 1 if and only if $\rho = \sigma$. The data processing inequality for fidelity states that for any quantum channel \mathcal{E} :

$$F(\mathcal{E}(\rho), \mathcal{E}(\sigma)) \ge F(\rho, \sigma)$$
 . (S4)

This is yet another way of saying that states map closer together under a quantum channel \mathcal{E} .

The second form arises from applying Eq. (S1) to the mutual information. Let $\mathcal{E}_A : \mathcal{B}(\mathcal{H}_A) \to \mathcal{B}(\mathcal{H}_C)$ be a quantum channel. The local operation $\mathcal{E}_A \otimes I_B$ maps a bipartite system AB to CB. The data processing inequality Eq. (S1) implies that we have $I(C:B) \leq I(A:B)$.

In these terms, our thermodynamic efficiency goal— $\Delta S_{\text{loc}} = 0$ —translates into determining conditions for equality—I(C:B) = I(A:B)—using the Petz recovery channel and channel fixed points.

Appendix B: Fixed Points and Reversible Computation

To understand our result on local channels, an illustrative starting point is a key result on fixed points of quantum channels [28, 29] that leads to a natural decomposition, as illustrated in Fig. S1.

Theorem 5 (Channel and Stationary State Decomposition). Suppose $\mathcal{E} : \mathcal{B}(\mathcal{H}_A) \to \mathcal{B}(\mathcal{H}_A)$ is a quantum channel, Hilbert space \mathcal{H}_A has a transient subspace \mathcal{H}_T , and there is a projective measurement $X = \{\Pi^{(x)}\}$ on \mathcal{H}_T^{\perp} with countable outcomes \mathcal{X} , such that $\mathcal{H}_A = \mathcal{H}_T \oplus \left(\bigoplus_x \mathcal{H}_A^{(x)}\right)$, where $\mathcal{H}_A^{(x)}$ is the support of $\Pi^{(x)}$. Then:

- 1. Subspace $\mathcal{H}_{A}^{(x)}$ is preserved by \mathcal{E} , in that for all $\rho \in \mathcal{B}\left(\mathcal{H}_{A}^{(x)}\right)$, we have $\mathcal{E}\left(\rho\right) \in \mathcal{B}\left(\mathcal{H}_{A}^{(x)}\right)$.
- 2. Subspace $\mathcal{H}_A^{(x)}$ further decomposes into an ergodic subspace $\mathcal{H}_{A_{\rm E}}^{(x)}$ and decoherence-free subspace $\mathcal{H}_{A_{\rm Df}}^{(x)}$:

$$\mathcal{H}_A^{(x)} = \mathcal{H}_{A_{\mathrm{E}}}^{(x)} \otimes \mathcal{H}_{A_{\mathrm{Df}}}^{(x)}$$

such that the Kraus operators of $\mathcal{E}|_{\mathcal{H}_T^{\perp}}$ decompose as [30]:

$$K^{(\alpha)} = \bigoplus_{x \in \mathcal{X}} K^{(\alpha,x)}_{A_{\mathrm{E}}} \otimes U^{(x)}_{A_{\mathrm{Df}}} \tag{S1}$$

and the map $\mathcal{E}_{A_{\mathrm{E}}}^{(x)}(\cdot) := \sum_{\alpha} K_{A_{\mathrm{E}}}^{(\alpha,x)} \cdot K_{A_{\mathrm{E}}}^{(\alpha,x)\dagger}$ has a unique invariant state $\pi_{A_{\mathrm{E}}}^{(x)}$. 3. Any subspace of \mathcal{H} satisfying the above two properties is, in fact, $\mathcal{H}_{A}^{(x)}$ for some $x \in \mathcal{X}$.

Furthermore, if ρ_A is any invariant state—that is, $\rho_A = \mathcal{E}_A(\rho_A)$ —then it decomposes as:

$$\rho_A = \bigoplus_{x \in \mathcal{X}} \Pr\left(x\right) \pi_{A_{\rm E}}^{(x)} \otimes \rho_{A_{\rm Df}}^{(x)} , \qquad (S2)$$

for any distribution $\Pr(x)$ and state $\rho_{A_{\mathrm{Df}}}^{(x)}$ satisfying $U_{A_{\mathrm{Df}}}^{(x)}\rho_{A_{\mathrm{Df}}}^{(x)}U_{A_{\mathrm{Df}}}^{(x)\dagger} = \rho_{A_{\mathrm{Df}}}^{(x)}$.

Figure S1 gives the geometric structure implied by the theorem. The ergodic subspace of quantum channel \mathcal{E} has two complementary decompositions. First, there is an orthogonal decomposition $\mathcal{H}_{\mathrm{T}}^{\perp} = \bigoplus_{x} \mathcal{H}_{A}^{(x)}$ induced by a projective measurement X whose values are conserved by \mathcal{E} 's action. This conservation is decoherent: only states compatible with X are stationary under the action of \mathcal{E} . X is called the *conserved measurement* of \mathcal{E} [31]. Then, each $\mathcal{H}_{A}^{(x)}$ has a tensor decomposition $\mathcal{H}_{A}^{(x)} = \mathcal{H}_{A_{\mathrm{E}}}^{(x)} \otimes \mathcal{H}_{A_{\mathrm{Df}}}^{(x)}$ into an *ergodic* (E) and a *decoherence-free* (Df) part. The decoherence-free subspace $\mathcal{H}_{A_{\mathrm{Df}}}^{(x)}$ undergoes only a unitary transformation [30]. The ergodic part $\mathcal{H}_{A_{\mathrm{E}}}^{(x)}$ is irreducibly mixed such that there is a single stationary state.

This result's contribution here is to aid in identifying when the data-processing inequality saturates. That is, using Thm. 5 and the Petz recovery channel, we derive necessary and sufficient constraints on the structures of ρ_A , σ_A , and \mathcal{E}_A for determining when $D(\mathcal{E}(\rho_A) || \mathcal{E}(\sigma_A)) = D(\rho_A || \sigma_A)$.

To achieve this, we recall a previously known result [22, 23], showing that it can be derived using only the Petz recovery map and Thm. 5. The immediate consequence is a novel proof.

Theorem 6 (Reversible Information Processing). Suppose for two states ρ_A and σ_A on a Hilbert space \mathcal{H}_A and a quantum channel $\mathcal{E} : \mathcal{B}(\mathcal{H}_A) \to \mathcal{B}(\mathcal{H}_A)$, we have:

$$D\left(\rho_A \| \sigma_A\right) = D\left(\rho_C \| \sigma_C\right) \;,$$

where $\rho_C = \mathcal{E}_A(\rho_A)$ and $\sigma_C = \mathcal{E}_A(\sigma_A)$. Then there is a measurement X on A with countable outcomes X and orthogonal decompositions $\mathcal{H}_{\mathcal{E}} = \bigoplus_x \mathcal{H}_A^{(x)}$ and $\mathcal{H}_C^{(x)}$ such that:

1. For all
$$\rho \in \mathcal{B}\left(\mathcal{H}_{A}^{(x)}\right)$$
 and $\mathcal{E}\left(\rho\right) \in \mathcal{B}\left(\mathcal{H}_{C}^{(x)}\right)$, the mapping $\mathcal{E}|_{\mathcal{H}_{A}^{(x)}}$ onto $\mathcal{B}\left(\mathcal{H}_{C}^{(x)}\right)$ is surjective.

2. Subspaces $\mathcal{H}^{(x)}_A$ further decompose into:

$$\begin{aligned} \mathcal{H}_A^{(x)} &= \mathcal{H}_{A_{\rm E}}^{(x)} \otimes \mathcal{H}_{A_{\rm Df}}^{(x)} \text{ and } \\ \mathcal{H}_C^{(x)} &= \mathcal{H}_{C_{\rm E}}^{(x)} \otimes \mathcal{H}_{C_{\rm Df}}^{(x)} , \end{aligned}$$

so that $\mathcal{H}_{C_{\mathrm{E}}}^{(x)}$ and $\mathcal{H}_{A_{\mathrm{Df}}}^{(x)}$ are unitarily equivalent and the Kraus operators decompose as:

$$L^{(\alpha)} = \bigoplus_{x \in \mathcal{X}} L^{(\alpha,x)}_{A_{\mathrm{E}}} \otimes U^{(x)}_{A_{\mathrm{Df}}} .$$
(S3)

Furthermore, states ρ_A and σ_A satisfy:

$$\rho_A = \sum_{x \in \mathcal{X}} \Pr\left(x;\rho\right) \pi_{A_{\rm E}}^{(x)} \otimes \rho_{A_{\rm Df}}^{(x)} \tag{S4a}$$

$$\sigma_A = \sum_{x \in \mathcal{X}} \Pr\left(x; \sigma\right) \pi_{A_{\mathrm{E}}}^{(x)} \otimes \sigma_{A_{\mathrm{Df}}}^{(x)}$$
(S4b)

for some $\pi_{A_{\rm F}}^{(x)}$. And, their images are:

$$\rho_C = \sum_{x \in \mathcal{X}} \Pr\left(x;\rho\right) \pi_{C_{\rm E}}^{(x)} \otimes \rho_{C_{\rm Df}}^{(x)} \tag{S5a}$$

$$\sigma_C = \sum_{x \in \mathcal{X}} \Pr\left(x; \sigma\right) \pi_{C_{\mathrm{E}}}^{(x)} \otimes \sigma_{C_{\mathrm{Df}}}^{(x)} , \qquad (S5b)$$

where $\pi_{C_{\rm E}}^{(x)} = \mathcal{E}_{A_{\rm E}}^{(x)} \left(\pi_{A_{\rm E}}^{(x)} \right), \ \rho_{C_{\rm Df}}^{(x)} = U_{A_{\rm Df}}^{(x)} \rho_{A_{\rm Df}}^{(x)} U_{A_{\rm Df}}^{(x)\dagger}, \ and \ \sigma_{C_{\rm Df}}^{(x)} = U_{A_{\rm Df}}^{(x)} \sigma_{A_{\rm Df}}^{(x)} U_{A_{\rm Df}}^{(x)\dagger}.$

We know that $\mathcal{N}_{\sigma} := \mathcal{R}_{\sigma} \circ \mathcal{E}$ must have both ρ_A and σ_A as stationary distributions. Let X be the conserved measurement of \mathcal{N}_{σ} . It induces the decompositions $\mathcal{H}_A = \mathcal{H}_T \oplus \left(\bigoplus_x \mathcal{H}_A^{(x)}\right)$ and $\mathcal{H}_A^{(x)} = \mathcal{H}_{A_E}^{(x)} \otimes \mathcal{H}_{A_{Df}}^{(x)}$, as well as the state decompositions Eq. (S4).

Now, we leverage the fact that $\mathcal{N}_{\sigma}|_{\mathcal{H}_{\pi}^{\perp}}$ has Kraus decomposition of the form Eq. (S1). The net effect is summed up by two constraints:

- 1. For each α , $K^{(\alpha)}$ maps each $\mathcal{H}_{A}^{(x)}$ to itself. 2. Let $K^{(\alpha,x)}$ be the block of $K^{(\alpha)}$ restricted to $\mathcal{H}_{A}^{(x)}$. Let M be a complete projective measurement on $\mathcal{H}_{A_{\text{Df}}}^{(x)}$ with basis $\{|m\rangle\}$ and define the transformed basis $\{|\tilde{m}\rangle = U^{(x)} |m\rangle\}$. Now, let $\mathcal{H}_{A}^{(x,m)} = \{|\psi\rangle \otimes |m\rangle : |\psi\rangle \in \mathcal{H}_{A_{\text{E}}}^{(x)}\}$ and similarly for $\mathcal{H}_{A}^{(x,\widetilde{m})}$. Then, $K^{(z,x)}$ maps $\mathcal{H}_{A}^{(x,m)}$ to $\mathcal{H}_{A}^{(x,\widetilde{m})}$. This holds for any measurement M.

Proving Eq. (S3) requires these two constraints. Each is a form of distinguishability criterion for the total channel $\mathcal{N}_{\sigma}|_{\mathcal{H}_{\pi}^{\perp}}$. Since \mathcal{N}_{σ} can tell certain orthogonal outcomes apart, so too must \mathcal{E} . Or else, \mathcal{R}_{σ} would "pull apart" nonorthogonal states. However, this is impossible for a quantum channel. By formally applying this notion to constraints 1 and 2 above, we recover Eq. (S3).

Let $L^{(\alpha)}$ be the Kraus operators of $\mathcal{E}|_{\mathcal{H}^{\perp}_{\tau}}$. Then $K^{(\alpha)} = \sigma^{1/2} L^{(\alpha)\dagger} \mathcal{E}(\sigma)^{1/2} L^{(z)}$. Now, if $L^{(\alpha)}$ did not map each $\mathcal{H}^{(x)}_A$ to some orthogonal subspace $\mathcal{H}_C^{(x)}$, then for some distinct x and x' there would be $|\psi\rangle \in \mathcal{H}_A^{(x)}$ and $|\phi\rangle \in \mathcal{H}_C^{(x')}$ such that $F(\mathcal{E}(|\psi\rangle\langle\psi|), \mathcal{E}(|\phi\rangle\langle\phi|)) > 0$; recall Eq. (S3) defines fidelity. However, we must also have $F(\mathcal{N}_{\sigma}(|\psi\rangle\langle\psi|), \mathcal{N}_{\sigma}(|\phi\rangle\langle\phi|)) = 0$ 0, which is impossible by Eq. (S4) since as applying \mathcal{R}_{σ} cannot reduce fidelity. So, it must be the case that $L^{(\alpha)}$ maps each $\mathcal{H}_A^{(x)}$ to some orthogonal subspace $\mathcal{H}_C^{(x)}$. This proves Claim 1 in Thm. 6.

Let $L^{(\alpha,x)}$ be the block of $L^{(z)}$ restricted to $\mathcal{H}_A^{(x)}$. Further, let M be a complete measurement on $\mathcal{H}_{A_{\mathrm{Df}}}^{(x)}$. Then $\mathcal{E}|_{\mathcal{H}_*^{(x)}}$. must map each $\mathcal{H}_{A}^{(x,m)}$ onto orthogonal subspaces $\mathcal{H}_{C}^{(x,m)}$ of $\mathcal{H}_{C}^{(x)}$, lest $\mathcal{N}_{\sigma}|_{\mathcal{H}_{A}^{(x)}}$ could not map each $\mathcal{H}_{A}^{(x,m)}$ to orthogonal spaces $\mathcal{H}_{A}^{(x,\widetilde{m})}$. This follows from fidelity, as in the previous paragraph.

Now, let $L^{(\alpha,x,m)}: \mathcal{H}_{A_{\mathrm{E}}}^{(x)} \to \mathcal{H}_{C}^{(x,m)}$, so that:

$$L^{(\alpha,x)} = \bigoplus_{m} L^{(\alpha,x,m)} \otimes \langle m |$$
.

Let N be another complete measurement on $\mathcal{H}_{A_{Df}}^{(x)}$ such that $|n\rangle = \sum_{m} w_{m,n} |m\rangle$, with $w_{m,n}$ a unitary matrix. And, let $n, n' \in \mathcal{N}$ be distinct. We have for any $|\psi\rangle$ that:

$$\mathcal{E}^{(x)}\left(\left|\psi\right\rangle\left\langle\psi\right|\otimes\left|n\right\rangle\left\langle n\right|\right) = \sum_{\alpha} \bigoplus_{m,m'} w_{m',n} w_{m,n}^* L^{(\alpha,x,m)}\left|\psi\right\rangle\left\langle\psi\right| L^{(\alpha,x,m')\dagger} \text{ and}$$
$$\mathcal{E}^{(x)}\left(\left|\psi\right\rangle\left\langle\psi\right|\otimes\left|n'\right\rangle\left\langle n'\right|\right) = \sum_{\alpha} \bigoplus_{m,m'} w_{m',n'} w_{m,n'}^* L^{(\alpha,x,m)}\left|\psi\right\rangle\left\langle\psi\right| L^{(\alpha,x,m')\dagger} .$$

Now, it must be that $\mathcal{E}^{(x)}(|\psi\rangle \langle \psi| \otimes |n\rangle \langle n|)$ and $\mathcal{E}^{(x)}(|\psi\rangle \langle \psi| \otimes |n'\rangle \langle n'|)$ are orthogonal. However:

$$\operatorname{Tr}\left[\mathcal{E}^{(x)}\left(\left|\psi\right\rangle\left\langle\psi\right|\otimes\left|n\right\rangle\left\langle n\right|\right)\mathcal{E}^{(x)}\left(\left|\psi\right\rangle\left\langle\psi\right|\otimes\left|n'\right\rangle\left\langle n'\right|\right)\right]=\sum_{\alpha,\alpha'}\sum_{m}w_{m,n}w_{m,n'}^{*}\left\langle\psi\right|L^{(\alpha,x,m)\dagger}L^{(\alpha,x,m)}\left|\psi\right\rangle \ .$$

This vanishes for arbitrary N and $|\psi\rangle$ only if $L^{(\alpha|x,m)\dagger}L^{(\alpha'|x,m)}$ is independent of m for each α and α' . This implies that $L^{(\alpha|x,m)} = W^{(m)}L_{\rm E}^{(\alpha|x)}$ for some unitary $W^{(m)}$.

All of which leads one to conclude that the $\mathcal{H}_{C}^{(x,m)}$ for each m must be unitarily equivalent. And so, the decomposition $\mathcal{H}_{C}^{(x)} = \bigoplus_{m} \mathcal{H}_{C}^{(x)}$ instead becomes a tensor product decomposition $\mathcal{H}_{C}^{(x)} = \mathcal{H}_{C_{\rm E}}^{(x)} \otimes \mathcal{H}_{C_{\rm Df}}^{(x)}$ and further that $L^{(\alpha,x)} = L_{A_{\rm E}}^{(\alpha,x)} \otimes V_{A_{\rm Df}}^{(x)}$.

Finally, the constraints Eq. (S5) follow from the form of \mathcal{E} and Eq. (S4). \Box

Theorem 6's main implication is that, when a channel \mathcal{E} acts, information stored in the conserved measurement and in the decoherence-free subspaces is recoverable. Two states that differ in terms of the conserved measurement and the decoherence-free subspaces remain different and do not grow more similar under \mathcal{E} 's action. Conversely, information stored in measurements not compatible with the conserved measurement or stored in the ergodic subspaces is irreversibly garbled by \mathcal{E} .

The next section uses this decomposition to study how locally acting channels impact correlations between subsystems. This directly drives the thermodynamic efficiency of local operations. Namely, for thermodynamic efficiency correlations must be stored specifically in the conserved measurement and decoherence-free subspaces of the local channel.

Appendix C: Quantum Sufficient Statistics and Reversible Local Operations

Let's review the definition of sufficient statistic in the main body. Let ρ_{AB} be a bipartite quantum state. A maximal local commuting measurement (MLCM) of A for B is any local measurement X on system A such that:

$$\rho_{AB} = \bigoplus_{x} \Pr(X = x) \rho_{AB}^{(x)} ,$$

where:

$$\Pr(X = x) = \operatorname{Tr}\left((\Pi_X^{(x)} \otimes I_B)\rho_{AB}\right)$$

and:

$$\Pr(X=x)\rho_{AB}^{(x)} = (\Pi_X^{(x)} \otimes I_B)\rho_{AB}(\Pi_X^{(x)} \otimes I_B)$$

$$\rho_{AB}^{(x)} \neq \sum_{y} (\Pi_Y^{(y)} \otimes I_B) \rho_{AB}^{(x)} (\Pi_Y^{(y)} \otimes I_B) \ .$$

We call the states $\left\{\rho_{AB}^{(x)}\right\}$ quantum correlation atoms.

Proposition 3 (MLCM uniqueness). Given a state ρ_{AB} , there is a unique MLCM of A for B. Suppose there were two distinct MLCMs, X and Y. Then:

$$\rho_{AB} = \sum_{y} \Pr(X = x) (\Pi_Y^{(y)} \otimes I_B) \rho_{AB} (\Pi_Y^{(y)} \otimes I_B) .$$

This can be written as:

$$\rho_{AB} = \bigoplus_{x} \sum_{y} \Pr(X = x) (\Pi_Y^{(y)} \otimes I_B) \rho_{AB}^{(x)} (\Pi_Y^{(y)} \otimes I_B) .$$

However, this means for each x:

$$(\Pi_Y^{(y)} \otimes I_B)\rho_{AB}^{(x)}(\Pi_Y^{(y)} \otimes I_B) = \rho_{AB}^{(x)} .$$

So, X is not a MLCM, giving a contradiction. \Box

It will be helpful in our study of quantum generators to have the following fact as well:

Proposition 4 (MLCM for a classical-quantum state). Given a classical-quantum state:

$$\rho_{AB} := \sum_{x} \Pr(x) |x\rangle \langle x| \otimes \rho_{B}^{(x)} ,$$

the MLCM is the most refined measurement Θ such that:

$$\rho_B^{(x)} = \sum_{\theta} \Pi^{(\theta)} \rho_B^{(x)} \Pi^{(\theta)}$$

for all x.

Given that Θ is a commuting local measurement, the question is whether it is maximal. If it is not maximal, though, there is a refinement Y that is also a commuting local measurement. By Θ 's definition, there is an x such that $\rho_B^{(x)} \neq \sum_y \Pi^{(y)} \rho_B^{(x)} \Pi^{(y)}$. This implies $\rho_{AB} \neq \sum_y (I \otimes \Pi^{(y)}) \rho_{AB}(I \otimes \Pi^{(y)})$, contradicting the assumption that Y is commuting local. \Box

Now, let ρ_{AB} be a bipartite quantum state and let X be the MLCM of A for B. We define the correlation equivalence relation $x \sim x'$ over values of X where $x \sim x'$ if and only if $\rho_{AB}^{(x)} = (U \otimes I_B)\rho_{AB}^{(x')}(U^{\dagger} \otimes I_B)$ for a local unitary U. Finally, we define the Minimal Local Sufficient Statistic (MLSS) $[X]_{\sim}$ as the equivalence class $[x]_{\sim} := \{x' : x' \sim x\}$ generated by the relation \sim between correlation atoms. Thus, our notion of sufficiency of A for B is to find the most informative local measurement and then coarse-grain its outcomes by unitary equivalence over their correlation atoms. The correlation atoms and the MLSS $[X]_{\sim}$ together describe the correlation structure of the system AB. Using this definition and Theorem 6 we can prove our result on reversible local operations.

Theorem 7 (Reversible local operations). Let ρ_{AB} be a bipartite quantum state and let $\mathcal{E}_A \otimes I_B$ be a local operation with $\mathcal{E}_A : \mathcal{B}(\mathcal{H}_A) \to \mathcal{B}(\mathcal{H}_C)$. Suppose X is the MLCM of ρ_{AB} and Y, that of $\rho_{CB} = \mathcal{E}_A \otimes I_B(\rho_{AB})$. The decomposition into correlation atoms is:

$$\rho_{AB} = \bigoplus_{x} \Pr_{A}(x) \rho_{AB}^{(x)} and$$
(S1)

$$\rho_{CB} = \bigoplus_{y} \Pr_{C}(y) \rho_{CB}^{(y)} .$$
(S2)

Then, I(A:B) = I(C:B) if and only if \mathcal{E}_A can be expressed by Kraus operators of the form:

$$K^{(\alpha)} = \bigoplus_{x,y} e^{i\phi_{xy\alpha}} \sqrt{\Pr(y,\alpha|x)} U^{(y|x)} , \qquad (S3)$$

where $\phi_{xy\alpha}$ is any arbitrary phase and $\Pr(y, \alpha|x)$ is a stochastic channel that is nonzero only when $\rho_{AB}^{(x)}$ and $\rho_{CB}^{(y)}$ are equivalent up to a local unitary operation $U^{(y|x)}$ that maps $\mathcal{H}_A^{(x)}$ to $\mathcal{H}_C^{(y)}$.

We can apply the Reversible Information Processing Theorem (Thm. 6) from the previous section here. This demands that there be a measurement X and a decomposition of the Hilbert space $\mathcal{H}_A = \mathcal{H}_{A_E} \otimes \mathcal{H}_{A_{Df}}$ such that:

$$\rho_{AB} = \sum_{x} \Pr_{A}(x) \rho_{(AB)_{\rm E}}^{(x)} \otimes \rho_{(AB)_{\rm Df}}^{(x)}$$
$$\rho_{A} \otimes \rho_{B} = \sum_{x} \Pr_{A}(x) \rho_{(AB)_{\rm E}}^{(x)} \otimes \left(\rho_{A_{\rm Df}}^{(x)} \otimes \rho_{B_{\rm Df}}^{(x)}\right)$$

such that $\mathcal{E}_A \otimes I_B$ conserves measurement X and acts decoherently on $(AB)_E$ and coherently on $(AB)_{Df}$. However, the local nature of $\mathcal{E}_A \otimes I_B$ makes it clear we can simplify this decomposition to:

$$\rho_{AB} = \sum_{x} \Pr_{A}(x) \rho_{A_{\rm E}}^{(x)} \otimes \left(\rho_{A_{\rm Df}B}^{(x)}\right)$$
$$\rho_{A} \otimes \rho_{B} = \sum_{x} \Pr_{A}(x) \rho_{A_{\rm E}}^{(x)} \otimes \left(\rho_{A_{\rm Df}}^{(x)} \otimes \rho_{B}\right) ,$$

where \mathcal{E}_A conserves the *local* measurement X on A and acts decoherently on $A_{\rm E}$ and acts as a local unitary $U_{A_{\rm Df}} \otimes I_B$ on $A_{\rm Df}B$.

Suppose now, however, that given X the variable Y_x is the diagonalizing measurement of $\rho_{A_{\rm E}}^{(x)}$ and Z_x is the MLCM of $\rho_{A_{\rm Df}B}^{(x)}$. The joint measurement XY_XZ_X —where X is measured first and then the other two measurements are determined with knowledge of its outcome—is the MLCM of ρ_{AB} . Note that for any x and z, the outcomes (x, y, z) and (x, y', z) are correlation equivalent: measurement Y_X is completely decoupled from system B. Then, the MLSS $\Sigma := [XY_XZ_X]_B$ is simply a function of X and Z_X .

Since XZ_X is conserved by the action of $\mathcal{E}_A \otimes I$ —where X is the conserved measurement, while Z_X is preserved through the unitary evolution—the MLSS Σ must be preserved and each correlation atom is transformed only by a local unitary. This results in the form Eq. (S3).

This proves that I(A:B) = I(C:B) implies Eq. (S3). The converse is straightforward to check. Let $\Sigma = [X]_B$ and let $I(A:B|\Sigma = \sigma)$ be the mutual information of $\rho_{AB}^{(x)}$ for any $x \in \sigma$. (This is the same for all such x by local unitary equivalence.) Then:

$$I(A:B) = \sum_{\sigma} \Pr(\Sigma = \sigma) I(A:B|\Sigma = \sigma)$$
.

Similarly, let $\Sigma' = [Y]_B$ and let $I(C: B|\Sigma' = \sigma')$ be the mutual information of $\rho_{CB}^{(y)}$ for any $y \in \sigma'$; then:

$$I(C:B) = \sum_{\sigma'} \Pr\left(\Sigma' = \sigma'\right) I\left(C:B|\Sigma' = \sigma'\right) \ .$$

Since \mathcal{E} isomorphically maps each σ to a unique σ' , such that $I(A:B|\Sigma=\sigma) = I(C:B|\Sigma'=\sigma')$ by unitary equivalence, we must have I(A:B) = I(C:B). \Box

Appendix D: Classical Generators

Recall the definitions from the main body. A classical generator is the physical implementation of a hidden Markov model (HMM) [40] $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$, where (here) \mathcal{S} is countable, \mathcal{X} is finite, and for each $x \in \mathcal{X}, \mathbf{T}^{(x)}$ is a matrix

with values given by a stochastic channel from \mathcal{S} to $\mathcal{S} \times \mathcal{Y}$, $T_{s's}^{(x)} := \Pr_{\mathfrak{G}}(s', x|s)$. We define generators to use *recurrent* HMMs, which means the total transition matrix $T_{s's} := \sum_{x} T_{s's}^{(x)}$ is irreducible. In this case, there is a unique stationary distribution $\pi_{\mathfrak{G}}(s)$ over states \mathcal{S} satisfying $\pi_{\mathfrak{G}}(s) > 0$, $\sum_{s} \pi_{\mathfrak{G}}(s) = 1$, and $\sum_{s} T_{s's} \pi_{\mathfrak{G}}(s) = \pi_{\mathfrak{G}}(s')$.

The following probability distributions are particularly relevant:

$$\Pr_{\mathfrak{G}}(x_{1}\dots x_{\ell}) := \sum_{s_{0}\dots s_{\ell}\in\mathcal{S}^{\ell+1}} T_{s_{\ell}s_{\ell-1}}^{(x_{\ell})}\dots T_{s_{1}s_{0}}^{(x_{1})}\pi(s_{0}) .$$
$$\Pr_{\mathfrak{G}}(x_{1}\dots x_{t}, s_{t}) := \sum_{s_{0}\dots s_{t-1}\in\mathcal{S}^{t}} T_{s_{t}s_{t-1}}^{(x_{t})}\dots T_{s_{1}s_{0}}^{(x_{1})}\pi_{\mathfrak{G}}(s_{0})$$

Also, recall the definition of a mergeable partition and retrodictive equivalence. Each partition $\mathcal{P} = \{\mathcal{P}_{\theta}\}$ of \mathcal{S} has a corresponding merged generator $\mathfrak{G}^{\mathcal{P}} = (\mathcal{P}, \mathcal{X}, \{\widetilde{T}_{\theta'\theta}^{(x)}\})$ with transition dynamics given by:

$$\widetilde{T}_{\theta'\theta}^{(x)} := \sum_{\substack{s' \in \mathcal{P}_{\theta'} \\ s \in \mathcal{P}_{\theta}}} \Pr_{\mathfrak{G}}\left(s', x|s\right) \pi(s|\theta) \ ,$$

where $\pi_{\mathfrak{G}}(s|\theta) = \pi_{\mathfrak{G}}(s)/\pi_{\mathfrak{GP}}(\theta)$ and $\pi_{\mathfrak{GP}}(\theta) = \sum_{s \in \mathcal{P}_{\theta}} \pi_{\mathfrak{G}}(s)$. A partition $\{\mathcal{P}_{\theta}\}$ is *mergeable* if the merged generator generates the same process as the original.

One example of a partition is the retrodictive equivalence class. Two states $s, s' \in S$ are considered equivalent $s \sim s'$ if $\Pr_{\mathfrak{G}}(x_1 \ldots x_t | s) = \Pr_{\mathfrak{G}}(x_1 \ldots x_t | s')$ for all words $x_1 \ldots x_t$. The equivalence class $[S_t]_{\sim}$ is the sufficient statistic of S_t for predicting the past symbols $X_1 \ldots X_t$. The set $\mathcal{P}_{\sim} := \{[s]_{\sim} : s \in S\}$ of equivalence classes is a partition on S that we index by σ .

Proposition 5. Given a generator $(S, \mathcal{X}, \{T_{s's}^{(x)}\})$, the partition $\mathcal{P} := \mathcal{P}_{\sim}$ induced by retrodictive equivalence is mergeable.

To see this, we follow a proof by induction. We first suppose that the distribution $\Pr_{\mathfrak{GP}}(x_1 \dots x_t, \sigma_t)$ is equal to the distribution:

$$\Pr_{\mathfrak{G}}(x_1 \dots x_t, \sigma_t) := \sum_{s_t \in \mathcal{P}_{\sigma_t}} \Pr_{\mathfrak{G}}(x_1 \dots x_t, s_t) ,$$

for some t. Then, noting that:

$$\Pr_{\mathfrak{G}}(x_1 \dots x_{t+1}, s_{t+1}) = \sum_{s_t} \Pr_{\mathfrak{G}}(x_1 \dots x_t | s_t) \pi_{\mathfrak{G}}(s_t) T_{s_{t+1}s_t}^{(x_{t+1})}$$
$$= \sum_{\substack{\sigma_t \\ s_t \in \mathcal{P}_{\sigma_t}}} \Pr_{\mathfrak{G}}(x_1 \dots x_t | \sigma_t) \pi_{\mathfrak{G}}(\sigma_t) \pi_{\mathfrak{G}}(s_t | \sigma_t) T_{s_{t+1}s_t}^{(x_{t+1})}$$

,

we have:

$$\Pr_{\mathfrak{G}} (x_1 \dots x_{t+1}, \sigma_{t+1}) = \sum_{\substack{\sigma_t \\ s_t \in \mathcal{P}_{\sigma_t} \\ s_{t+1} \in \mathcal{P}_{\sigma_{t+1}}}} \Pr_{\mathfrak{G}^{\mathcal{P}}} (x_1 \dots x_{t-1}, \sigma_{t-1}) \widetilde{T}_{\sigma_{t+1}\sigma_t}^{(x_{t+1})}$$
$$= \Pr_{\mathfrak{G}^{\mathcal{P}}} (x_1 \dots x_{t+1}, \sigma_{t+1}) .$$

Now, for t = 1, we have:

$$\Pr_{\mathfrak{G}^{\mathcal{P}}}(x_{1},\sigma_{1}) = \sum_{\sigma_{0}} \pi_{\mathfrak{G}^{\mathcal{P}}}(\sigma_{0}) \widetilde{T}_{\sigma_{1}\sigma_{0}}^{(x_{1})}$$
$$= \sum_{\sigma_{0}} \sum_{\substack{s_{0} \in \mathcal{P}_{\sigma_{0}}\\s_{1} \in \mathcal{P}_{\sigma_{1}}}} \pi_{\mathfrak{G}^{\mathcal{P}}}(\sigma_{0}) \pi_{\mathfrak{G}}(s_{0}|\sigma_{0}) T_{s_{1}s_{0}}^{(x_{1})}$$
$$= \sum_{s_{1}} \Pr_{\mathfrak{G}}(x_{1},s_{1}) .$$

Then, by induction, $\Pr_{\mathfrak{G}^{\mathcal{P}}}(x_1 \dots x_t, \sigma_t) = \Pr_{\mathfrak{G}}(x_1 \dots x_t, \sigma_t)$ for all t. By summing over σ_t , we have $\Pr_{\mathfrak{G}^{\mathcal{P}}}(x_1 \dots x_t) = \Pr_{\mathfrak{G}}(x_1 \dots x_t)$ for all t as well. \Box

Using this result and the Reversible Local Operations Theorem, we can establish the following.

Theorem 8. A generator $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ satisfies $I(S_t : X_1 \dots X_t) = I(S_{t+1}X_{t+1} : X_1 \dots X_t)$ for all t if and only if the retrodictively state-merged generator $\mathfrak{G}^{\mathcal{P}} = (\mathcal{P}_{\sim}, \mathcal{X}, \{\widetilde{T}_{\sigma'\sigma}^{(x)}\})$ satisfies $\widetilde{T}_{\sigma'\sigma}^{(x)} \propto \delta_{\sigma,f(\sigma',x)}$ for some function $f : \mathcal{S} \times \mathcal{X} \to \mathcal{S}$.

To see this, recall from Cor. 1 that $I(S_t: X_1 \dots X_t) = I(S_{t+1}X_{t+1}: X_1 \dots X_t)$ only if $\Pr_{\mathfrak{G}}(s_{t+1}, x_{t+1}|s_t) > 0$ implies:

$$\Pr_{\mathfrak{G}}(x_1 \dots x_t | s_{t+1}, x_{t+1}) = \Pr_{\mathfrak{G}}(x_1 \dots x_t | s_t) .$$

Now, note that:

$$\Pr_{\mathfrak{G}} (x_1 \dots x_{t+1} | s_{t+1}) = \Pr_{\mathfrak{G}} (x_1 \dots x_t | s_{t+1}, x_{t+1}) \Pr(x_{t+1} | s_{t+1})$$
$$= \Pr_{\mathfrak{G}} (x_1 \dots x_t | s_t) \Pr(x_{t+1} | s_{t+1}) .$$

Rearranging and using the retrodictive equivalence partitions $\sigma_t := [s_t]_{\sim}$ and $\sigma_{t+1} := [s_{t+1}]_{\sim}$, we have:

$$\Pr_{\mathfrak{GP}}(x_1 \dots x_t | \sigma_t) = \frac{\Pr_{\mathfrak{GP}}(x_1 \dots x_{t+1} | \sigma_{t+1})}{\Pr_{\mathfrak{GP}}(x_{t+1} | \sigma_{t+1})} .$$
(S1)

Define a function $f: \mathcal{S} \times \mathcal{X} \to \mathcal{S}$ as follows. For a given σ' and x, let $f(\sigma', x)$ be the equivalence class such that:

$$\Pr_{\mathfrak{GP}}(\cdot|f(\sigma',x)) = \frac{\Pr_{\mathfrak{GP}}(\cdot,x|\sigma')}{\Pr_{\mathfrak{GP}}(x|\sigma')}$$

Such an equivalence class $f(\sigma', x)$ must exist by Eq. (S1). It is unique since, by definition, equivalence classes σ have unique distributions $\Pr(\cdot|\sigma)$. Then $\sigma_t = f(\sigma_{t+1}, x_{t+1})$ is a requirement for $\Pr_{\mathfrak{G}^{\mathcal{P}}}(\sigma_{t+1}, x_{t+1}|\sigma_t) > 0$. If we then take the merged generator $\mathfrak{G}^{\mathcal{P}}$, we must have $\widetilde{T}_{\sigma'\sigma}^{(x)} > 0$ only when $\sigma = f(\sigma', x)$.

Conversely, suppose that the retrodictively state-merged generator $\mathfrak{G}^{\mathcal{P}} = (\mathcal{P}_{\sim}, \mathcal{X}, \{\widetilde{T}_{\sigma'\sigma}^{(x)}\})$ satisfies $\widetilde{T}_{\sigma'\sigma}^{(x)} \propto \delta_{\sigma,f(\sigma',x)}$ for some function $f: \mathcal{S} \times \mathcal{X} \to \mathcal{S}$. Then, for a given s_t , its equivalence class $[s_t]_{\sim}$ is always a function of the next state and symbol: $[s_t]_{\sim} = f([s_{t+1}]_{\sim}, x_{t+1})$. This implies the Markov chain $[S_t]_{\sim} - X_{t+1}S_{t+1} - X_0 \dots X_t$. However, we also have the chain $X_{t+1}S_{t+1} - [S_t]_{\sim} - X_0 \dots X_t$ from the basic Markov property of the generator. Therefore, we must have:

$$I(S_{t+1}X_{t+1} : X_1 \dots X_t) = I([S_t]_{\sim} : X_1 \dots X_t)$$

= $I(S_t : X_1 \dots X_t)$.

Last, we connect this result to previous literature on the efficiency of retrodictors by showing that our conditions for efficiency imply that the generator is retrodictive. Recall that any generator whose state S_t is a sufficient statistic of \vec{X}_t for $X_1 \dots X_t$ is called a *retrodictor*.

Proposition 6. A generator $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ whose retrodictively state-merged generator $\mathfrak{G}^{\mathcal{P}} = (\mathcal{P}_{\sim}, \mathcal{X}, \{\widetilde{T}_{\sigma'\sigma}^{(x)}\})$ satisfies $\widetilde{T}_{\sigma'\sigma}^{(x)} \propto \delta_{\sigma,f(\sigma',x)}$ for some function $f : \mathcal{S} \times \mathcal{X} \to \mathcal{S}$ is a retrodictor.

This follows since $\mathfrak{G}^{\mathcal{P}}$, being co-unifilar and already retrodictively state-merged, must be the reverse ϵ -machine with states Σ_t . It is clear from the proof of Prop. 5 that for all t, we have the Markov chain $(X_1 \dots X_t) - \Sigma_t - S_t$. Since Σ_t is also the minimum sufficient statistic of \overrightarrow{X}_t for $X_1 \dots X_t$, we must have the Markov chain $(X_1 \dots X_t) - \overrightarrow{X}_t - S_t$. Combined with the Markov chain $(X_1 \dots X_t) - S_t - \overrightarrow{X}_t$, we conclude that S_t is a sufficient statistic of \overrightarrow{X}_t for $X_1 \dots X_t$. \Box

Appendix E: Forward and Reverse q-Machines

1. q-Machines and Their Time Reversals

A q-machine is constructed from a forward ϵ -machine $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ by choosing any set of phases $\{\phi_{xs} : x \in \mathcal{X}, s \in \mathcal{S}\}$ and constructing an encoding $\{|\psi_s\rangle : s \in \mathcal{S}\}$ of the memory states \mathcal{S} into a Hilbert space \mathcal{H}_S , and a set of Kraus operators $\{K^{(x)} : x \in \mathcal{X}\}$ on said Hilbert space. The phases, encoding states and the Kraus operators satisfy the formula:

$$K^{(x)} |\psi_{s}\rangle = e^{i\phi_{xs}} \sqrt{T_{f(s,x),s}^{(x)}} |\psi_{f(s,x)}\rangle \quad .$$
(S1)

This expression implicitly defines the Kraus operators given the encoding $\{|\psi_s\rangle\}$. The encoding, in turn, is determined up to a unitary transformation by the following constraint on their overlaps:

$$\langle \psi_r | \psi_s \rangle = \sum_{x \in \mathcal{X}} e^{i(\phi_{xs} - \phi_{xr})} \sqrt{T_{r',r}^{(x)} T_{s',s}^{(x)}} \langle \psi_{r'} | \psi_{s'} \rangle \quad , \tag{S2}$$

where r' = f(r, x) and s' = f(s, x) [49].

Now, let $\Omega_{rs} := \langle \psi_r | \psi_s \rangle$ be the unique solution to Eq. (S2), which is effectively an eigenvalue equation for a superoperator. Notice that this is entirely determined by the phases $\{\phi_{xs}\}$ and the dynamics $\{T_{s's}^{(x)}\}$ of the original ϵ -machine. It can be computed without any reference to encoding states or Kraus operators. Indeed, once Ω_{rs} is determined, the encoding states and Kraus operators can be explicitly constructed.

Let $\sqrt{\pi_r \pi_s} \Omega_{rs} = \sum_{\alpha} U_{r\alpha} U_{s\alpha}^* \lambda_{\alpha}$ be the singular value decomposition of $\sqrt{\pi_r \pi_s} \Omega_{rs}$ into a unitary $U_{i\alpha}$ and singular values $\lambda_{\alpha} > 0$. Suppose $\alpha = 1, \ldots, d$. Then given any computational basis $\{|\alpha\rangle : \alpha = 1, \ldots, d\}$, we can construct:

$$|\psi_s\rangle = \sum_{\alpha} \sqrt{\frac{\lambda_{\alpha}}{\pi_s}} U_{s\alpha}^* |\alpha\rangle \text{ and}$$
(S3)

$$K^{(x)} = \sum_{\alpha,\beta,s} e^{i\phi_{xs}} U^*_{s'\beta} U_{s\alpha} \sqrt{\frac{\lambda_{\beta}\pi_s}{\lambda_{\alpha}\pi_{s'}}} T^{(x)}_{s's} |\beta\rangle \langle\alpha| \quad .$$
(S4)

It is easy to check that $\langle \psi_r | \psi_s \rangle = \Omega_{rs}$ and that Eq. (S1) is satisfied by this construction. Notice that:

$$\rho_{\pi} = \sum_{s} \pi_{s} \left| \psi_{s} \right\rangle \left\langle \psi_{s} \right| = \sum_{\alpha} \lambda_{\alpha} \left| \alpha \right\rangle \left\langle \alpha \right| \; .$$

So the computational basis α is the diagonal basis of the stationary state ρ_{π} .

This explicit construction is useful for time-reversing the q-machine. Let \mathfrak{G} be a reverse ϵ -machine. The time-reverse of a generator \mathfrak{G} is the generator $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{\widetilde{T}_{s's}^{(x)}\})$ where $\widetilde{T}_{s's}^{(x)} = \pi_s T_{s's}^{(x)}/\pi_{s'}$. \mathfrak{G} is the forward ϵ -machine of the reverse process. From it, we can construct a q-machine, with Kraus operators expressed by $\{\widetilde{K}^{(x)}\}$. Recall that the generated process of the q-machine is given by

$$\Pr_{\widetilde{\mathfrak{G}}}(x_1\dots x_t) := \operatorname{Tr}\left[\widetilde{K}^{(x_t\dots x_1)}\rho_{\pi}\widetilde{K}^{(x_t\dots x_1)\dagger}\right]$$

This is the time-reverse of the process generated by \mathfrak{G} , expressed in the equation $\Pr_{\widetilde{\mathfrak{G}}}(x_1 \dots x_t) = \Pr_{\mathfrak{G}}(x_t \dots x_1)$. In

terms of the q-machine, we can write:

$$\Pr_{\mathfrak{G}}(x_1 \dots x_t) := \operatorname{Tr}\left[\widetilde{K}^{(x_t)} \dots \widetilde{K}^{(x_1)} \rho_{\pi} \widetilde{K}^{(x_1)\dagger} \widetilde{K}^{(x_t)\dagger}\right]$$
$$= \operatorname{Tr}\left[K^{(x_t)} \dots K^{(x_1)} \rho_{\pi} K^{(x_1)\dagger} \dots K^{(x_t)\dagger}\right],$$

where $K^{(x)} = \rho_{\pi}^{1/2} \widetilde{K}^{(x)\dagger} \rho_{\pi}^{-1/2}$. This is, essentially, the Petz reversal of the POVM $\{\widetilde{K}^{(x)}\}$, and it constitutes a formal time-reversal of the quantum process [50].

Computing $K^{(x)}$ is straightforward using Eq. (S3), as this gives the Kraus operators in the diagonal basis of ρ_{π} , where it is easiest to compute $\rho_{\pi}^{1/2}$ and its inverse. We have:

$$K^{(x)} = \sum_{\alpha,\beta,s} e^{-i\phi_{xs}} U_{s'\beta} U^*_{s\alpha} \sqrt{\frac{\pi_s}{\pi_{s'}}} \widetilde{T}^{(x)}_{s's} \left| \alpha \right\rangle \left\langle \beta \right| \; .$$

Now, take the basis $|\psi_s\rangle = \sum_{\alpha} U_{s\alpha}^* |\alpha\rangle$. In this basis:

$$K^{(x)} = \sum_{s'} e^{-i\phi_{xs'}} \sqrt{T^{(x)}_{s's}} \left| \psi_{s'} \right\rangle \left\langle \psi_s \right| \quad$$

Thus, we see that the basis $\{|\psi_s\rangle\}$ and Kraus operators $\{K^{(x)}\}$ form a reverse q-machine as described in the main body—one that compresses the reverse ϵ -machine \mathfrak{G} .

Note that the stationary state of a time-reversed q-machine is just the stationary state of the original q-machine—this is not altered under time reversal. However, we find a new expression for the stationary state, in terms of the basis $\{|\psi_s\rangle\}$:

$$\rho_{\pi} = \sum_{r,s,\alpha} \lambda_{\alpha} U_{r\alpha}^* U_{s\alpha} |\psi_s\rangle \langle\psi_r|$$
$$= \sum_{r,s} \sqrt{\pi_r \pi_s} \Omega_{sr} |\psi_s\rangle \langle\psi_r| .$$

So, ρ_{π} is generally not diagonal in the basis $\{|\psi_s\rangle\}$. The extent to which ρ commutes with $\{|\psi_s\rangle\}$ is the extent to which Ω_{rs} is block-diagonal.

2. Efficiency of *q*-Machines

To establish our first theorem relating memory to efficiency we turn to forward q-machines. First, we must prove a result regarding the synchronization of q-machine states.

Proposition 7 (Synchronization of q-machines). Let $\rho_{x_1...x_t} := \Pr_{\mathfrak{G}}(x_1...x_t)^{-1} K^{(x_t...x_1)} \rho_{\pi} K^{(x_t...x_1)\dagger}$ be the state of the q-machine's memory system after seeing the word $x_1...x_t$. Further, let $\hat{s}(x_1...x_t) := \operatorname{argmax}_{s_t} \Pr(s_t|x_1...x_t)$ be the most likely memory state after seeing the word $x_1...x_t$ and let $F(x_1...x_t) := F(\rho_{x_1...x_t}, |\psi_{\hat{s}}\rangle \langle \psi_{\hat{s}}|)$ be the fidelity between the quantum state of the memory system and the most likely encoded state. Then, there exist $0 < \alpha < 1$ and K > 0 such that, for all t:

$$\Pr\left(F(x_1\dots x_t) < 1 - \alpha^t\right) \le K\alpha^t$$

This shows that the quantum state of the memory system converges to a single encoded state with high probability.

This follows straightforwardly from a similar statement about ϵ -machines [38, 56, 57]. Let $Q(x_1 \dots x_t) := 1 - \Pr(\hat{s}|x_1 \dots x_t)$ be the probability of *not* being in the most likely state after word $x_1 \dots x_t$. Then there exist $0 < \alpha < 1$ and K > 0 such that, for all t:

$$\Pr\left(Q(x_1\dots x_t) > \alpha^t\right) \le K\alpha^t$$

Note that, from the Kraus operator definition, the quantum state of the memory system after word $x_1 \ldots x_t$ is:

$$\rho_{x_1...x_t} = \sum_{s_t} \Pr\left(s_t | x_1 \dots x_t\right) | \psi_{s_t} \rangle \left\langle \psi_{s_t} \right|$$

Now, the fidelity can be computed as:

$$F(x_1 \dots x_t) = \operatorname{Tr} \left[\sqrt{|\psi_{\hat{s}}\rangle \langle \psi_{\hat{s}}| \rho_{x_1 \dots x_t} |\psi_{\hat{s}}\rangle \langle \psi_{\hat{s}}|} \right]^2$$

= $\langle \psi_{\hat{s}}|\rho_{x_1 \dots x_t}|\psi_{\hat{s}}\rangle$
= $\operatorname{Pr} (\hat{s}|x_1 \dots x_t) + \sum_{s_t \neq \hat{s}} \operatorname{Pr} (s_t|x_1 \dots x_t) |\langle \psi_{s_t}|\psi_{\hat{s}}\rangle|^2$
 $\geq 1 - Q (x_1 \dots x_t) .$

And so, from the synchronization theorem for ϵ -machines, we have:

$$\Pr\left(F(x_1\dots x_t) < 1 - \alpha^t\right) \le K\alpha^t$$

for all t. \Box

The second notion we must introduce is a simple partition that may be constructed on the memory states of the ϵ -machine for a given quantum implementation. Specifically, we define the maximal commuting partition on S to be the most refined partition $\{\mathcal{B}_{\theta}\}$ such that the overlap matrix $\langle \psi_r | \psi_s \rangle$ is block-diagonal. That is, $\{\mathcal{B}_{\theta}\}$ is such that $\langle \psi_r | \psi_s \rangle = 0$ if $r \in \mathcal{B}_{\theta}$ and $s \in \mathcal{B}_{\theta'}$ for $\theta \neq \theta'$. From this partition we construct the maximally commuting local measurement required to define sufficient statistics.

Proposition 8. Let $\rho_{\mathfrak{G}}(t)$ be the state of the system $A_1 \ldots A_t S_t$ at time t. Let Θ be the projective measurement on \mathcal{H}_S corresponding to the MCP of the quantum generator. Then, for sufficiently large t, Θ is the MLCM of S_t . Similarly, $X_t \Theta$ is the MLCM of $A_t S_t$.

By Prop. 4, the MLCM must leave each $\rho_{x_1...x_t}$ unchanged, for all t. This is true for Θ . The question is if any nontrivial refinement, say Y, of Θ can do so. Now, realize that for any $\epsilon > 0$, there is sufficiently large t, so that for each state s there must be at least one word $x_1 \ldots x_t$ satisfying $F(\rho_{x_1...x_t}, |\psi_s\rangle \langle \psi_s|) > 1 - \epsilon$. Then, for sufficiently large t, it must be the case that there exists a word $x_1 \ldots x_t$ such that Y modifies $\rho_{x_1...x_t}$, because Y (by virtue of being a refinement of the maximal commuting partition) cannot commute with all the $|\psi_s\rangle \langle \psi_s|$. Therefore, Θ is the maximal commuting local measurement. That $X_t\Theta$ is the MLCM of A_tS_t follows from similar considerations. \Box We can now prove the following.

Theorem 9 (Maximally-efficient quantum generator). Let $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ be a given process' ϵ -machine. Suppose we build from it a quantum generator with encoding $\{\psi_s\}$ and Kraus operators $\{K^{(x)}\}$. Let $\mathcal{B} := \{\mathcal{B}_{\theta}\}$ be the MCP of \mathcal{S} . Then the quantum generator has $\Delta S_{\text{loc}} = 0$ if and only if the partition \mathcal{B} is trivially maximal—in that $|\mathcal{B}_{\theta}| = 1$ for each θ —and the retrodictively state-merged generator $\mathfrak{G}^{\mathcal{B}}$ of \mathfrak{G} is co-unifilar.

We define $\rho_{\mathfrak{G}}^{(\theta)}(t) := \Pi^{(\theta)}\rho_{\mathfrak{G}}(t)\Pi^{(\theta)}$, so that $\rho_{\mathfrak{G}}(t) = \bigoplus_{\theta} \rho_{\mathfrak{G}}^{(\theta)}(t)$. For the equivalence relation we say that $\theta \sim \theta'$ if $\rho_{\mathfrak{G}}^{(\theta)}(t)$ and $\rho_{\mathfrak{G}}^{(\theta')}(t)$ are unitarily equivalent on the subsystem S. This defines a partition $\mathcal{P} = \{[\theta]_{\sim}\}$ over the set $\{\mathcal{B}_{\theta}\}$. The quantum channel:

$$\mathcal{E}(\rho) := \sum_{x} K^{(x)} \rho K^{(x)\dagger} \otimes |x\rangle \langle x|$$

has Kraus operators $L^{(x)} := K^{(x)} \otimes |x\rangle$. The MLCM of $A_t S_t$ is $X_t \Theta_t$. Now, Thm. 1 demands that the Kraus operators $L^{(x)}$ have the form:

$$L^{(x)} = \bigoplus_{\theta,\theta'} \sqrt{\Pr(\theta', x|\theta)} U^{(x,\theta'|\theta)}$$
$$= \bigoplus_{\theta,\theta'} \sqrt{\Pr(\theta', x|\theta)} |x\rangle \otimes U^{(x)}_{\theta \mapsto \theta'} .$$

This implies that the quantum-generator Kraus operators have the form:

$$K^{(x)} = \bigoplus_{\theta,\theta'} \sqrt{\Pr(\theta', x|\theta)} U^{(x)}_{\theta \mapsto \theta'} .$$
(S5)

The values $\Pr(\theta', x | \theta)$ must be positive only when $\theta' x \sim \theta$.

Now, we note that this means there is a machine $(\mathcal{B}, \mathcal{X}, \{\hat{T}^{(x)}_{\theta'\theta}\})$ with transition probabilities $\hat{T}^{(x)}_{\theta'\theta} := \Pr(\theta', x|\theta)$ that generates the process. By the previous paragraph's conclusion, a merged machine must have the property that merging its retrodictively equivalent states results in a co-unifilar machine.

However, there is more to consider here. We assumed starting with a process' ϵ -machine. This means each state s_t must have a unique prediction of the future $\Pr(x_{t+1} \dots x_{\ell} | s_t)$. In fact, we can relate these predictions to the partition:

$$\Pr(x_{t+1}\dots x_{\ell}|s_t) = \operatorname{Tr}\left[K^{(x_{\ell}\dots x_{t+1})}|\psi_s\rangle \langle\psi_s| K^{(x_{\ell}\dots x_{t+1})\dagger}\right]$$
$$= \sum_{\theta_{t+1}\dots\theta_{\ell}} \Pr(\theta_{t+1}, x_{t+1}|\theta(s_t))\dots \Pr(\theta_{\ell}, x_{\ell}|\theta_{\ell-1}).$$

Which, we see, only depends on the partition index θ . Then any two $s_t, s'_t \in \mathcal{B}_{\theta}$ with $s_t \neq s'_t$ have the same future prediction, which is not possible for an ϵ -machine. Then, it must be the case that each partition has only one element: $|\mathcal{B}_{\theta}| = 1$ for all θ . The remainder of the theorem follows directly. \Box

3. Efficiency of Reverse q-Machines

To establish a similar theorem for the reverse q-machine, we must build similar prerequisite notions. It will be necessary, first, to be very precise about all the time-reversals at play here.

The reverse q-machine is built from the reverse ϵ -machine \mathfrak{G} of a process $\operatorname{Pr}_{\mathfrak{G}}$. This generates strings of random variables $X_1 \ldots X_L$. The time-reversed generator $\widetilde{\mathfrak{G}}$ is the forward ϵ -machine of the process $\operatorname{Pr}_{\widetilde{\mathfrak{G}}}$ and generates strings of random variables $\widetilde{X}_1 \ldots \widetilde{X}_t$, where $\widetilde{X}_j = X_{t-j+1}$. It is from \mathfrak{G} that we build the reverse q-machine.

However, further consider that the process $\Pr_{\mathfrak{G}}$ with random variables $X_1 \ldots X_t$ can itself be generated by some forward ϵ -machine, say $\mathfrak{F} = \left(\mathcal{P}, \mathcal{X}, \{R_{p'p}^{(x)}\}\right)$. There is a useful result [52] that we can always write:

$$\Pr_{\mathfrak{S}}(s_t|x_1\dots x_t) = \sum_{q_t} \Pr_{\mathcal{C}}(s_t|q_t) \Pr_{\mathfrak{F}}(q_t|x_1\dots x_t)$$
(S6)

for some channel $\Pr_{\mathcal{C}}(s_t|q_t)$. Combined with the synchronization result for ϵ -machines, the following result for reverse ϵ -machines arises straightforwardly.

Proposition 9 (Mixed-state synchronization of retrodictors). Let \mathfrak{G} be a reverse ϵ -machine and let \mathfrak{F} be the forward ϵ -machine of the same process. For a length-L word $x_1 \ldots x_t$, let $\hat{p}(x_1 \ldots x_t) := \operatorname{argmax}_{p_t} \operatorname{Pr}_{\mathfrak{F}}(q_{=}t|x_1 \ldots x_t)$. Let $D(x_1 \ldots x_t) = \frac{1}{2} \|\operatorname{Pr}_{\mathfrak{G}}(\cdot|x_1 \ldots x_t) - \operatorname{Pr}_{\mathcal{C}}(\cdot|\hat{p})\|_1$. Then there exist $0 < \alpha < 1$ and K > 0 such that, for all t:

$$\Pr\left(D(x_1\dots x_t) > \alpha^t\right) \le K\alpha^t \,.$$

This shows that for a sufficiently long word $x_1 \dots x_t$, the distribution over states of the reverse ϵ -machine converges to one of a finite number of options (namely, the $\Pr_{\mathcal{C}}(s|q)$ for different q).

Note that:

$$\Pr_{\mathfrak{G}}(s_t|x_1\ldots x_t) = \Pr_{\mathcal{C}}(s_t|\hat{p}) \Pr_{\mathfrak{F}}(\hat{p}|x_1\ldots x_t) + \sum_{p_t\neq\hat{p}} \Pr_{\mathcal{C}}(s_t|p_t) \Pr_{\mathfrak{F}}(p_t|x_1\ldots x_t)$$

So:

$$\Pr_{\mathcal{C}}(s_t|\hat{p}) - \Pr(s_t|\hat{p}) Q(x_1 \dots x_t) \leq \Pr_{\mathfrak{G}}(s_t|x_1 \dots x_t)$$
$$\leq \Pr_{\mathcal{C}}(s_t|\hat{p}) + (1 - \Pr(s_t|\hat{p})) Q(x_1 \dots x_t) ,$$

where $Q(x_1 \dots x_t) = 1 - \Pr_{\mathfrak{F}}(p_t | x_1 \dots x_t)$. Then the result follows from the forward ϵ -machine synchronization theorem. \Box

Now, we prove a synchronization theorem for the reverse q-machine.

Proposition 10 (Pure-state synchronization of retrodictors). Let \mathfrak{G} be a reverse ϵ -machine and let \mathfrak{F} be the forward ϵ -machine of the same process. For a length-L word $w = x_1 \dots x_t$, let $\hat{p}(x_1 \dots x_t) := \operatorname{argmax}_{p_t} \operatorname{Pr}_{\mathfrak{F}}(p_L|x_1 \dots x_t)$. Define the state:

$$|\Psi_{x_1...x_t}\rangle := \sum_{s_t} \sqrt{\Pr_{\mathcal{C}}(s_t|\hat{p})} e^{i\phi_{s_tw}} |\psi_{s_t}\rangle ,$$

where $\phi_{s_tw} = \sum_{j=1}^t \phi_{s_jx_j}$ (recall each s_j is determined by s_{j+1} and x_{j+1}). Last, let $\rho_{x_1...x_t} := \Pr(x_1...x_t)K^{(x_1...x_t)}\rho_{\pi}K^{(x_1...x_t)\dagger}$ and let $F(x_1...x_t) = F(\rho_{x_1...x_t}, |\Psi_{x_1...x_t}\rangle \langle \Psi_{x_1...x_t}|)$. Then there exist $0 < \alpha < 1$ and K > 0 such that, for all t:

$$\Pr\left(F(x_1\dots x_t) < 1 - \alpha^t\right) \le K\alpha^t$$

This shows that for a sufficiently long word $x_1 \dots x_t$, the reverse q-machine state converges on a pure state whose amplitudes are determined by the mapping $\Pr_{\mathcal{C}}(s|p)$ from \mathfrak{F} states to \mathfrak{G} states.

First, recall that $x_1 \ldots x_t$ corresponds to the reverse of the sequence $\tilde{x}_1 \ldots \tilde{x}_t$, generated by the forward ϵ -machine \mathfrak{G} of the reverse process. We have the relation $\Pr_{\mathfrak{G}}(\tilde{s}_t | \tilde{x}_1 \ldots \tilde{x}_t) = \Pr_{\mathfrak{G}}(s_0 | x_1 \ldots x_t)$ where $\tilde{s}_0 = s_t$; representing the initial state of \mathfrak{G} and the final state of \mathfrak{G} under time reversal. By the forward ϵ -machine synchronization theorem, then the distribution $\Pr_{\mathfrak{G}}(s_0 | x_1 \ldots x_t)$ is concentrated on a particular state \hat{s} with high probability. Now:

$$\begin{split} \rho_{x_1...x_t} &= \sum_{\substack{s_0 \\ s_t, r_t}} \sqrt{\Pr\left(s_t, s_0 | x_1 \dots x_t\right) \Pr\left(r_t, s_0 | x_1 \dots x_t\right)} e^{(i\phi_{s_tw} - i\phi_{r_tw})} \left|\psi_{s_t}\right\rangle \left\langle\psi_{r_t}\right| \\ &= (1 - Q) \sum_{\substack{s_t, r_t}} \sqrt{\Pr\left(s_t | x_1 \dots x_t, \hat{s}\right) \Pr\left(r_t | x_1 \dots x_t, \hat{s}\right)} e^{(i\phi_{s_tw} - i\phi_{r_tw})} \left|\psi_{s_t}\right\rangle \left\langle\psi_{r_t}\right| \\ &+ Q \sum_{\substack{s_0 \neq \hat{s} \\ s_t, r_t}} \sqrt{\Pr\left(s_t | x_1 \dots x_t, s_0\right) \Pr\left(r_t | x_1 \dots x_t, s_0\right)} e^{(i\phi_{s_tw} - i\phi_{r_tw})} \left|\psi_{s_t}\right\rangle \left\langle\psi_{r_t}\right| \ , \end{split}$$

where $Q := 1 - \Pr(\hat{s} | x_1 \dots x_t).$

To handle the term $\Pr(s_t|x_1 \dots x_t, \hat{s})$, we use the co-unifilarity properties of the reverse ϵ -machine. Note that:

$$\Pr\left(s_t | x_1 \dots x_t, \hat{s}\right) = \frac{\Pr\left(\hat{s} | x_1 \dots x_t, s_t\right) \Pr\left(s_t | x_1 \dots x_t\right)}{1 - Q} .$$

By co-unifilarity, $\Pr(\hat{s}|x_1 \dots x_t, s_t)$ is either 0 or 1 depending on the value of s_t . The synchronization theorem then demands that $\Pr(s_t|x_1 \dots x_t) < Q$ for those s_t which assign $\Pr(\hat{s}|x_1 \dots x_t, s_t)$ a zero value. Due to this, along with the above equation:

$$\rho_{x_1\dots x_t} = \sum_{s_t, r_t} \sqrt{\Pr\left(s_t | x_1 \dots x_t\right) \Pr\left(r_t | x_1 \dots x_t\right)} e^{\left(i\phi_{s_t w} - i\phi_{r_t w}\right)} \left|\psi_{s_t}\right\rangle \left\langle\psi_{r_t}\right| + O(Q) ,$$

where O(Q) contains all the terms so far which scale with Q. Then $\rho_{x_1...x_t} = |\Psi_w\rangle \langle \Psi_w| + O(Q)$. The rest follows as in Prop. 7. \Box

With this, we have sufficient information to determine the MCLM of the system $S_t A_t \dots A_1$ for sufficiently long t. Let π_s be the stationary distribution for \mathfrak{S} 's memory state, and let λ_q be the same for \mathfrak{F} . Consider the channel $\Pr(s'|s) = \sum_q \Pr(s'|q) \Pr(s|q) \lambda_q / \pi_s$. From it we can construct the ergodic partition $\{\mathcal{B}_\theta\}$ which is defined as the most refined partition such that $\Pr(s'|s) > 0$ only if $\theta(s) = \theta(s')$. This partition represents information about the state of \mathfrak{S} which is recoverably encoded in the state of \mathfrak{F} . If, on the one hand, the partition is trivially coarse-grained ($|\{\mathcal{B}_\theta\}| = 1$), then no information about \mathfrak{S} is recoverably encoded. If, on the other, the partition is trivially maximal ($|\mathcal{B}_\theta| = 1$ for all θ), then the state of \mathfrak{S} is actually a function of the state of \mathfrak{F} . Consequently, in this extreme case, \mathfrak{S} is unifilar in addition to being unifilar.

Proposition 11. Let $\rho_{\mathfrak{F}}(t)$ be the state of the system $A_1 \ldots A_t S_t$ at time t. Let Θ be the projective measurement on \mathcal{H}_S corresponding to the ergodic partition described above. Then, for sufficiently large t, Θ is the MLCM of S_t . Similarly, $X_t \Theta$ is the MLCM of $A_t S_t$.

By Prop. 4, the MLCM must leave each $\rho_{x_1...x_t}$ unchanged, for all t. This is true for Θ . The question is if any nontrivial refinement, say Y, of Θ can do so. Realize that for any $\epsilon > 0$, there is sufficiently large t, so that for each state s there must be at least one word $x_1 \ldots x_t$ satisfying $F(\rho_{x_1...x_t}, |\Psi_{x_1...x_t}\rangle \langle \Psi_{x_1...x_t}|) > 1 - \epsilon$. Then, for sufficiently large t, it must be the case that there exists a word $x_1 \ldots x_t$ such that Y modifies $\rho_{x_1...x_t}$, because Y (by virtue of being a refinement of Θ) cannot commute with all the $|\psi_s\rangle \langle \psi_s|$. Therefore, Θ is the maximal commuting local measurement. That $X_t\Theta$ is the MLCM of A_tS_t follows from similar considerations. \Box

Theorem 10 (Maximally-efficient reverse q-machine). Let $\mathfrak{G} = (\mathcal{S}, \mathcal{X}, \{T_{s's}^{(x)}\})$ be a given process' reverse ϵ -machine. Suppose we build from it a reverse q-machine with encoding $\{|\psi\rangle_s\}$ and Kraus operators $\{K^{(x)}\}$. Let $\mathcal{B} := \{\mathcal{B}_{\theta}\}$ be the MCP of \mathcal{S} . Then the quantum generator has $\Delta S_{\text{loc}} = 0$ if and only if the partition \mathcal{B} is trivially maximal—in that $|\mathcal{B}_{\theta}| = 1$ for each θ —and the predictively state-merged generator $\mathfrak{G}^{\mathcal{B}}$ of \mathfrak{G} is unifilar.

We define $\rho_{\mathfrak{G}}^{(\theta)}(t) := \Pi^{(\theta)} \rho_{\mathfrak{G}}(t) \Pi^{(\theta)}$, so that $\rho_{\mathfrak{G}}(t) = \bigoplus_{\theta} \rho_{\mathfrak{G}}^{(\theta)}(t)$. For the equivalence relation we say that $\theta \sim \theta'$ if $\rho_{\mathfrak{G}}^{(\theta)}(t)$ and $\rho_{\mathfrak{G}}^{(\theta')}(t)$ are unitarily equivalent on the subsystem S. This defines a partition $\mathcal{P} = \{[\theta]_{\sim}\}$, labeled by σ , over the set $\{\mathcal{B}_{\theta}\}$.

The quantum channel:

$$\mathcal{E}\left(\rho\right) := \sum_{x} K^{(x)} \rho K^{(x)\dagger} \otimes \left|x\right\rangle \left\langle x\right|$$

has Kraus operators $L^{(x)} := K^{(x)} \otimes |x\rangle$. The MLCM of $A_t S_t$ is $X_t \Theta_t$. Theorem 1 demands that the Kraus operators $L^{(x)}$ have the form:

$$L^{(x)} = \bigoplus_{\theta,\theta'} \sqrt{\Pr(\theta', x|\theta)} U^{(x,\theta'|\theta)}$$
$$= \bigoplus_{\theta,\theta'} \sqrt{\Pr(\theta', x|\theta)} |x\rangle \otimes U^{(x)}_{\theta \mapsto \theta'}$$

This implies that the quantum generator's Kraus operators have the form:

$$K^{(x)} = \bigoplus_{\theta,\theta'} \sqrt{\Pr(\theta', x|\theta)} U^{(x)}_{\theta \mapsto \theta'} .$$
(S7)

The values $\Pr(\theta', x|\theta)$ must be positive only when $\theta' x \sim \theta$. This would imply that the resulting merged machine is retrodictive. However, since the states S are those of the reverse ϵ -machine, they cannot be further merged into a retrodictive machine. It must then be the case that the partition Θ is trivially maximal. Consequently, it must be the case that \mathfrak{F} is predictive. \Box

The consequences of this theorem are augmented by the following statement about partition $\{\mathcal{B}_{\theta}\}$ and the stationary distribution ρ_{π} .

Proposition 12 (Block-diagonality of stationary distribution.). Let Θ be the projective measurement on \mathcal{H}_S corresponding to the ergodic partition described above, with projectors $\{\Pi_{\theta}\}$. Then $\rho_{\pi} = \sum_{\theta} \Pi_{\theta} \rho_{\pi} \Pi_{\theta}$.

For, we can express the stationary distribution as:

$$\rho_{\pi} = \sum_{x_1 \dots x_t} \rho_{x_1 \dots x_t}$$

at arbitrarily long t. Taking t sufficiently large, each $\rho_{x_1...x_t}$ is arbitrarily close to $|\Psi_{x_1...x_t}\rangle \langle \Psi_{x_1...x_t}|$. Since all these pure states commute with Θ , so does ρ_{π} . \Box

The primary implication of this theorem, then, is that trivial maximality of Θ implies ρ_{π} is diagonal in the basis $\{|\psi_s\rangle\}$, which itself implies that the reverse q-machine does not achieve any nonzero memory compression.