Thermodynamics of complexity and pattern manipulation

Vedral, Vlatko; Gu, Mile; Garner, Andrew J. P.; Thompson, Jayne

2017

Garner, A. J. P., Thompson, J., Vedral, V., & Gu, M. (2017). Thermodynamics of complexity and pattern manipulation. *Physical Review E*, 95(4), 042140.

<https://hdl.handle.net/10356/87402>

<https://doi.org/10.1103/PhysRevE.95.042140>

© 2017 American Physical Society (APS). This paper was published in *Physical Review E* and is made available as an electronic reprint (preprint) with permission of American Physical Society (APS). The published version is available at: [<http://dx.doi.org/10.1103/PhysRevE.95.042140>]. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper is prohibited and is subject to penalties under law.

Downloaded on 26 Apr 2023 22:01:11 SGT

Thermodynamics of complexity and pattern manipulation

Andrew J. P. Garner,^{1,2,*} Jayne Thompson,¹ Vlatko Vedral,^{3,1,4,2} and Mile Gu^{5,6,1,2,†}

¹Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, 117543, Singapore

²Center for Quantum Information, Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing 100084, China

³Atomic and Laser Physics, University of Oxford, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, United Kingdom

⁴Department of Physics, National University of Singapore, 3 Science Drive 2, Singapore 117543

⁵School of Physical and Mathematical Sciences, Nanyang Technological University, 639673, Singapore

⁶Complexity Institute, Nanyang Technological University, 639673, Singapore

(Received 10 October 2016; revised manuscript received 30 November 2016; published 24 April 2017)

Many organisms capitalize on their ability to predict the environment to maximize available free energy and reinvest this energy to create new complex structures. This functionality relies on the *manipulation of patterns*—temporally ordered sequences of data. Here, we propose a framework to describe pattern manipulators—devices that convert thermodynamic work to patterns or vice versa—and use them to build a “pattern engine” that facilitates a thermodynamic cycle of pattern creation and consumption. We show that the least heat dissipation is achieved by the provably simplest devices, the ones that exhibit desired operational behavior while maintaining the least internal memory. We derive the ultimate limits of this heat dissipation and show that it is generally nonzero and connected with the pattern’s intrinsic *crypticity*—a complexity theoretic quantity that captures the puzzling difference between the amount of information the pattern’s past behavior reveals about its future and the amount one needs to communicate about this past to optimally predict the future.

DOI: [10.1103/PhysRevE.95.042140](https://doi.org/10.1103/PhysRevE.95.042140)

I. INTRODUCTION

The manipulation of patterns is as important to living organisms as it is for computation. Living things capitalize on structure in their environment for available energy and use this energy to generate new complex structures. Similarly, a crucial task in the modern era of big data is to identify patterns in large data sets in order to make predictions about future events—often at great energetic cost. Here, we consider the thermodynamic costs intrinsic to this sort of pattern manipulation and ask: is there a preferred method by which this manipulation should be done? Our intuition is that *simpler is better*, a longstanding tenant of natural philosophy known as Occam’s razor. To formalize this, we first qualify what is meant both by *simpler* and by *better*.

In complexity science, computational mechanics formalizes what is *simpler* in the context of pattern manipulation [1–3]. The premise is that everything we observe in the environment can be considered to be a pattern—a temporal sequence of data exhibiting certain statistical structure. Much of science then deals with building models that can explain such statistics—machines that take information from past observations and use it to generate statistically coinciding conditional future predictions. Given two machines that exhibit the same pattern of behavior, the one that stores less information from the past is considered simpler, the motivation being that it better isolates indicators of future behavior. The simplest such machine then defines exactly how much memory is required to produce a given pattern and thus quantifies the pattern’s intrinsic structure. Known as *statistical complexity*, this measure has been applied to quantify structure in diverse contexts [4–6].

Meanwhile, in thermodynamics, *better* originally described heat engines that produce more work with less wasted heat. This carries through to modern thermodynamics: the best approach for a given task being the one that minimizes the expenditure of a limited resource [7,8] (e.g., work or hard-to-create states). Since information is physical [9–12], patterns, which are correlated information, are also physical and hence subject to the laws of thermodynamics. In this context, a pattern may be treated as an *information reservoir*—a source of free energy encoded in correlations (e.g., Refs. [9,12–19]). *Pattern manipulators*, which convert a pattern to useful work or vice versa, are thus a type of heat engine. The pattern manipulator that effects a prescribed change in a pattern with the minimal heat dissipation can thus be regarded as *better*.

Here, we derive the fundamental thermodynamic limits for the manipulation of patterns by devices operating in a cycle (see Fig. 1). We do not rederive the second law, but (believing it highly likely to be true) consider the implications that it places on the work dissipation intrinsic to *all* possible patterns manipulators. Our approach is to connect pattern manipulators in the context of thermodynamics to predictive models in the context of computational mechanics—observing that the creation or consumption of a given pattern involves retaining enough of the past to correctly anticipate its expected future statistics. We show that simplest patterns manipulators (i.e., ones that store the least information about the past of the pattern) results in the least dissipation—and thus simpler is thermodynamically better.

We describe the simplest, most thermodynamically efficient causal pattern manipulators—those whose memory requirement is given by the pattern’s statistical complexity—and show that they, remarkably, still must dissipate some excess heat. We show that this heat dissipation is lower-bounded by the *crypticity* of the pattern [20], a hitherto complexity—theoretic property quantifying the puzzling difference between the amount of information the past of a pattern reveals about

*ajpgarner@nus.edu.sg

†cqtmileg@nus.edu.sg

FIG. 1. Cycle of pattern generation and extraction. A tape moves through the system in a clockwise manner. A generator expends work to write a pattern to the tape. The extractor then uses the pattern on this tape to extract work. To run cyclically, each device maintains *prescient* memory that keeps track of the pattern. In this article, we identify the dissipative work costs. We find that the simplest generator has the best thermodynamic performance; but surprisingly for the extractor, the choice of memory has no thermodynamic consequence.

its future, and the amount of information one needs to communicate about the past of the pattern in order to predict its future. These bounds apply to any model in any physical framework that can implement pattern manipulation tasks described in a manner consistent with Landauer’s principle and computational mechanics. Our work thus highlights the many thermodynamic consequences of complexity in pattern manipulation.

II. PATTERNS AS A RESOURCE

Knowing a system’s internal state has thermodynamic consequence. This knowledge can be used to perform work (drive a mechanical task), as illustrated by the Szilárd engine thought experiment [11]: A box has a single particle inside, on the left- or right-hand side. A movable barrier inserted in the box’s center acts as a piston that expands as the particle pushes against it. If an agent knows which side of the barrier the particle is on, she can couple the barrier (e.g., via a pulley) to raise a weight. As the piston expands, it lifts the weight and generates an amount of work $k_B T \ln 2$ at temperature T by drawing in the same amount of heat from its surroundings.

Knowledge about patterns may also be exploited. Consider a sequence of Szilárd engines arranged in a linear configuration on a conveyer belt (or more abstractly, symbols on a tape), indexed sequentially by $t \in \mathbb{Z}$. These engines are prepared such that the particle in engine t is on the same side of the barrier as the in engine $t - 1$ with probability $p \neq \frac{1}{2}$. Suppose

now that an agent attempts to extract work from these engines in sequential order. An agent unaware of this pattern would only be able to correctly predict the particle’s location half of the time and, hence, will extract less work than an agent who knows the pattern and couples her pulley accordingly. As such, the ability to predict grants thermodynamic advantage.

This is a manifestation of Maxwell’s *dæmon*—an apparently paradoxical conversion of heat into work that is only resolved by accepting that information is physical and hence subject to the laws of thermodynamics [9–12]. For the single Szilárd engine, we must also account for the cost of resetting the agent’s memory about the particle’s location—this knowledge must be thought of as a resource. Likewise, since it is more thermodynamically useful that the sequence follows a pattern than be uncorrelated, the pattern itself must also be considered as a resource. Producing a pattern hence requires an investment of work. Moreover, any physical device that generates (or exploits) a pattern contains some memory about what has been observed in the pattern so far, in order to accurately generate (or anticipate) upcoming parts of the pattern. Any thermodynamic costs of maintaining this internal memory must also be accounted for. One thus can think of a tape with a pattern as one thinks of a rising weight (or any other form of battery) and treat it as an *information reservoir* [9,12–16]. This approach has been used recently as a bridge from computational mechanics to modern thermodynamics [17–19].

The quantitative link between information, entropy, and heat dissipation is given by Landauer’s principle [11]: the minimum work cost of any information-processing task is proportional to the total change in information entropy¹ (just like macroscopic thermodynamics, where the minimum work required to slowly change between two states of the same internal energy is proportional to the change in thermodynamic entropy).

Landauer’s principle sets a lower bound on the work cost of performing an information processing task. In this article, we shall take Landauer’s principle as our starting point. Our results have meaning in any physical framework that: (1) provides a way of mapping random variables to physical states and (2) has definitions of heat and work exchange, such that any allowed reconfiguration of these mapped states has a work cost bounded by Landauer’s principle, as applied to the associated random variables. A list of such frameworks satisfying these criteria includes the trajectory formalism [21–26] (see *Physical Example* in the *Appendix*), the resource theory of thermodynamics (e.g., Refs. [7,8]), single-shot statistical mechanics (e.g., Refs. [27–29]), among others.

In these frameworks, Landauer’s principle is not an additional imposition on top of the physics but rather manifests from microdynamical behavior—it is an emergent law, much like the second law of thermodynamics. It is not necessary to know the particulars of the microdynamics to derive or apply the results in this article—so long as the framework obeys Landauer’s principle, any model in that framework

¹Classically, given by the *Shannon entropy* $H(X) = -\sum_i P(X = x_i) \log_2 P(X = x_i)$. Using a base 2 logarithm gives units of *bits*.

that implements the pattern manipulations presented in the following sections will be bound by our results.

III. THE FRAMEWORK OF PATTERNS

Computational mechanics provides a formal framework for describing patterns [1–3,20]. Consider a sequence of physical systems indexed by a parameter $t \in \mathbb{Z}$, each having a configuration space $\{x\}$. A general pattern on such a sequence of systems is defined by a bi-infinite sequence of random variables $\vec{X} = \dots X^{t-1} X^t X^{t+1} \dots$, where X^t governs the configuration of system t . Here, one generally considers devices that observe the systems in some sequential order, such that system t is observed at time t . Thus, $\vec{X}^t = \dots X^{t-1} X^t$ govern the configurations of all systems that could have been observed in the past, and $\vec{X}^t = X^{t+1} X^{t+2} \dots$ governs the configurations of all systems that could be observed in the future. A statistical description of the *pattern* is given by the probability distribution $P(\vec{X}^t, \vec{X}^t)$, where each instance of such a pattern taking on a particular configuration $\vec{x} = \dots x^{t-1} x^t x^{t+1} \dots$ occurs with probability $P(\vec{X} = \vec{x})$.

When there is a meaningful pattern, the past and future are correlated. For a particular instance of the pattern where $\vec{x}$ is observed, the statistics of future observations are given by $P(\vec{X}^t | \vec{X}^t = \vec{x})$. The mutual information² $I(\vec{X}^t; \vec{X}^t) = H(\vec{X}^t) - H(\vec{X}^t | \vec{X}^t)$ then quantifies how useful this past knowledge is for predicting the future. In computational mechanics, this value is known as the *excess entropy* E (see Definition 13 of Ref. [2]).

We make the simplifying assumption that the pattern is a *stationary*—the statistics $P(\vec{X}^t, \vec{X}^t)$ are invariant under time translation. This does not mean that every output x^t in the sequence is identical, or that the pattern is Markovian, but rather that the statistics of X^{t+1} onwards, given a past sequence, have no explicit time dependence: $P(\vec{X}^t | \vec{X}^t = \vec{x}) = P(\vec{X}^{t'} | \vec{X}^{t'} = \vec{x})$ for all t, t' , and $\vec{x}$. Hence, we can often omit the superscript t .

IV. CYCLIC PATTERN MANIPULATION

The classical expressions of thermodynamic laws (e.g., Kelvin’s statement that a device cannot convert heat into work with no other effect [30]) concern cyclic behavior—processes that leave the system in a state allowing for repetition with the same thermodynamic consequences. Without a full cycle in mind, there is the danger that the thermodynamic benefit of a process may come at the expense of consuming an unaccounted-for resource. A cycle does not require the *microstate* of the system to return to its original value. Consider a piston of gas expanding and compressing: it does not matter

if the individual molecules have moved to new locations by the end of the cycle, as long as the important thermodynamic variables—the pressure and volume—return to their original values.

We now establish a framework for understanding the thermodynamics of generating and consuming patterns as a cyclic process (Fig. 1). Consider a sequence of physical systems indexed by a parameter $t \in \mathbb{Z}$, each having a configuration space $\{x\}$. We assume that the default configuration for these systems to be uncorrelated, such that the configuration of each system is governed by some default random variable X_{dft} . The mathematical results of this article will hold for any choice of X_{dft} , though for clarity, we shall frame our discussion as if X_{dft} represents the uniformly random distribution.

For each run of the cycle, we act on a moving window of length k on this sequence. We index the first system in this window as $t + 1$, and the final system as $t + k$. Between each run of the cycle, the window advances by k systems, such that the subsequent cycle acts on systems $t + k + 1$ to $t + 2k$. We refer to k as the *stride* of the cycle.

In each cycle, the tape is acted on by two machines (Fig. 1) associated with the same pattern $P(\vec{X}, \vec{X})$. Each machine is in contact with a thermal reservoir (i.e., heat bath at inverse temperature $\beta = \frac{1}{k_B T}$) and a battery for storing free energy (e.g., a rising weight). The first machine, a *generator*, does work in order to act on the systems in the window $t + 1, t + 2, \dots, t + k$, taking their configurations from $X_{\text{dft}}^{\otimes k}$ to configurations governed by the random variables $X^{t+1} X^{t+2} \dots X^{t+k}$ from the pattern. The second machine, an *extractor*, resets these k systems back to their uncorrelated default state $X_{\text{dft}}^{\otimes k}$, outputting work as it does so. Collectively, we refer to these devices as *pattern manipulators*.

Each of these two devices can operate independently. By itself, the generator operating *ad infinitum* encodes an arbitrarily long section of the pattern onto a sequence of physical systems that were initially in the default configuration. This will require an investment of work and produces a pattern as a resource. Likewise, the extractor consumes a section of the pattern of arbitrary length, extracting work and resetting the systems back to their default configuration. However, to account for all thermodynamic resources, it is helpful to consider them operating together as the cycle described above (or in Fig. 1).

V. PRESCIENT MEMORY

For any device to generate a particular pattern $\vec{X}$ in time-sequential order, it must contain an internal memory (with some configuration space $\{r\}$) that records some information about what it has generated before. Consider any pattern with two possible pasts $\vec{x}$ and $\vec{x}'$ that yield differing conditional future statistics; i.e., $P(\vec{X} | \vec{X} = \vec{x}) \neq P(\vec{X} | \vec{X} = \vec{x}')$. A machine that generates such a pattern must behave differently depending on whether it has generated $\vec{x}$ or $\vec{x}'$ so far. Thus, the state of the machine’s memory $\{r\}$ must be dependent on $\vec{x}$. Likewise, an extractor must also adjust its future actions based on what past $\vec{x}$ it has observed so far to best harness the

²The *mutual information* between random variables A and B with domains $\mathcal{A}$ and $\mathcal{B}$, respectively, is given $I(A; B) = H(A) - H(A | B)$, where $H(A | B) := - \sum_{b \in \mathcal{B}} P(B) \sum_{a \in \mathcal{A}} P(A = a | B = b) \log_2 P(A = a | B = b)$ is the *conditional entropy* of A given B . Using logarithms of base 2, all have units of bits.

free energy in $\vec{x}$. Thus, it must also field some $\vec{x}$ -dependent memory.

Each particular strategy for recording this past information can be captured by some mapping f that describes the state of the memory, r , depends on $\vec{x}$. f is referred to as an *encoding map* and defines a particular strategy in which a given generator or extractor keeps track of the past. The encoding map then induces a probability distribution over r , governed by some random variable R .

To generate a desired pattern governed by $P(\vec{X}, \vec{X})$, the memory R must be *prescient* with respect to the pattern. That is, for all possible $\vec{x}$, R satisfies $P(\vec{X} | \vec{X} = \vec{x}) = P(\vec{X} | R = r)$, for any r assigned by $f(\vec{x})$. In words: for the purpose of inferring or generating the pattern's future statistics behavior $\vec{X}$, knowing the state of *prescient memory* r is as useful as knowing the entire semi-infinite string of past outcomes $\vec{x}$. It follows that for any prescient memory R , the mutual information $I(R; \vec{X}) = I(\vec{X}; \vec{X})$. We shall consider only memory that acts in a *causal way*, such that the state of the memory at time t does not contain any information about the future of the pattern $\vec{X}$, that is not already contained in the past. (In the language of computational mechanics, we do not allow the memory to have *oracular information* [31]).

Given any pattern, there are clearly many encoding maps that generate memory states satisfying prescience. The most obvious choice is the identity map, which represents a generator that remembers every single step of the pattern it has generated up to the present. Such choices are clearly wasteful. Formally, the Shannon entropy of the memory, $H(R)$ would coincide with $H(\vec{X})$, and be unboundedly large—even when generating patterns with no structure (e.g., a completely random sequence). This has motivated complexity theorists to look for more efficient encodings—corresponding to generators whose corresponding memory, $H(R)$, is minimized.

The *causal states* [1,2] represent the most efficient prescient encoding. They coincide with the equivalence classes defined by the equivalence relation $\sim_\epsilon$, where $\vec{x} \sim_\epsilon \vec{x}'$ if and only if $P(\vec{X} | \vec{X} = \vec{x}) = P(\vec{X} | \vec{X} = \vec{x}')$. The optimal machine does not store which past $\vec{x}$ has occurred, but rather which equivalence class s that $\vec{x}$ belongs to—the rationale being that two pasts in the same equivalence class do not need to be distinguished as they have coinciding conditional futures. This motivates a particular encoding map ϵ that deterministically takes each $\vec{x}$ to a particular causal state s whenever $x \in s$. The resulting random variable S over causal states is uniquely defined for each pattern $\vec{X}$ and represents the prescient memory with minimal entropy [1,32]. The memory required to store these states, given by the Shannon entropy $C_\mu = H(S)$, is known as the pattern's *statistical complexity* and quantifies the pattern's intrinsic structure.

In computational mechanics, generators and predictors that store only the causal states are considered the simplest such devices and are known as ϵ -machines. Furthermore, any pattern can be described by dynamics on causal states (see Ref. [33] or our subsequent example in Fig. 4), and significant work exists on inferring causal states and the

resulting ϵ -machines from raw observational data [1,4,5,32]. Throughout this article, we use S to represent the random variable governing the pattern's causal state and R to represent the random variable governing the state of some generic choice of prescient memory. R could be in one-to-one correspondence with S (since ϵ is a specific choice of encoding map f), but in general it does not have to be.

VI. PRESCIENT PATTERN MANIPULATORS

To operate continually in a cycle, a pattern manipulator's memory must be prescient before and after the device has acted on the pattern. Suppose a generator at time t starts with memory in configuration $R^t = r_i$. After encoding k systems into the pattern $X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k}$, the final state of memory $R^{t+k} = r_j$ must satisfy $P(\vec{X}^{t+k} | R^{t+k} = r_j) = P(\vec{X}^{t+k} | R^t = r_i, X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k})$. We call this condition *maintaining prescience*, since it implies that R^{t+k} must also be prescient for the entire history of the pattern to that point, including the variate $x^{t+1} \dots x^{t+k}$ just produced. Likewise, if an extractor encounters and resets $X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k}$ on the tape, it must also update its memory from R^t to R^{t+k} in a way that satisfies the same condition.

For a generator and extractor that maintain prescient memory R and R' , respectively, if the initial configurations r_i and r'_i satisfy $P(\vec{X}^t | R^t = r_i) = P(\vec{X}^t | R'^t = r'_i)$, then the final configurations r_j and r'_j (after k steps of the pattern have been written and subsequently reset) will satisfy $P(\vec{X}^{t+k} | R^{t+k} = r_j) = P(\vec{X}^{t+k} | R'^{t+k} = r'_j)$. This follows almost immediately from the definition of maintaining prescience (see Lemma 1 in the Appendix for a proof). In words: at the beginning of each run of the cycle, the portion of the pattern that the generator anticipates to next produce remains in alignment with the portion of the pattern that the extractor anticipates to next consume.

We have thus specified the action of pattern manipulators in terms of the initial and final states of their related information variables (i.e., state of the tape and state of the internal prescient memory). In the following sections, we provide bounds on the minimum work costs of *any* process in any framework consistent with Landauer's principle that implements manipulations according to this specification.

VII. INVESTING WORK TO GENERATE A PATTERN

For any given pattern, there is a family of generators, characterized by their choice of prescient memory R and by the number of steps k of the pattern that they generate at once. In Theorem 2 of the Appendix, we prove that for any such generator, the work investment $W_{\text{gen}}^k(R)$ required to generate k steps of the pattern when the tape is subject to a degenerate Hamiltonian (that is, all configurations of the tape have equal energy), and the tape-memory system is coupled to a heat bath with inverse temperature $\beta = \frac{1}{k_B T}$, is bounded by

$$\begin{aligned} \beta W_{\text{gen}}^k(R) \geq & k[H(X_{\text{diff}}) - H(X^{t+1} | S^t)] \\ & + [H(R^t | X^{t+1} \dots X^{t+k} R^{t+k}) \\ & - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k})], \end{aligned} \quad (1)$$

FIG. 2. One method of writing a pattern to a tape. The various choices of symbols on a tape can each be associated with a different energy level of a system (drawn as black horizontal lines whose relative height indicates relative energy). The statistical state of the symbol is a probability distribution (drawn as gray bars) over these configurations. By changing the Hamiltonian of the tape while remaining in contact with a thermal reservoir, the statistics can be altered to $P(X^{t+1}|S^t)$. At this point, the system is isolated from the heat reservoir and the Hamiltonian is adiabatically removed. The whole procedure requires an investment of work proportional to the reduction in the state's entropy. (See also Sec. 7 in the Appendix.)

where S^t are the *causal states* of the pattern. The bound holds regardless of what physical mechanism is used to generate the pattern and relies only on the assumption that the generator obeys Landauer's principle. Equality is achieved when the theoretically optimal mechanism is used. If the tape's Hamiltonian is not degenerate, the minimum work investment requires an additional contribution $\Delta E := k[E(X) - E(X_{\text{dflt}})]$, where $E(X)$ is the expectation value of a tape system's energy when configured into the pattern X^t , and $E(X_{\text{dflt}})$ is the expectation value of the tape system's energy in its default state X_{dflt} .

We can divide the work cost of generating a pattern into two contributions: one from changing the entropy (*and thus the free energy*) of the systems on the tape, and one from updating the internal memory. The cost W_{tape}^k associated with writing k symbols onto the tape (at inverse temperature $\beta = \frac{1}{k_B T}$) is given by the first line of Eq. (1):

$$\beta W_{\text{tape}}^k \geq k[H(X_{\text{dflt}}) - H(X^{t+1}|S^t)]. \quad (2)$$

This value is determined by the distribution over causal states S (recall that this is unique for any given pattern [2]) rather than the specific device-dependent internal memory R . As such, W_{tape}^k has no dependence on the choice of prescient memory, and scales trivially with the stride k , and is therefore an intrinsic property of the pattern rather than of the machine generating it (see Fig. 2 for one possible physical realization of this cost using the framework of, e.g., Ref. [27], or the “physical example” in Sec. 7 of the Appendix). (If the systems onto which the pattern is encoded do not have a degenerate Hamiltonian, the additional energy term ΔE should also be associated with the tape, and incorporated into W_{tape}^k).

The remaining contribution W_{diss}^k corresponds to the cost of updating the internal memory from R^t to R^{t+k} so that the generator maintains prescience. This is bounded by

$$\beta W_{\text{diss}}^k(R) \geq H(R^t | X^{t+1} \dots X^{t+k} R^{t+k}) - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}), \quad (3)$$

which is always nonnegative for all k , and all choices of memory R (see Lemma 5). Figure 3 illustrates one possible realization of this limit, but we stress that the bound given in

FIG. 3. One method of updating the generator's memory. (The case for unifilar R is shown.) A blank pure ancilla state R_{dflt} is updated at no cost to R^{t+k} , conditioned on the values of the initial internal state R^t and systems on the tape $X^{t+1} \dots X^{t+k}$. The old state R^t is then *decorrelated* from R^{t+1} and $X^{t+1} \dots X^{t+k}$, and the mutual information used to reduce the entropy of the internal state (transforming it into $\tilde{R}$). Finally, $\tilde{R}$ is *reset* back to the blank ancilla state at work cost $H(R^t | X^{t+1} \dots X^{t+k} R^{t+k})$, so that the generator's internal state is ready to produce the next part of the pattern. (See also Sec. 7 in the Appendix.)

Eq. (1) is general for any method of prescient pattern generation. In particular, any nondegeneracy within the internal memory's Hamiltonian does not play a role in this quantity, since the process is assumed to be stationary: the average memory state remains the same at each time-step and so the memory's average energy does not change from step to step.

The first term in Eq. (3) represents the cost of erasing the previous state R^t , offset by the mutual information between the new state of the memory R^{t+k} and the patterned outputs $X^{t+1} \dots X^{t+k}$.

The second term reflects the effect of (non-)unifilarity. Memory R is defined as *unifilar* (see, e.g., Ref. [34]) if $H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}) = 0$. If the update is not unifilar,³ according to Landauer's principle we can recover a portion of the work cost associated with the memory's change in entropy. However, the first term in Eq. (3) may be expanded as $H(R^t | X^{t+1} \dots X^{t+k} R^{t+k}) = H(X^{t+1} \dots X^{t+k} | R^t) - H(X^{t+1} \dots X^{t+k} | R^{t+k}) + H(R^{t+k} | X^{t+1} \dots X^{t+k} R^t)$, allowing us to alternatively express the bound on dissipated work as

$$\beta W_{\text{diss}}^k(R) \geq H(X^{t+1} \dots X^{t+k} | R^t) - H(X^{t+1} \dots X^{t+k} | R^{t+k}), \quad (4)$$

[see Lemma 4(iii) in the Appendix]. Here, it can be seen that the nonunifilar term has been explicitly canceled. Whatever work might have been gained by introducing randomness into R^{t+k} is entirely canceled out by the cost of resetting this randomness in the previous state R^t . Hence, unifilarity does not *per se* play

³Causal states are automatically unifilar [2], but we do not need to make this assumption for the broader class of memory considered in this article.

a role in deciding the thermodynamic advantage of prescient memory.

Rather, we see that the dissipative cost is proportional to the difference between R^t 's predictive power to guess the next k symbols, and R^{t+k} 's *retrodictive* power to remember the preceding k symbols. Failure to predict increases the first entropy and hence the amount of dissipation. On the other hand, failure to retrodict increases the second entropy and lessens the total dissipation.

VIII. EXTRACTING WORK FROM A PATTERN

Let us now evaluate how much work we can extract by consuming a pattern, by considering the prescient *extractor*. Recall that the extractor takes k systems on the tape from a configuration according to the pattern $X^{t+1} \dots X^{t+k}$ into the default configuration $X_{\text{dft}} \dots X_{\text{dft}}$ and must update its internal memory from R^t to R^{t+k} . In order to fully account for all the changes in entropy in our system, again we consider both the tape on which the pattern is written and the internal memory of the extractor. In the Appendix (see Theorem 3), we prove that the maximum work-output when the tape Hamiltonian is degenerate is bounded by

$$\beta W_{\text{out}}^k \leq k[H(X_{\text{dft}}) - H(X^{t+1} | S^t)]. \quad (5)$$

The work output is entirely proportional to the change in entropy of the k symbols on the tape and has no dependence on the choice of internal memory R . If the tape Hamiltonian is not degenerate, such as would typically be the case in experimental applications, the above acquires the additional term $-\Delta E$, corresponding to a change in the expectation value of the tape's energy that is equal and opposite to the term in generation. Unlike with the generator, it does not matter what sort of memory is used for extraction. Here the cost of updating the memory appears to be zero.

This curiosity may be explained by carefully considering the operational difference between generators and extractors. First, note that these two processes are not exactly the reverse of each other in terms of initial and final states: in both generator and extractors, the memory advances in the same direction from R^t to R^{t+k} in each cycle. Next, observe that at the end of the extraction, there is only one copy of the pattern's relevant past information (encoded in the extractor's internal memory), whereas in generation, this information is retained in both the generator's internal memory and on the tape. The extractor can move the information from the systems on the tape $X^{t+1} \dots X^{t+k}$ into its internal memory R^{t+k} , but the generator must copy this information. Moving information is a logically reversible process, whereas copying information is not—and it is *logical irreversibility* that lies behind the dissipative costs in computation [9]. This subtle, but important, distinction reveals to us why updating the memory must dissipate heat during pattern generation, but not pattern extraction.

IX. SIMPLER IS THERMODYNAMICALLY BETTER

We are now in a position to consider the whole thermodynamic cycle, as illustrated in Fig. 1. What is the minimal amount of heat dissipation in such a cyclic process? Suppose

a generator and extractor make use of a particular encoding map resulting in prescient memory governed by R . In the limit where both devices are implemented optimally at the same inverse temperature β , such that the inequalities (1) and (5) are both saturated, we find that $W_{\text{gen}}^k(R) - W_{\text{out}}^k = \mathcal{W}_{\text{diss}}^k(R)$, where [as per inequality (4)]:

$$\beta \mathcal{W}_{\text{diss}}^k(R) = H(X^{t+1} \dots X^{t+k} | R^t) - H(X^{t+1} \dots X^{t+k} | R^{t+k}). \quad (6)$$

Physically, this represents the minimal amount of heat we must dissipate given a particular choice of prescient memory, given that we generate k steps of the pattern per cycle. We now state two immediate results:

Result I: $\mathcal{W}_{\text{diss}}^k > 0$ whenever $I(R; \vec{X}) > I(\vec{X}; \vec{X})$: If our choice of prescient states stores more information about the past of a pattern than the information that the past contains about the future, then *any* thermodynamic cycle based on such states will be wasteful.

Result II: Out of all the possible prescient states R to use as our internal memory, the simplest ones—corresponding to the causal states S —dissipate the least heat, i.e., minimize $\mathcal{W}_{\text{diss}}^k(R)$. Thus, simpler is thermodynamically better.

We briefly outline the reasoning here, leaving the formal proofs for the Appendix. Consider Eq. (4). The first term is the same for *all* choices of prescient memory, since they all contain equal capacity to reduce uncertainty about future outputs. However, the second term takes its largest value when R^{t+k} contains the least information about the preceding k symbols. This suggests the the less information we retain about the past, the less heat is dissipated. The minimal dissipation is then satisfied by using the causal states.

X. THERMODYNAMICS AND COMPLEXITY

Naïvely, one may expect that it is always possible (in principle) to set $\mathcal{W}_{\text{diss}}^k(R) = 0$ by choosing prescient states such that $I(R; \vec{X}) = I(\vec{X}; \vec{X})$. That is, given that the past of a pattern contains b bits about its future, we can build a fully reversible cycle of pattern generation and extraction by ensuring our pattern manipulators' memories contain no more than b bits about the pattern's past.

Surprisingly, a fundamental result in computational mechanics implies this is not true. Given a generic pattern $P(\vec{X}, \vec{X})$, its statistical complexity $C_\mu = H(S) = I(\vec{X}; S)$ is generally strictly greater than the amount of information its past contains about its future, $I(\vec{X}; \vec{X})$. That is, for most patterns, even the most-efficient prescient memory—the causal states—contain superfluous information about the past [20,35]. For these patterns, $I(R; \vec{X}) > I(\vec{X}; \vec{X})$ for all choices of prescient memory R . This yields us our third result:

Result III: A thermodynamically reversible cycle of pattern generation and extraction is impossible for any pattern where $C_\mu > I(\vec{X}; \vec{X})$.

In the Appendix [see Lemma 4(iv)], we also prove that the minimum dissipative cost may be expressed as a difference in

mutual information:

$$\beta \mathcal{W}_{\text{diss}}^k(R) = I(X^{t+1} \dots X^{t+k}; R^{t+k}) - I(X^{t+1} \dots X^{t+k}; R^t). \quad (7)$$

In the instantaneous limit $k = 1$, this recovers a mathematical quantity similar to that introduced by Still *et al.* [36] as the *useless instantaneous nostalgia*. This coincidence is particularly striking, given that the two results are derived using very different mathematical frameworks.

Increasing the number of steps of the pattern processed per cycle k never increases the amount of dissipation *per step of the pattern* (Theorem 9). (Cf. the dissipation *per cycle*, which certainly must not decrease with the step size—Theorem 6.) In the limit of large k , the minimum dissipation of produce a string of k parts of the pattern tends toward $I(R; \vec{X}) - I(\vec{X}; \vec{X})$. Any attempts using smaller k to produce the same number amount of the pattern will dissipate at least as much work as this. Moreover, when most efficient memory—the causal states S —are used, this limiting quantity becomes the *crypticity* of a process $\chi = C_\mu - I(\vec{X}; \vec{X})$, which captures the minimal amount of superfluous information that *any* predicative model of the given pattern must unavoidable store [20,33,34,37], coinciding with a previous result of Wiesner *et al.* [35].

We hope that future work will also bridge the results of this article with specific physical frameworks, such as presented in Ref. [13]. This requires further development in the nascent field of continuous time computational mechanics [38], to identify whether the process in [13] is indeed an example of prescient pattern manipulator (as presented in this text), or if its internal state instead describes another nonprescient type of memory (e.g., containing oracular information [31]).

XI. GENERAL FRAMEWORK OF PATTERN MANIPULATION

The two types of device—generators and extractors—that we have presented in this article are building blocks that can be combined into complex pattern manipulators. The simplest example is the cycle of generation followed by extraction at the same temperature (Fig. 1). This configuration could be considered as charging a battery that stores energy in the form of a pattern, to be later released by the extractor. Since the contributions from writing and consuming the pattern cancel out, the net cost is from the work dissipation whilst updating the generator’s memory. Using the simplest internal memory ensures that the least work is wasted.

However, the cycle presented in Fig. 1 could be modified such that the generator and extractor act at *different* temperatures. In the scenario where the extractor bath temperature T_E is greater than the generator bath temperature T_G and $H(X^{t+1} | S^t) < H(X_{\text{dft}})$, and the tape is subject to a degenerate Hamiltonian, the system will function as a *heat engine* with efficiency η given:

$$\eta = 1 - \frac{T_G}{T_E} - \frac{T_G \chi_R(k)}{T_E k [H(X_{\text{dft}}) - H(X^{t+1} | S^t)]}, \quad (8)$$

where $\chi_R(k) = H(R^t | X^{t+1} \dots X^{t+k} R^{t+k}) - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k})$ is the right-hand side of Eq. (6). It follows that the Carnot efficiency $\eta_C = 1 - \frac{T_E}{T_G}$ can only be achieved

FIG. 4. Perturbed coin pattern ϵ -machine ($k = 1$). The “perturbed coin” pattern may be produced by a weighted random walk on this network. The nodes represent the two configurations of the ϵ -machines internal memory $R \in \{s_h, s_t\}$ (i.e., the *causal states*). The directed edges represent the effect of a single time step: a (possible) change in memory configuration, and the setting of the system on the tape to some value. The labels $P | \mathbf{x}$ gives the probability P of transitioning to the particular memory configuration R^{t+1} (conditional on the initial memory configuration R^t), while configuring the system at index $t + 1$ on the tape into the state $\mathbf{x} \in \{H, T\}$.

when the statistical complexity $C_\mu = I(\vec{X}; \vec{X})$ such that the pattern has zero crypticity.

XII. SIMPLE EXAMPLE

Let us explicitly evaluate the thermodynamics of a simple pattern described by the *perturbed coin* process. Envision a coin in a box that takes one of two possible configurations: $\{x\} = \{H, T\}$ (standing for “heads” and “tails,” respectively). At each discrete time step, the box is perturbed such that with probability $p \in (0, 1)$ the coin flips from H to T (or vice versa). The pattern $\vec{X}$ consists of bi-infinite string of random variables describing which side of the coin faces up at each time step.

It is clear that this pattern is Markovian: the statistics of future outputs $\vec{X}$ depend only on the very last output X^t , corresponding to state of coin at the current time step. The past thus divides into two causal states, denoted s_H and s_T , corresponding to the two possible values of X^t . The statistics of the pattern can then be represented by transitions between these causal states (see Fig. 4). By symmetry, $P(S = s_H) = P(S = s_T) = \frac{1}{2}$, and so the process has *statistical complexity* $C_\mu = 1$ bit.

Here we study the thermodynamic quantities when running a thermodynamic cycle at different strides k using a variety of different machines: (1) The simplest pattern manipulators, that stores only the causal states, (2) the “last two” machines that stores the last two values the coin took (i.e., setting $R = X^{t-1} X^t$ as per Fig. 5), and (3) a more general “last N ” machine that stores the last N outputs (i.e., setting $R = X^{t-N+1} \dots X^{t-1} X^t$). The cases $N = 1$ and $N = 2$ correspond to the ϵ -machine and “last two” machine, respectively.

The change in free energy for k -symbols of the pattern [Eq. (2)], extractable work when consuming k -symbols of the pattern [Eq. (5)], and unavoidable heat dissipation are summarized in Table I. Detailed calculations can be found in Sec. 6 of the Appendix. We make several observations:

(1) The free energy change of the tape from Eq. (2) and the extractable work from Eq. (5) are always equal and opposite. Although this cost must be accounted for when running

FIG. 5. Perturbed coin process “last two” machine ($k = 1$). The random walk on this network produces exactly the same pattern as the ϵ -machine (Fig. 4). This machine’s memory corresponds to the last two outcomes of the pattern.

the generator or extractor in isolation, the two cancel in a thermodynamic cycle.

(2) Choosing different X_{dft} affects the amount of work required to write to the tape, and the amount of work generated when the tape is consumed, but does not change the total heat dissipated per cycle. Suppose that X_{dft} is maximally mixed, such that $H(X_{\text{dft}}) = 1$ bit. In this case, work must be invested to write the pattern to the tape (except when $p = \frac{1}{2}$). If instead X_{dft} is always heads, such that $H(X_{\text{dft}}) = 0$, then we would instead extract $h(p) := -p \log p - (1-p) \log(1-p)$ units of work from each time step (at the cost of reducing the free energy of the tape). Another interesting choice would be to set $H(X_{\text{dft}}) = h(p)$, such that the free energy of the tape does not change during generation or consumption, and only the dissipative terms remain. In all cases, the total amount of heat dissipated over the entire cycle is the same.

(3) Simpler is better. Out of all possible last- N machines, the one that minimizes the dissipation per cycle occurs when $N = 1$, corresponding to the simplest machines that store only the current causal state of the past. Heat dissipation is monotonically nondecreasing with N . Note, however, there are subtleties: For a given stride k , the penalty for remembering

more of the past ‘saturates’ when $N = k$. This is because when the stride $k < N$, part of the memory can be updated in a logically reversible way (see appendix for details).

(4) If one wants to create as much of the pattern as possible with minimal heat dissipation per output symbol, it is generally advantageous to use large strides. For example, when using ϵ -machines, the heat dissipation remains fixed at $h(p)$ for all k . Thus, that generating k outputs per cycle will reduce the dissipation per symbol by a factor of k .

Finally, the ultimate limits on heat dissipation can be immediately evaluated by noting that the crypticity of the pattern is nonzero. In particular the pattern has crypticity $\chi = C_\mu - I(\vec{X}, \vec{X}) = h(p)$. This bounds the minimal dissipation regardless of stride or what machine is used.

In Sec. 7 of the Appendix, we provide a discussion of this pattern manipulation, as it could be implemented in a more immediately physical framework. There, we provide one mechanism for implementing pattern manipulations at work costs that saturate the bounds of the first row of Table I, and we consider explicitly the efficiency of a heat engine constructed from a cycle of generation followed by extraction.

XIII. DISCUSSION AND OUTLOOK

Here we have presented a formal framework for treating patterns as thermodynamic resources that require energy to synthesize, which can then in turn be used to perform work. This involved the study of generators that convert work into a pattern and extractors that consume patterns to perform useful work. These components were then combined to construct a full thermodynamic cycle of pattern generation and consumption. We then identified the dissipative work costs involved within such a cycle and related it to the complexity of the generators and extractors used. We found that *simpler is thermodynamically better*; the less past information retained, the less dissipation. Optimality is achieved when we retain just enough information to be able to replicate desired operational behavior. The resultant unavoidable dissipation is then given by the crypticity of the pattern. These relations present thermodynamic interpretations for a fundamental complexity-theoretic property.

TABLE I. Minimum thermodynamic costs for perturbed coin process. Here $h(p) := -p \log p - (1-p) \log(1-p)$ denotes the binary entropy. To arrive at the value in these columns, we calculate that $H(X^{t+1} | S^t) = h(p)$ by noting that both H and T states are equally likely to occur, and both have an uncertainty of $h(p)$ as to the value of the next output. It can be seen immediately from the form of Eqs. (2) and (5) that these values scale linearly with k : the number of pattern elements written or consumed per cycle.

Machine name	Tape cost βW_{tape}^k Ineq. (2)	Extractable work βW_{out}^k Ineq. (5)	Dissipation βW_{diss}^k Ineqs. (3), (4), (7), or (6)	Dissipation per output $\beta W_{\text{diss}}^k / k$
ϵ -machine ($k = 1$)	$[H(X_{\text{dft}}) - h(p)]$	$[H(X_{\text{dft}}) - h(p)]$	$h(p)$	$h(p)$
ϵ -machine ($k = 2$)	$2[H(X_{\text{dft}}) - h(p)]$	$2[H(X_{\text{dft}}) - h(p)]$	$h(p)$	$\frac{h(p)}{2}$
ϵ -machine ($k \geq 3$)	$k[H(X_{\text{dft}}) - h(p)]$	$k[H(X_{\text{dft}}) - h(p)]$	$h(p)$	$\frac{h(p)}{k}$
“Last two” ($k = 1$)	$[H(X_{\text{dft}}) - h(p)]$	$[H(X_{\text{dft}}) - h(p)]$	$h(p)$	$h(p)$
“Last two” ($k = 2$)	$2[H(X_{\text{dft}}) - h(p)]$	$2[H(X_{\text{dft}}) - h(p)]$	$2h(p)$	$h(p)$
“Last two” ($k \geq 3$)	$k[H(X_{\text{dft}}) - h(p)]$	$k[H(X_{\text{dft}}) - h(p)]$	$2h(p)$	$\frac{2}{k}h(p)$
“Last N ” ($k \leq N$)	$k[H(X_{\text{dft}}) - h(p)]$	$k[H(X_{\text{dft}}) - h(p)]$	$kh(p)$	$h(p)$
“Last N ” ($k > N$)	$k[H(X_{\text{dft}}) - h(p)]$	$k[H(X_{\text{dft}}) - h(p)]$	$Nh(p)$	$\frac{N}{k}h(p)$

One can extend this framework to consider alternative scenarios where generators and extractors act in parallel on different patterns. For example, consider a scenario where an extractor consumes one pattern, and the energy released is then used to power a generator that writes a different pattern. This could approximate the actions of a living organism: for example, a lion metabolizes the structure of an antelope (destroying it in the process) and uses the energy released to build more lion. Using the simplest internal memory for generation grants the advantage that less antelope needs to be consumed in order to produce the same amount of lion.

A foundational question of interest is whether this unavoidable heat dissipation is truly fundamental, or could it somehow be surpassed with more exotic information processing. Recent research indicates that quantum processors can generate certain statistical patterns of behavior more simply [39], resulting in a push to generalize computational mechanics into the quantum regime [40–44]. Could this simplicity also yield thermodynamic advantage? Any such results would present exciting new thermodynamic signatures of quantumness.

ACKNOWLEDGMENTS

We are grateful for funding from the John Templeton Foundation Grant No. 53914 “*Occam’s Quantum Mechanical Razor: Can Quantum theory admit the Simplest Understanding of Reality?*”; the Foundational Questions Institute (“Physics of the Observer” large grant); the National Research Foundation, Prime Minister’s Office, Singapore, under its Competitive Research Programme (CRP Award No. NRF-CRP14-2014-02); the Ministry of Education in Singapore, the Academic Research Fund Tier 3 Grant No. MOE2012-T3-1-009; the Australian Research Council (Project No. DP160101911); the National Basic Research Program of China Grants No. 2011CBA00300 and No. 2011CBA00302; the National Natural Science Foundation of China Grants No. 11450110058, No. 61033001, and No. 61361136003; the 1000 Talents Program of China; and the Oxford Martin School. We thank Blake Pollard, Merlin Seller, and Whei-Yeap Suen for illuminating discussions.

APPENDIX

1. Prescient pattern manipulators

Lemma 1 (Cycle remains in step). Consider a generator with prescient memory R and an extractor with prescient memory R' . At time t , let these memories initially be in configurations r and r' respectively, satisfying $P(\vec{X}^t | R = r) = P(\vec{X}^t | R' = r')$. After k symbols are produced by the generator and then subsequently consumed by the extractor, the two devices remain “in step” such that $P(\vec{X}^{t+k} | R = \tilde{r}) = P(\vec{X}^{t+k} | R' = \tilde{r}')$ for the updated configurations $\tilde{r}$ and $\tilde{r}'$.

Proof. This follows almost immediately from the fact that both devices maintain the prescience of their memory. When a prescient generator produces k steps of the pattern $X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k}$, it updates its memory to a new prescient state with configuration $\tilde{R} = \tilde{r}$ that satisfies

$$\begin{aligned} P(\vec{X}^{t+k} | \tilde{R} = \tilde{r}) \\ = P(\vec{X}^{t+k} | R = r, X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k}). \end{aligned} \quad (\text{A1})$$

The generator can only produce sequences $x^{t+1} \dots x^{t+k}$ that satisfy $P(X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k} | R = r) > 0$. Similarly, when the prescient extractor receives a sequence, it also updates its internal memory to some state $\tilde{R}' = \tilde{r}'$ such that

$$\begin{aligned} P(\vec{X}^{t+k} | \tilde{R}' = \tilde{r}') \\ = P(\vec{X}^{t+k} | R' = r', X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k}). \end{aligned} \quad (\text{A2})$$

We now show that if the devices are initially in step, such that $P(\vec{X}^t | R = r) = P(\vec{X}^t | R' = r')$, the above updates involving the same string will keep them in step.

The semi-infinite future of a pattern expected at the initial time t , $\vec{X}^t$, can be split into a finite string of length k , $X^{t+1} \dots X^{t+k}$, and another semi-infinite string $\vec{X}^{t+k} = X^{t+k+1} \dots$. Noting that both r and r' initially predict the same statistics for $\vec{X}^t$ (and hence also for any finite substring thereof), for any realization of the future $\vec{x}^t = x^{t+1} \dots x^{t+k} \vec{x}^{t+k}$, it follows from the probability chain rule on $\vec{X}^t$ conditioned on the initial state of both devices:

$$\begin{aligned} P(\vec{X}^t = \vec{x}^t | R = r) &= P(\vec{X}^t = \vec{x}^t | R' = r') \\ &= P(X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k} | R = r) \\ &\quad \cdot P(\vec{X}^{t+k} = \vec{x}^{t+k} | R = r, X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k}), \\ &= P(X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k} | R' = r') \\ &\quad \cdot P(\vec{X}^{t+k} = \vec{x}^{t+k} | R' = r', X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k}), \end{aligned} \quad (\text{A3})$$

and since also

$$\begin{aligned} P(X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k} | R = r) \\ = P(X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k} | R' = r'), \end{aligned} \quad (\text{A4})$$

we can then conclude that

$$\begin{aligned} P(\vec{X}^{t+k} = \vec{x}^{t+k} | R = r, X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k}) \\ = P(\vec{X}^{t+k} = \vec{x}^{t+k} | R' = r', X^{t+1} \dots X^{t+k} = x^{t+1} \dots x^{t+k}), \end{aligned} \quad (\text{A5})$$

for arbitrary $x^{t+1} \dots x^{t+k}$, and all possible futures $\vec{x}^{t+k}$. Thus, substituting in Eqs. (A1) and (A2):

$$P(\vec{X}^{t+k} | \tilde{R} = \tilde{r}) = P(\vec{X}^{t+k} | \tilde{R}' = \tilde{r}'). \quad (\text{A6})$$

The devices hence remain in step after each cycle. ■

2. Proofs of thermodynamic costs

Theorem 2 (Work cost of generation). For a generator of a pattern $\vec{X}$ that maintains prescient memory R , the work cost W_{gen}^k to configure k systems according to the pattern is bounded by

$$\begin{aligned} \beta W_{\text{gen}}^k(R) &\geq k[H(X_{\text{dft}}) - H(X^{t+1} | S^t)] \\ &\quad + H(R^t | X^{t+1} \dots X^{t+k} R^{t+k}) \\ &\quad - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}). \end{aligned} \quad (\text{A7})$$

Proof. We prove this using information theoretical means. In the generation process, the generator's internal memory is initially distributed according to R^t , and the k systems on the tape according to $X_{\text{dft}} \dots X_{\text{dft}}$. After generation, the final state of the generator's internal memory is R^{t+k} , and the k systems on the tape are in states $X^{t+1} \dots X^{t+k}$.

As such, a device-independent lower bound on the cost of generation may be found by considering Landauer's principle as applied to these states

$$\beta W_{\text{gen}}(R) \geq H(R^t X_{\text{dft}} \dots X_{\text{dft}}) - H(R^{t+k} X^{t+1} \dots X^{t+k}). \quad (\text{A8})$$

First, we note that $H(R^t X^{t+1} \dots X^{t+k} R^{t+k})$ may be expanded using the chain rule in two different ways:

$$\begin{aligned} H(R^t X^{t+1} \dots X^{t+k} R^{t+k}) &= H(R^t) + H(X^{t+1} | R^t) + \dots \\ &\quad + H(X^{t+k} | R^t X^{t+1} \dots X^{t+k-1}) \\ &\quad + H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}) \\ &= H(R^{t+k}) + H(X^{t+k} | R^{t+k}) + \dots \\ &\quad + H(X^{t+1} | R^{t+k} X^{t+k} \dots X^{t+2}) \\ &\quad + H(R^t | R^{t+k} X^{t+k} \dots X^{t+1}). \end{aligned} \quad (\text{A9})$$

This allows us to re-express the final state entropy $H(R^{t+k} X^{t+1} \dots X^{t+k})$ in the form

$$\begin{aligned} H(R^{t+k} X^{t+1} \dots X^{t+k}) &= H(R^{t+k}) + H(X^{t+k} | R^{t+k}) + \dots \\ &\quad + H(X^{t+1} | R^{t+k} X^{t+k} \dots X^{t+2}), \\ &= H(R^t) + H(X^{t+1} | R^t) + \dots \\ &\quad + H(X^{t+k} | R^t X^{t+1} \dots X^{t+k-1}) \\ &\quad + H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}) \\ &\quad - H(R^t | R^{t+k} X^{t+k} \dots X^{t+1}). \end{aligned} \quad (\text{A10})$$

Moreover, since the internal memory R is prescient about the future of X , and all choices of prescient memory must correspond to a fine-graining of the causal states S^t (see Lemma 7 in Ref. [2]). As such $R^t X^{t+1} \dots X^{t+j}$ is sufficient to perfectly determine S^{t+j} , which is in turn as exactly as useful as $R^t X^{t+1} \dots X^{t+j}$ for predicting values of X^{t+j+1} onwards. Thus, the term $H(X^{t+j} | S^{t+j-1})$ may be substituted for every addend of the form $H(X^{t+j} | R^t X^{t+1} \dots X^{t+j-1})$ in Eq. (A10). Moreover, by stationarity, $H(X^{t+1} | S^t) = H(X^{t+2} | S^{t+1}) = \dots$, and hence the sum over k of these terms may be combined into the single expression $kH(X^{t+1} | S^t)$. As such, we can re-express the final entropy following generation as

$$\begin{aligned} H(R^{t+k} X^{t+1} \dots X^{t+k}) &= H(R^t) + kH(X^{t+1} | S^t) \\ &\quad + H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}) \\ &\quad - H(R^t | R^{t+k} X^{t+k} \dots X^{t+1}). \end{aligned} \quad (\text{A11})$$

Meanwhile, the initial entropy $H(R^t X_{\text{dft}} \dots X_{\text{dft}})$ can straightforwardly be written

$$H(R^t X_{\text{dft}} \dots X_{\text{dft}}) = H(R^t) + kH(X_{\text{dft}}), \quad (\text{A12})$$

since there are no correlations within the initial state.

Hence, from Landauer's principle, the difference between these two quantities yields the minimum work exchange required to generate a pattern and update the generator's internal memory:

$$\begin{aligned} \beta W_{\text{gen}}^k(R) &\geq k[H(X_{\text{dft}}) - H(X^{t+1} | S^t)] \\ &\quad + H(R^t | X^{t+1} \dots X^{t+k} R^{t+k}) \\ &\quad - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}). \end{aligned} \quad (\text{A13})$$

Theorem 3 (Work output from extraction). No extractor with prescient internal memory R can extract from k symbols of a pattern $\vec{X}$, more work W_{out}^k than the following bound:

$$\beta W_{\text{out}}^k \leq k[H(X_{\text{dft}}) - H(X^{t+1} | S^t)]. \quad (\text{A14})$$

The choice of internal memory plays no limiting role in determining the amount of work extractable.

Proof. As with Theorem 2, we consider the fundamental limitations placed by Landauer's principle. A cyclically operating prescient extractor must result in the transformation of the tape from $X^{t+1} \dots X^{t+k}$ to $X_{\text{dft}} \dots X_{\text{dft}}$, while updating its internal memory from R^t to R^{t+k} . As such, the bound from Landauer's principle is

$$\beta W_{\text{out}}^k \leq H(R^{t+k} X_{\text{dft}} \dots X_{\text{dft}}) - H(R^t X^{t+1} \dots X^{t+k}). \quad (\text{A15})$$

In this case, the final entropy $H(R^{t+k} X_{\text{dft}} \dots X_{\text{dft}})$ is easy to calculate, since there is no correlation between any of the variables:

$$H(R^{t+k} X_{\text{dft}} \dots X_{\text{dft}}) = H(R^{t+k}) + kH(X_{\text{dft}}). \quad (\text{A16})$$

On the other hand, the entropy of the initial state $H(R^t X^{t+1} \dots X^{t+k})$ is lower than if the variables were independent, since R^t contains some information that can be used to infer $X^{t+1} \dots X^{t+k}$. We can use the chain rule expansion to see that this entropy is

$$\begin{aligned} H(R^t X^{t+1} \dots X^{t+k}) &= H(R^t) + H(X^{t+1} | R^t) + \dots \\ &\quad + H(X^{t+k} | R^t X^{t+1} \dots X^{t+k-1}). \end{aligned} \quad (\text{A17})$$

Using exactly the same logic as in Theorem 2, each term of the form $H(X^{t+j} | R^t X^{t+1} \dots X^{t+j-1})$ can be substituted with a term $H(X^{t+j} | S^{t+j-1})$ (because the memory is prescient), and from stationarity these terms are all the same, and so may be collectively replaced by $kH(X^{t+1} | S^t)$. Thus, the initial entropy is

$$H(R^t X^{t+1} \dots X^{t+k}) = H(R^t) + kH(X^{t+1} | S^t). \quad (\text{A18})$$

Finally, from stationarity we note that $H(R^t) = H(R^{t+k})$ and so computing the difference in these entropies, and applying Landauer's principle, we arrive at the bound

$$\beta W_{\text{out}}^k \leq k[H(X_{\text{dft}}) - H(X^{t+1} | S^t)], \quad (\text{A19})$$

which has no explicit dependence on R . ■

3. Dissipative work term

As discussed in the article, since the work cost in the first line of Eq. (15) of Ref. [2] has no dependence on the choice of

memory used, but is entirely a function of the pattern on the tape, and moreover is exactly equal to the energy that may be recovered according to Ref. [3], we may naturally divide the cost into two parts:

$$\beta W_{\text{tape}}^k \geq k[H(X_{\text{dft}}) - H(X^{t+1} | S^t)], \quad (\text{A20})$$

$$\beta W_{\text{diss}}^k(R) \geq H(R^t | X^{t+1} \dots X^{t+k} R^{t+k}) - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}). \quad (\text{A21})$$

In an optimally implemented cycle of generation and extraction of the same pattern using the same type of memory R in both devices (where both run at the same temperature), the change in work is given by the quantity that saturates Eq. (A21). For notational convenience, we denote this minimum difference as $\chi_R(k)$,

$$\chi_R(k) := H(R^t | X^{t+1} \dots X^{t+k} R^{t+k}) - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}), \quad (\text{A22})$$

such that Theorem 2 may be rephrased as

$$\beta W_{\text{diss}}^k(R) \geq \chi_R(k). \quad (\text{A23})$$

We can garner some additional physical insight about $\chi_R(k)$, by expressing it in a few alternative forms, listed below:

Lemma 4 (Equivalent forms of $\chi_R(k)$). For prescient memory R , the following expressions are equal:

(1) Memory reset-nonunifilarity:

$$\chi_R(k) = H(R^t | X^{t+1} \dots X^{t+k} R^{t+k}) - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}), \quad (\text{A24})$$

(2) Past memory-future memory uncertainty:

$$\chi_R(k) = H(R^t | X^{t+1} \dots X^{t+k}) - H(R^{t+k} | X^{t+1} \dots X^{t+k}). \quad (\text{A25})$$

(3) Predictive-retrodictive uncertainty:

$$\chi_R(k) = H(X^{t+1} \dots X^{t+k} | R^t) - H(X^{t+k} \dots X^{t+1} | R^{t+k}). \quad (\text{A26})$$

(4) Retrodictive-predictive information:

$$\chi_R(k) = I(X^{t-k+1} \dots X^t; R^t) - I(X^{t+k} \dots X^{t+1}; R^t). \quad (\text{A27})$$

(5) Pattern-memory-memory-pattern block entropy:

$$\chi_R(k) = H(X^{t+1} \dots X^{t+k} R^t) - H(X^{t+k} \dots X^{t+1} R^{t+k}). \quad (\text{A28})$$

(6) Difference in memory retrodictability:

$$\chi_R(k) = H(R^t | \vec{X}) - H(R^{t+k} | \vec{X}). \quad (\text{A29})$$

Proof. (i) We take the form of 1 as our initial definition of $\chi_R(k)$:

$$\chi_R(k) := H(R^t | X^{t+1} \dots X^{t+k} R^{t+k}) - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}). \quad (\text{A30})$$

(ii) We may expand $H(R^t X^{t+1} \dots X^{t+k} R^{t+k})$ in two ways:

$$\begin{aligned} & H(R^t X^{t+1} \dots X^{t+k} R^{t+k}) \\ &= H(X^{t+1} \dots X^{t+k}) + H(R^t | X^{t+1} \dots X^{t+k}) \end{aligned}$$

$$\begin{aligned} & + H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}) \\ &= H(X^{t+1} \dots X^{t+k}) + H(R^{t+k} | X^{t+1} \dots X^{t+k}) \\ & + H(R^t | R^{t+k} X^{t+1} \dots X^{t+k}). \end{aligned} \quad (\text{A31})$$

Since the terms $H(X^{t+1} \dots X^{t+k})$ cancel, we see

$$\begin{aligned} & H(R^t | R^{t+k} X^{t+1} \dots X^{t+k}) - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}) \\ &= H(R^t | X^{t+1} \dots X^{t+k}) - H(R^{t+k} | X^{t+1} \dots X^{t+k}). \end{aligned} \quad (\text{A32})$$

The top term is i, the bottom ii, hence these are equal.

(iii) A different expansion of the joint entropy is

$$\begin{aligned} & H(R^t X^{t+1} \dots X^{t+k} R^{t+k}) \\ &= H(R^t) + H(X^{t+1} \dots X^{t+k} | R^t) \\ & + H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}) \\ &= H(R^{t+k}) + H(X^{t+k} \dots X^{t+1} | R^{t+k}) \\ & + H(R^t | R^{t+k} X^{t+k} \dots X^{t+1}). \end{aligned} \quad (\text{A33})$$

Using stationarity to set $H(R^t) = H(R^{t+k})$,

$$\begin{aligned} & H(R^t | R^{t+k} X^{t+1} \dots X^{t+k}) - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}) \\ &= H(X^{t+1} \dots X^{t+k} | R^t) - H(X^{t+k} \dots X^{t+1} | R^{t+k}). \end{aligned} \quad (\text{A34})$$

The top term is i, the bottom is iii, hence these expressions are equal.

(iv) From the definition of mutual information $I(A; B) = H(A) - H(A | B)$, and using stationarity, we can re-express $\chi_R(k)$ as

$$\begin{aligned} & H(X^{t+1} \dots X^{t+k} | R^t) - H(X^{t+k} \dots X^{t+1} | R^{t+k}) \\ &= I(X^{t+1} \dots X^{t+k}; R^{t+k}) - I(X^{t+k} \dots X^{t+1}; R^t), \\ &= I(X^{t-k+1} \dots X^t; R^t) - I(X^{t+k} \dots X^{t+1}; R^t). \end{aligned} \quad (\text{A35})$$

The top term is iii, the bottom is iv, hence these expressions are equal.

(v) Again, we expand

$$\begin{aligned} & H(R^t X^{t+1} \dots X^{t+k} R^{t+k}) \\ &= H(R^t X^{t+1} \dots X^{t+k}) + H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}) \\ &= H(R^{t+k} X^{t+k} \dots X^{t+1}) + H(R^t | R^{t+k} X^{t+k} \dots X^{t+1}), \end{aligned} \quad (\text{A36})$$

such that

$$\begin{aligned} & H(R^t | R^{t+k} X^{t+1} \dots X^{t+k}) - H(R^{t+k} | R^t X^{t+1} \dots X^{t+k}) \\ &= H(X^{t+1} \dots X^{t+k} R^t) - H(X^{t+k} \dots X^{t+1} R^{t+k}). \end{aligned} \quad (\text{A37})$$

The top term is i, the bottom is v, hence these expressions are equal.

(vi) Showing this last form is equivalent is slightly more involved. It may be proven by adapting and generalizing Theorem 1 of Mahoney *et al.* [37] beyond causal states into general (possibly nonunifilar) memory R .

For some $l > k$, consider the expansions of the two terms:

$$\begin{aligned} & H(R^t X^{t+1} \dots X^{t+l}) = H(X^{t+1} \dots X^{t+l}) \\ & + H(R^t | X^{t+1} \dots X^{t+l}), \end{aligned} \quad (\text{A38})$$

$$H(R^{t+k} X^{t+1} \dots X^{t+l}) = H(X^{t+1} \dots X^{t+l}) + H(R^{t+k} | X^{t+1} \dots X^{t+l}). \quad (\text{A39})$$

The difference between these two terms is

$$\begin{aligned} \Delta &:= H(R^t X^{t+1} \dots X^{t+l}) - H(R^{t+k} X^{t+1} \dots X^{t+l}) \\ &= H(R^t X^{t+1} \dots X^{t+k}) \\ &\quad + H(X^{t+k+1} \dots X^{t+l} | R^t X^{t+1} \dots X^{t+k}) \\ &\quad - H(R^{t+k} X^{t+1} \dots X^{t+k}) \\ &\quad - H(X^{t+k+1} \dots X^{t+l} | R^{t+k} X^{t+1} \dots X^{t+k}). \quad (\text{A40}) \end{aligned}$$

Both the conditional entropy terms are the conditional entropy of the future of the pattern (steps $t+k+1$ to $t+l$) with respect to the part of the pattern (and memory) that they are conditioned on (steps t to $t+k$). This allows us to use the fact that R is prescient to argue that $H(\vec{X}^{t+k} | R^{t+k} X^{t+1} \dots X^{t+k}) = H(\vec{X}^{t+k} | R^{t+k})$, since R^{t+k} already contains all the information to predict $\vec{X}^{t+k}$, the additional $X^{t+1} \dots X^{t+k}$ are redundant—a property of prescient memory known as *causal shielding* [45]. Similarly, $H(\vec{X}^{t+k} | R^t X^{t+1} \dots X^{t+k}) = H(\vec{X}^{t+k} | R^{t+k})$ (see discussion within Theorem 2). Thus, these two terms are the same, and we can hence simplify the expression to

$$\begin{aligned} \Delta &= H(R^t | X^{t+1} \dots X^{t+l}) - H(R^{t+k} | X^{t+1} \dots X^{t+l}) \\ &= H(R^t X^{t+1} \dots X^{t+k}) - H(R^{t+k} X^{t+1} \dots X^{t+k}). \quad (\text{A41}) \end{aligned}$$

Thus, $\Delta = \chi_R(k)$ in the form given by v. Since this equality is true for all $l > k$, we can then take the limit $l \rightarrow \infty$ (such that $X^{t+1} \dots X^{t+l} \rightarrow \vec{X}$) and arrive at

$$\chi_R(k) = H(R^t | \vec{X}) - H(R^{t+k} | \vec{X}), \quad (\text{A42})$$

proving v and vi are equivalent. ■

4. Memory k -step crypticity

The crypticity of a pattern [34,37] is given $\chi = H(S^t | \vec{X})$. This property can be generalized to a *memory crypticity* χ_R , defined

$$\chi_R = H(R^t | \vec{X}) - H(R^t | \vec{X}). \quad (\text{A43})$$

The second term subtracts any uncertainty in the memory state having observed the entire sequence of pattern to date. For causal states $R = S$, $\chi_S = \chi$, since $H(S^t | \vec{X}) = 0$.

Using the form of $\chi_R(k)$ in Lemma 4(ii), we see $\lim_{k \rightarrow \infty} \chi_R(k) = \chi_R$. This motivates the naming of the quantity $\chi_R(k)$ as the *memory k -step crypticity*. Moreover, in any of the forms listed in ii when $R = S$, $\chi_S(k) =: \chi(k)$, the pattern's intrinsic k -step crypticity (defined in [37] as the “ k -cryptic approximation”).

We also remark that by treating the $k = 0$ case as an invitation to completely omit the terms “ $X^{t+1} \dots X^{t+k}$ ” in Lemma 4, then $\chi_R(0) = 0$ in every form. (It is established similarly in Ref. [37] that $\chi(0) = 0$). Physically (i.e., when taken with Theorem 2), this is a statement that any “generator” that produces nothing and does not change its memory is theoretically allowed to dissipate no work, regardless of what memory it has.

Despite not necessarily corresponding to causal states, nor necessarily updating in a unifilar manner, the memory k -step crypticity $\chi_R(k)$ shares a few useful properties with the pattern's intrinsic k -step crypticity $\chi(k)$, which we now prove:

Lemma 5 ($\chi_R(k)$ is nonnegative and nondecreasing). For any choice of prescient memory R , $\chi_R(k) \geq 0$ for all $k \geq 0$, and $\chi_R(k) \geq \chi_R(k')$ for all $k \geq k'$.

Proof. First, recall from Theorem 4(vi) that

$$\chi_R(k) = H(R^t | \vec{X}^t) - H(R^{t+k} | \vec{X}^t). \quad (\text{A44})$$

The entropy $H(R^{t+k} | \vec{X}^t) \geq H(R^{t+k+1} | \vec{X}^t)$. This is most obviously seen by using stationarity to rewrite the two terms as

$$H(R^{t+k} | \vec{X}^t) = H(R^t | X^{t-k+1} \dots X^t \vec{X}^t), \quad (\text{A45})$$

$$H(R^{t+k+1} | \vec{X}^t) = H(R^t | X^{t-k} X^{t-k+1} \dots X^t \vec{X}^t). \quad (\text{A46})$$

Since the latter term is conditioned on the same variables as the former, plus an additional variable X^{t-k} , it cannot be higher. This nonincreasing property then implies by induction $H(R^{t+k} | \vec{X}) \leq H(R^t | \vec{X})$. Hence, we see that

$$\chi_R(k) = H(R^t | \vec{X}) - H(R^{t+k} | \vec{X}) \geq 0 \quad (\text{A47})$$

for all k .

Since $\chi_R(k) = H(R^t | \vec{X}^t) - H(R^{t+k} | \vec{X}^t)$ is the difference between a constant term and a nonincreasing term, it follows that $\chi_R(k)$ is nondecreasing. ■

In the context laid out in our article, the above lemma may be interpreted physically:

Theorem 6. For any generator with prescient memory R , the dissipative work cost of generation $W_{\text{diss}}^k(R)$ is *always* nonnegative. In the special case where the process is implemented optimally (i.e., at the limit from Landauer's principle), a generator that produces k steps of the pattern never dissipates less work than a generator that produces $k' < k$ steps of the pattern.

Proof. Recall that

$$\beta W_{\text{diss}}^k(R) \geq \chi_R(k). \quad (\text{A48})$$

From Lemma 5, $\chi_R(k) \geq 0$, hence

$$\beta W_{\text{diss}}^k(R) \geq 0. \quad (\text{A49})$$

When implemented optimally (i.e., at the limit from Landauer's principle), this inequality $\beta W_{\text{diss}}^k(R) \geq \chi_R(k)$ is saturated. In this regime, the monotonic non-decreasing nature of $\chi_R(k)$ (Lemma 5) guarantees monotonic nondecreasing dissipation. ■

We can also show that $\chi_R(k)$ is convex upward for $k \geq 0$. To do this, we first prove the following lemma:

Lemma 7 (Monotonically decreasing retrodiction). If prescient memory R can be used by a generator to produce a pattern, the quantity

$$\rho(k) := \begin{cases} 0 & k = 0, \\ H(X^t | R^t) & k = 1, \\ H(X^t | X^{t+1} \dots X^{t+k-1} R^{t+k-1}) & k > 1 \end{cases} \quad (\text{A50})$$

is nonincreasing with respect to k .

Proof. For causal states, an elegant proof of this presented as Theorem 2 of Ref. [34]. However, as generic memory R is not necessarily unifilar, their method cannot be generalized here. Instead, we supply an alternative proof that relies on the data-processing inequality, which states

$$I(A; B) \geq I(A; f(B)) \quad (\text{A51})$$

for random variables A and B , and a map f that acts on them. (Intuitively, post-processing local data can not increase its nonlocal correlations).

If R is memory that is physically used by a generator to produce a pattern one step at a time, there must be some function $f: X_{\text{dft}} R^t \mapsto R^{t+1} X^{t+1}$ corresponding to the update. Thus, from the data-processing inequality,

$$I(X^t X_{\text{dft}}; R^t) \geq I(X^t X_{\text{dft}}; X^{t+1} R^{t+1}). \quad (\text{A52})$$

Since X_{dft} is uncorrelated with everything, we can omit it from both sides of the equation,

$$I(X^t; R^t) \geq I(X^t; X^{t+1} R^{t+1}), \quad (\text{A53})$$

which in terms of conditional entropy is

$$H(X^t | R^t) \leq H(X^t | X^{t+1} R^{t+1}). \quad (\text{A54})$$

Hence, $\rho(2) \geq \rho(1)$.

Closely related to the update function f is the function $f': X^{t+1} \dots X^{t+k-1} X_{\text{dft}} R^{t+k-1} \mapsto X^{t+1} \dots X^{t+k} R^{t+k}$, which describes the update in the presense of a stretch of the tape $X^{t+1} \dots X^{t+k}$, which remains undisturbed and moreover does not affect the choice of update. This is then a tensor product of the identity function on $X^{t+1} \dots X^{t+k-1}$ with function f on $X_{\text{dft}} R^{t+k-1}$; and so if f is a valid positive map, so too must be f' . Thus, we may once more use the data-processing inequality:

$$\begin{aligned} I(X^t; X^{t+1} \dots X^{t+k-1} X_{\text{dft}} R^{t+k-1}) \\ \geq I(X^t; X^{t+1} \dots X^{t+k-1} X^{t+k} R^{t+k}). \end{aligned} \quad (\text{A55})$$

Taking into account that X_{dft} has no correlations and writing in terms of conditional entropies,

$$\begin{aligned} H(X^t | X^{t+1} \dots X^{t+k-1} R^{t+k-1}) \\ \leq H(X^t | X^{t+1} X^{t+k} R^{t+k}). \end{aligned} \quad (\text{A56})$$

By induction, this proves the claim for $k > 1$. Trivially, $0 \leq H(X^t | R^t)$, because the term is an entropy. Together, these statements show that ρ monotonically increases for all $k \geq 0$. ■

Lemma 8 (Memory-crypticity is convex upwards). $\chi_R(k)$ is convex upwards with respect to k . That is,

$$\chi_R(k) - \chi_R(k-1) \geq \chi_R(k') - \chi_R(k'-1) \quad (\text{A57})$$

when $k \leq k'$.

Proof. We can expand $\chi_R(k)$ in the form of Lemma 4(v):

$$\begin{aligned} \chi_R(k) &= H(X^{t+1} \dots X^{t+k} R^t) - H(X^{t+k} \dots X^{t+1} R^{t+k}) \\ &= H(X^{t+k-1} \dots X^{t+1} R^t) \\ &\quad + H(X^{t+k} | X^{t+k-1} \dots X^{t+1} R^t) \\ &\quad - H(X^{t+2} \dots X^{t+k} R^{t+k}) \\ &\quad - H(X^{t+1} | X^{t+2} \dots X^{t+k} R^{t+k}). \end{aligned} \quad (\text{A58})$$

Using stationarity and rearranging:

$$\begin{aligned} \chi_R(k) &= H(X^{t+k-1} \dots X^{t+1} R^t) \\ &\quad - H(X^{t+1} \dots X^{t+k-1} R^{t+k-1}) \\ &\quad + H(X^{t+k} | X^{t+k-1} \dots X^{t+1} R^t) \\ &\quad - H(X^t | X^{t+1} \dots X^{t+k-1} R^{t+k-1}) \\ &= \chi_R(k-1) + H(X^{t+k} | X^{t+k-1} \dots X^{t+1} R^t) \\ &\quad - H(X^t | X^{t+1} \dots X^{t+k-1} R^{t+k-1}). \end{aligned} \quad (\text{A59})$$

Using the prescience of R , the second term $H(X^{t+k} | X^{t+k-1} \dots X^{t+1} R^t) = H(X^{t+k} | R^{t+k-1})$ and, from stationarity, is equal to $H(X^{t+1} | R^t)$. Thus,

$$\begin{aligned} \chi_R(k) &= \chi_R(k-1) + H(X^{t+1} | R^t) \\ &\quad - H(X^t | X^{t+1} \dots X^{t+k-1} R^{t+k-1}). \end{aligned} \quad (\text{A60})$$

Thus, we find by induction

$$\chi_R(k) = kH(X^{t+1} | R^t) - \Gamma(k), \quad (\text{A61})$$

where

$$\begin{aligned} \Gamma(k) &= H(X^t | R^t) + \sum_{i=1}^{k-1} H(X^t | X^{t+1} \dots X^{t+i} R^{t+i}) \\ &= \sum_{i=0}^k \rho(i), \end{aligned} \quad (\text{A62})$$

where $\rho(i)$ is the expression from Lemma 7, which has been established to be positive and monotonically nondecreasing. The sum $\Gamma(k)$ of such monotonically nondecreasing terms must hence be convex downward for $k \geq 0$. Since $\chi_R(k)$ is then the difference between a linear contribution and a convex downward contribution, it must be convex upward.

This property remains true when we extend the domain of $\chi_R(k)$ to include $\chi_R(0) = 0$ ($\Gamma(0) = \rho(0) = 0$). ■

When taken together with Theorem 2, the above property has the following physical implication:

Theorem 9. For any generator optimally implementing prescient memory R , the minimum work investment *per step* of the pattern produced is never larger if the generator produces a larger string of the pattern in any given run. That is, for a given choice of prescient memory R , it is thermodynamically better to produce as much of the pattern as possible at once.

Proof. Since $\chi_R(k)$ is convex upwards (Lemma 8), including its extension to $\chi_R(0) = 0$, this immediately implies

$$\frac{\chi_R(k)}{k} \geq \frac{\chi_R(k')}{k'} \quad \text{when } 1 \leq k \leq k'. \quad (\text{A63})$$

Taken together with Theorem 2, this immediately implies the claim. ■

We can now show that for any choice of R , either $\chi_R(k) = 0$ for all $k \geq 1$, or $\chi_R(k) \neq 0$ for all $k \geq 1$. That is, it is not possible for only some of $\chi_R(k \geq 1)$ to be zero.

Lemma 10 (All zero). For a given choice of R , if $\chi_R(k) = 0$ for any $k \geq 1$, then $\chi_R(k) = 0$ for all $k \geq 1$. Also, if $\chi_R = \lim_{k \rightarrow \infty} \chi_R(k) = 0$, then $\chi_R(k) = 0$ for all $k \geq 1$. Moreover, if $\chi_R(k) = 0$ for any $k \geq 1$, then $\chi_R = \lim_{k \rightarrow \infty} \chi_R(k) \rightarrow 0$.

Proof. First consider $k' < k$. From Lemma 5, χ_R is nondecreasing and positive. Hence, $\chi_R(k) = 0$ implies that $\chi_R(k') = 0$ for all $k' \leq k$.

Now consider for $k' > k$. From Theorem 9,

$$\frac{\chi_R(k)}{k} \geq \frac{\chi_R(k')}{k'} \text{ when } k \leq k', \quad (\text{A64})$$

and since $\chi_R(k) = 0$,

$$0 \geq \frac{\chi_R(k')}{k'}, \quad (\text{A65})$$

implying that $\chi_R(k') \leq 0$. But Lemma 5 states $\chi_R(k') \geq 0$, and hence we may conclude that $\chi_R(k') = 0$. Thus, if $\chi_R(k) = 0$ for one value of $k \geq 1$, $\chi_R(k) = 0$ for all values of $k \geq 1$.

The second sentence of the claim is simpler to prove. Again, Lemma 5 states that χ_R is nondecreasing and positive. If its limit $\chi_R = \lim_{k \rightarrow \infty} \chi_R(k) = 0$, then $\chi_R(k) = 0$ for all k , since this limit is approached from below.

Finally, to show the last sentence of the claim, we more carefully pick at ineq. (A63) of Theorem 9. Writing as

$$k \frac{\chi_R(k')}{k'} \geq \chi_R(k), \quad k \times 0 \geq \chi_R(k), \quad (\text{A66})$$

where $\frac{\chi_R(k')}{k'}$ is strictly 0 for all k (not, say, some finite expression that becomes vanishingly small with k). Only then may we safely take the limit $k \rightarrow \infty$, such that $\lim_{k \rightarrow \infty} (0 \times k) = 0$ and hence $\lim_{k \rightarrow \infty} \chi_R(k) \leq 0$. It thus follows that $\chi_R = 0$. ■

Corollary 11 (All nonzero). If for some $k \geq 1$, $\chi_R(k) > 0$, then $\chi_R(j) > 0$ for all $j \geq 1$, and $\chi_R = \lim_{k \rightarrow \infty} \chi_R(k) > 0$. Likewise, if $\chi_R = \lim_{k \rightarrow \infty} \chi_R(k) > 0$, then $\chi_R(j) > 0$ for all $j \geq 1$.

Proof. This follows by contradiction with Lemma 10. Suppose $\chi_R(k) > 0$ but either for some $k' \geq 1$, $\chi_R(k') = 0$ or $\chi_R = 0$. Lemma 10 states that this implies $\chi_R(k) = 0$, immediately leading to contradiction. Likewise, if $\chi_R > 0$, but that there was some finite $k \geq 1$ such that $\chi_R(k) = 0$, from Lemma 10 $\chi_R = 0$, leading to contradiction. ■

5. Proofs of main results

This allows us make our first main result:

Result I (Excessive information causes dissipation). $W_{\text{diss}}^k > 0$ whenever $I(R^t; \vec{X}) > I(\vec{X}; \vec{X})$.

Proof. Using the form of χ_R from Lemma 4(iv)

$$\begin{aligned} \beta W_{\text{diss}}^k &\geq \chi_R(k) \\ &= I(X^{t-k+1} \dots X^t; R^t) - I(X^{t+k} \dots X^{t+1}; R^t). \end{aligned} \quad (\text{A67})$$

In the limit of $k \rightarrow \infty$,

$$\lim_{k \rightarrow \infty} \chi_R(k) = I(\vec{X}; R^t) - I(\vec{X}; R^t), \quad (\text{A68})$$

and since R is prescient, we can replace the last term with $I(\vec{X}; \vec{X})$:

$$\lim_{k \rightarrow \infty} \chi_R(k) = I(\vec{X}; R^t) - I(\vec{X}; \vec{X}). \quad (\text{A69})$$

Suppose $I(\vec{X}; R^t) \neq I(\vec{X}; \vec{X})$ such that $\chi_R \neq 0$. Then by Corollary 11, $\chi_R(k) > 0$ for all k . Putting this into Theorem 2, it then immediately follows that $\beta W_{\text{diss}}^k > 0$ for all k . ■

Result II (Simpler is thermodynamically better). For generating k steps of any given pattern $\vec{X}$, the generator's dissipative work cost is minimized by choosing prescient memory R to be in one-to-one correspondence with the pattern's causal states S .

Proof. Consider then Theorem 2, namely $\beta W_{\text{diss}}^k(R) \geq \chi_R(k)$. If implemented at the theoretical optimal limit dictated by Landauer's principle, then equality holds. It then follows that to achieve the optimal thermodynamic performance, one should minimize $\chi_R(k)$.

From Lemma 4(iii),

$$\chi_R(k) = H(X^{t+1} \dots X^{t+k} | R^t) - H(X^{t+k} \dots X^{t+1} | R^{t+k}). \quad (\text{A70})$$

The definition of prescience for R^t tells us that $P(\vec{X} | R^t) = P(\vec{X} | \vec{X})$, and hence for any string of length k , $H(X^{t+1} \dots X^{t+k} | R^t) = H(X^{t+1} \dots X^{t+k} | \vec{X})$. For all choices of memory R (including when it is in one-to-one correspondence with the causal states S) this quantity is the same, and the minimization can be performed entirely by maximizing the second term of Eq. (A70).

Prescience implies an important condition on memory R : namely that it encodes a *refinement* of causal states: no two past histories $\vec{x}$ and $\vec{x}'$ can be mapped to the same state r if they belong to two separate causal states. This property immediately follows from our definition of prescience: if $P(\vec{X} | \vec{X} = \vec{x}) \neq P(\vec{X} | \vec{X} = \vec{x}')$, it is clearly impossible for $P(\vec{X} | R = r)$ to be equal to both. (See also Lemma 7 of Ref. [2].) This refinement property implies the existence of a deterministic map $\Phi: \{r\} \rightarrow \{s\}$, such that $\Phi(R) = S$. Then, we can apply the data processing inequality $I(X^{t+1} \dots X^{t+k}; R^{t+k}) \leq I(X^{t+1} \dots X^{t+k}; \Phi(R^{t+k}))$. Hence,

$$I(X^{t+1} \dots X^{t+k}; R^{t+k}) \leq I(X^{t+1} \dots X^{t+k}; S^{t+k}). \quad (\text{A71})$$

Expanding these mutual informations gives

$$\begin{aligned} &H(X^{t+1} \dots X^{t+k}) - H(X^{t+1} \dots X^{t+k} | R^{t+k}) \\ &\leq H(X^{t+1} \dots X^{t+k}) - H(X^{t+1} \dots X^{t+k} | S^{t+k}), \\ &H(X^{t+1} \dots X^{t+k} | S^{t+k}) \geq H(X^{t+1} \dots X^{t+k} | R^{t+k}). \end{aligned} \quad (\text{A72})$$

Thus, we see that using memory in one-to-one correspondence with the causal states S minimizes Eq. (A70). That is for any pattern,

$$\chi_R(k) \geq \chi_S(k) \text{ for all } R, k. \quad (\text{A73})$$

This proves the claim of the result. ■

We remark that this lower bound is not trivially saturated for all R . This is proven by the example in the article; where choosing R to not correspond to causal states in general resulted in increased dissipation.

Result III. A thermodynamically reversible cycle of pattern generation and extraction is impossible for any pattern where $C_\mu > I(\vec{X}; \vec{X})$, where C_μ is the pattern's *statistical complexity*.

Proof. This is a corollary of Results I and II. Consider using memory in one-to-one correspondence with the causal states.

In the limiting case of $k \rightarrow \infty$,

$$\lim_{k \rightarrow \infty} \beta W_{\text{diss}}^k \geq \chi_S = I(\bar{X}; S) - I(\bar{X}; \bar{X}). \quad (\text{A74})$$

Because causal states can be synchronized, $I(\bar{X}; S) = H(S) = C_\mu$, hence

$$\chi_S = C_\mu - I(\bar{X}; \bar{X}). \quad (\text{A75})$$

When $C_\mu > I(\bar{X}; \bar{X})$, $\chi_S > 0$, and it follows from Result I and II that for any k there will be some dissipation with this choice of memory. Result II then tells us that no other choice of memory can do better than this. It hence follows that when $C_\mu > I(\bar{X}; \bar{X})$, any attempt to generate a pattern will result in some work dissipation. ■

6. Perturbed coin example

Example (Further details of worked example). For a “last N ” generator that produces k steps of the *perturbed coin* pattern, using as its internal memory configurations in one-to-one correspondence with the last N outputs of the sequence, the amount of dissipation W_{diss}^k is bounded from below by

$$\beta W_{\text{diss}}^k \geq \begin{cases} kh(p) & \text{when } k \leq N \\ Nh(p) & \text{when } k > N, \end{cases} \quad (\text{A76})$$

where $h(p)$ is the *binary entropy* defined as

$$h(p) := -p \log p - (1-p) \log (1-p). \quad (\text{A77})$$

Proof. From Theorem 2, $\beta W_{\text{diss}}^k \geq \chi_R(k)$. We now derive the form of $\chi_R(k)$ for the last- N machine.

We begin with the form in Lemma 4(iii):

$$\chi_R(k) = H(X^{t+1} \dots X^{t+k} | R^t) - H(X^{t+k} \dots X^{t+1} | R^{t+k}). \quad (\text{A78})$$

Let us evaluate both terms.

First,

$$\begin{aligned} H(X^{t+1} \dots X^{t+k} | R^t) \\ = H(X^{t+1} | R^t) + H(X^{t+2} | R^t X^{t+1}) \\ + \dots + H(X^{t+k} | R^t X^{t+1} \dots X^{t+k-1}). \end{aligned} \quad (\text{A79})$$

Since R is prescient, and all the machines of the “last N ” type are unifilar (such that $H(R^{t+j} | R^t X^{t+1} \dots X^{t+j}) = 0$ for any integer j), it follows (using stationarity in the second step) that

$$\begin{aligned} H(X^{t+1} \dots X^{t+k} | R^t) \\ = H(X^{t+1} | R^t) + H(X^{t+2} | R^{t+1}) \\ + \dots + H(X^{t+k} | R^{t+k-1}), \\ = kH(X^{t+1} | R^t). \end{aligned} \quad (\text{A80})$$

Moreover, since R is prescient,

$$H(X^{t+1} \dots X^{t+k} | R^t) = H(X^{t+1} \dots X^{t+k} | S^t) = kh(p). \quad (\text{A81})$$

Now we evaluate the second term $H(X^{t+1} \dots X^{t+k} | R^{t+k})$. When $k \leq N$, since R encodes all the available information

about the last N outputs, $H(X^{t+1} \dots X^{t+k} | R^{t+k}) = 0$. Thus,

$$\chi_R(k) = kh(p) \quad \text{when } k \leq N. \quad (\text{A82})$$

Next, consider the case where $k > N$. We expand

$$\begin{aligned} H(X^{t+1} \dots X^{t+k} | R^{t+k}) \\ = H(X^{t+k-N} \dots X^{t+k} | R^{t+k}) \\ + H(X^{t+1} \dots X^{t+k-N-1} | X^{t+k-N} \dots X^{t+k} R^{t+k}) \\ = H(X^{t+1} \dots X^{t+k-N-1} | X^{t+k-N} \dots X^{t+k} R^{t+k}), \end{aligned} \quad (\text{A83})$$

where the first term was eliminated because R^{t+k} contains all the information about the preceding N outputs of the pattern.

For short-hand, we write $j = k - N + 1$. We can show that knowledge of R^{t+k} gives no further information beyond that in $X^{t+j} \dots X^{t+k}$ for the purpose of determining the values of $X^{t+1} \dots X^{t+j-1}$. This is seen by performing the following expansion:

$$\begin{aligned} H(X^{t+1} \dots X^{t+j-1} R^{t+k} | X^{t+j} \dots X^{t+k}) \\ = H(R^{t+k} | X^{t+j} \dots X^{t+k}) \\ + H(X^{t+1} \dots X^{t+j-1} | X^{t+j} \dots X^{t+k} R^{t+k}) \\ = H(X^{t+1} \dots X^{t+j-1} | X^{t+j} \dots X^{t+k}) \\ + H(R^{t+k} | X^{t+1} \dots X^{t+k}). \end{aligned} \quad (\text{A84})$$

Since $k > N$, both $H(R^{t+k} | X^{t+j} \dots X^{t+k}) = 0$ and $H(R^{t+k} | X^{t+1} \dots X^{t+k}) = 0$, since both strings have (at least) the N required outputs to fix the value of R^{t+k} . Thus, $H(X^{t+1} \dots X^{t+j-1} | X^{t+j} \dots X^{t+k} R^{t+k}) = H(X^{t+1} \dots X^{t+j-1} | X^{t+j} \dots X^{t+k})$.

We can now use the Markovian nature of the perturbed coin pattern to further simplify the above expression. For any Markovian process, by definition $H(X^{t+l} | X^{t+1} \dots X^{t+l-1}) = H(X^{t+l} | X^{t+l-1})$ for any $l \in \mathbb{Z}^+$. From this condition, it is possible to obtain the inverse statement for any $1 < n < m$:

$$\begin{aligned} H(X^{t+1} \dots X^{t+n-1} | X^{t+n} \dots X^{t+m}) \\ = H(X^{t+1} \dots X^{t+n-1} | X^{t+n}). \end{aligned} \quad (\text{A85})$$

We use this to further simplify

$$\begin{aligned} H(X^{t+1} \dots X^{t+j-1} | X^{t+j} \dots X^{t+k}) \\ = H(X^{t+j-1} | X^{t+j} \dots X^{t+k}) + \dots \\ + H(X^{t+1} | X^{t+2} \dots X^{t+k}) \\ = H(X^{t+j-1} | X^{t+j}) + \dots + H(X^{t+1} | X^{t+2}) \\ = (j-1)H(X^t | X^{t+1}), \end{aligned} \quad (\text{A86})$$

where the final step follows from the stationary nature of the process. Noting that $j-1 = k-N$, and directly calculating from the description of the perturbed coin process that $H(X^t | X^{t+1}) = h(p)$, we thus calculate that when $k > N$, the second term of Eq. (A78) is

$$H(X^{t+1} \dots X^{t+k} | R^{t+k}) = (k-N)h(p). \quad (\text{A87})$$

Taking the difference of Eq. (A81) and the above yields

$$\chi_R(k) = Nh(p) \quad \text{when } k > N. \quad (\text{A88})$$

From Theorem 2, $\beta W_{\text{diss}}^k \geq \chi_R(k)$. Thus, Eqs. (A82) and (A88) hence prove the claim. ■

FIG. 6. Logical reversibility in the last N machine. (Shown here for a last $N = 4$ machine updating with stride $k = 2$.) The last N machine's memory can be viewed as a window on the pattern of width N that moves forwards by k steps per cycle. The darker region indicates information stored in R^t , and the lighter region the information stored in R^{t+k} . Information corresponding to steps X^{t-1} and X^t (dashed region) exists in the machine's memory both before and after update. The part of the update pertaining to this information can be performed in a logical reversible manner.

That dissipation is limited by k for $k < N$ and N for $k \geq N$ may be understood in terms of the logical reversibility. The memory R of a last N machine may be visualized (Fig. 6) as a moving window of length N scanning over a pattern, advancing by k steps per update. When $k < N$, parts of this window overlap before and after the update: there is shared information contained within both R^t and within R^{t+k} pertaining to steps $X^{t-k+1} \dots X^t$ that can be updated using a logically reversible operation. On the other hand, when $k \geq N$, this overlap between R^t and R^{t+k} completely vanishes, and so the entire memory must be updated by logically irreversible processes.

One implementation of this would be to represent the last N machine's memory by a compound structure $R^t = X^{t-N+1} \dots X^{t-N+k} X^{t-N+k+1} \dots X^t$ with N registers, each storing information about one step of the pattern. When updating by $k < N$, one can first cyclically permute the contents of the registers (equivalently, relabel their indices) such that $\tilde{R}^t = X^{t-N+k+1} \dots X^t X^{t-N+1} \dots X^{t-N+k}$. This is an intrinsically reversible operation that requires no work investment. As the desired final memory state is $R^{t+k} = X^{t-N+k+1} \dots X^t X^{t+1} \dots X^{t+k}$, the first $N - k$ registers of the memory in $\tilde{R}^t$ already have their correct values, and work only needs to be invested to bring the final k registers of the memory up to date. On the other hand, when $k \geq N$, every register in the memory will require updating, and the cost will hence be bounded by the size of the memory N rather than the size of the update k .

7. Physical example: The trajectory formalism

For illustrative purposes, we present a physical model for the cycle of generation and extraction of the perturbed coin pattern. We shall employ a subset of the *trajectory formalism* (see, e.g., Refs. [21–28] among many), but stress that this framework is just one arbitrary choice from many thermodynamic models. The bounds derived in this article, being information-theoretic in origin, hold for any model of thermal interaction which defines heat and work in a way that is consistent with Landauer's principle.

We provide a few key details of this framework. Consider a system with a finite number of well-defined energy levels

(i.e., with a Hamiltonian $\mathcal{H} = \{E_1 \dots E_N\}$) and a (classical) state, reflecting the occupation probabilities $\{P_1 \dots P_N\}$ that the system is in a particular energy level. The energy of the system may change in one of two ways: (1) changes in the Hamiltonian, at fixed occupation probability; (2) changes in the occupation probability under a fixed Hamiltonian. We express this energy change differentially as

$$dU = \sum_{i=1}^N P_i dE_i + \sum_{i=1}^N E_i dP_i. \quad (\text{A89})$$

The first type of energy change can be induced by some choice of time-varying external parameter (i.e., force), and we shall require that the change in force is independent of the system's state to ensure that there is no unaccounted-for feedback.

We will place more restrictive conditions on the allowed energy exchanges of the second type. Namely, we only admit *thermalizing interactions*—such that all transformations on the state necessarily take it closer to the Gibbs state for a given Hamiltonian (that is, the state where $P_i = \frac{e^{-\beta E_i}}{\sum_j e^{-\beta E_j}}$ for all i). This ensures consistency with the second law of thermodynamics.⁴

For the purpose of the example in this article, we shall only consider *quasistatic* protocols, in which whenever the heat bath is coupled to the system, the system is allowed to reach perfect thermal equilibrium, and so remains in the Gibbs state associated with the Hamiltonian. This will trivially satisfy this requirement. There is some indication [26] that the cost of a quasistatic protocol bounds the actual finite-time cost reasonably tightly.

When all the above conditions are met, Eq. (A89) becomes somewhat like the *first law*, and the trajectory formalism allows us to associate the first term $\sum_i P_i dE_i$ with *work* and the second term $\sum_i E_i dP_i$ with *heat*.⁵ To find the total work cost of a protocol, one typically integrates over a series of infinitesimal contributions.

For the purpose of calculating costs of a pattern-manipulating protocol within this framework, we now prove the following lemma:

Lemma 12 (Work cost in a two-level system). In the trajectory formalism, for a two-level system that initially and finally is subject to a degenerate Hamiltonian, there is a quasistatic procedure that transforms it from state $(q, 1 - q)$ with binary entropy $h(q)$ to the state $(p, 1 - p)$ with binary entropy $h(p)$, at a work cost given by

$$W = k_B T [h(q) - h(p)], \quad (\text{A90})$$

⁴If additional behavioral constraints are imposed, the higher moments of the heat and work distributions' statistics can also be made to match physically expected behavior (e.g., imposing *detailed balance* ensures consistency with *fluctuation theorems* [46,47]). A general description of models that allow this is in the trajectory formalism is in Appendix A of Ref. [28].

⁵See, e.g., Refs. [48–50] for discussion of the additional considerations that must be taken into account before one can also make this equivalence in the quantum regime. In this article, we do not need to assert a quantum definition of work.

when the system has access to a thermal reservoir at temperature T .

Proof. We shall constructively provide a mechanism with this cost consisting of three stages. Let E_1 and E_2 be the values of the first and second energy levels respectively, such that initially $E_1 = E_2 = 0$.

First, in thermal isolation, we change the Hamiltonian such that $E_2 = E_q$ where E_q satisfies $q = \frac{1}{1+e^{-\beta E_q}}$. This produces a Hamiltonian where $(q, 1-q)$ is the associated Gibbs state. Since the second level was initially populated with probability $(1-q)$, changing the Hamiltonian has a work cost (respectively, gain if $q < \frac{1}{2}$) of $W_1 = (1-q)E_q$.

Next, we connect the system to a thermal bath. This has no effect on the occupation probabilities or Hamiltonian (and thus no associated work or heat cost). We slowly change the second energy level to a value $E_2 = E_p$ that satisfies $p = \frac{1}{1+e^{-\beta E_p}}$. When this is done quasistatically, and noting that $dE_1 = 0$ at all times, we find the total work exchange in this stage of the protocol is

$$\begin{aligned} W_2 &= \int_{E_q}^{E_p} P_2 dE_2 = \int_{E_q}^{E_p} \frac{e^{-\beta E_2}}{1+e^{-\beta E_2}} dE_2 \\ &= \frac{1}{\beta} (\ln p - \ln q), \end{aligned} \quad (\text{A91})$$

where we made the substitution $u = 1 + e^{-\beta E_2}$ to solve the integral.

For the third and final stage of the protocol, we disconnect the system [now in state $(p, 1-p)$] from the thermal bath, and lower the second energy level back down to $E_2 = 0$. This induces a work exchange of $W_3 = -(1-p)E_p$. Summing the contributions from the three parts of the protocol $W = W_1 + W_2 + W_3$, the total work exchange is

$$W = (1-q)E_q - (1-p)E_p + \frac{1}{\beta} (\ln p - \ln q). \quad (\text{A92})$$

However, since $E_q = \frac{1}{\beta} [\ln q - \ln(1-q)]$ (and likewise for p and E_p), we may rewrite the work cost

$$\begin{aligned} W &= \frac{1}{\beta} [p \ln p + \ln(1-p) - q \ln q - (1-q) \ln(1-q)], \\ &= k_B T [h(q) - h(p)], \end{aligned} \quad (\text{A93})$$

proving the claim. $\blacksquare$

It follows from conservation of energy (or can be shown directly using a calculation similar to the lemma above) that the heat transferred into the heat bath during this transaction must be equal to the work invested, since the average internal energy has not changed between the initial and final states.

Unlike the other lemmata in this article, here we have not assumed that Landauer's principle holds—the difference in entropies has appeared emergently from the protocol within the framework. However, we also have not explicitly proved Landauer's principle, since this would require a minimization over all possible protocols. Rather, what is shown is the existence of a mechanism within the trajectory framework that performs the above transformation in a manner that saturates Landauer's bound. If we consider the special case of $q = \frac{1}{2}$ and $p = 1$, then the above lemma yields the famous bit-reset

cost of $k_B T \ln 2$. Likewise by setting $q = 1$ and $p = \frac{1}{2}$, we have the Szilard engine output $-k_B T \ln 2$.

Pattern manipulation within the trajectory formalism. With this in mind, we may now provide a model for the generator of the perturbed coin pattern. Let us analyze the simplest possible example, where the pattern generator that writes one step at a time and has an internal memory configuration corresponding to the two causal states s_H and s_T (i.e., behaves according to Fig. 4). We shall make an extra “ancilla” bit of memory available to the generator, but must take care that it has explicitly been reset by the end of the procedure. As in Fig. 1, we make baths at temperatures T_G and T_E available to the generator and extractor, respectively, and assume that the energy investment required to make the worklike energy exchanges comes from some mutually available work reservoir (battery). This process may be viewed as an elaboration of the concepts mentioned in briefly in Figs. 2 and 3, but for the specialized case of a particular pattern, running with particular memory, at a particular stride.

Let the system on the tape be a two-level system, initially configured according to a predefined default distribution $X_{\text{dft}} = (q, 1-q)$. The first stage of the protocol will be to set the system from X_{dft} to an intermediate state $(p, 1-p)$, where the value of p is the “swap probability” determined by the perturbed coin process (i.e. has exactly the same meaning as in Fig. 4). We immediately see that the work cost of this (interacting with a bath at inverse temperature β_G) is what we have just calculated in Lemma 12: $\frac{1}{\beta_G} [H(X_{\text{dft}}) - h(p)]$. This cost corresponds to $W_{\text{gen}}^1(S)$ of Eq. (1).

Now, let us consider the system and the memory together. The memory has some state $S' = s_H$ or s_T . A controlled-not (CNOT) operation may be applied to the state of the system on the tape, such that if the memory was in state $S' = s_H$, nothing happens [X remains configured as $(p, 1-p)$], but if $S' = s_T$ the occupation probabilities are flipped putting the tape system into the state $(1-p, p)$. It is well-established that such purely reversible operations can be implemented at no net cost [9] (one can think of it much like a relabelling of energy levels). After this operation, the system on the tape will have been encoded with statistics X^{t+1} appropriate to the pattern.

Next, we must update the memory, to ensure that upcoming tape systems can also be set into the correct statistics (including appropriate correlations with X^{t+1}). To do this, we take the “ancilla” bit [initially in pure state $(1,0)$], and apply another reversible CNOT operation on it, controlled by the state of the patterned tape X^{t+1} . Next, we (reversibly) swap the state of the ancilla bit with our main memory bit. Since these two procedures are both reversible, they do not contribute to any work or heat costs.

Let us summarize the random variables describing the state of all three systems at this point in time: the main memory is configured according to S^{t+1} , the ancilla to S' and the tape to X^{t+1} . To finish the generation procedure in a manner that accounts for all potentially useful thermal resources, we must reset the ancilla from S' back to its initial pure state. In the simple Markovian example of the perturbed coin pattern, $S^{t+1} = X^{t+1}$, and all the useful information the tape contains regarding how to reset the ancilla is already encoded in the main memory [entropically:

$H(S' | S^{t+1} X^{t+1}) = H(S' | S^{t+1})$. This means we can at this point emit the patterned tape from the generating device, and consider the cost of resetting the ancilla from S' to $(1,0)$ only using knowledge in the main memory (S^{t+1}).

Clearly, except when $p = \frac{1}{2}$, there is some correlation between S' and S^{t+1} . Thus, the first stage of our memory reset is to decorrelate the two systems, by applying another reversible CNOT gate on the ancilla, controlled by the state of the main memory. This will set the ancilla into the state $(1 - p, p)$ independent of the current value of the memory (this works because of the symmetry of the perturbed coin pattern; there was a probability $(1 - p)$ that S^{t+1} and S' are the same, and a probability p that they were different).

Now, we may use the protocol in Lemma 12 to take the ancilla from state $(1 - p, p)$ back to the pure state $(1,0)$ at a work cost of $k_B T_G h(p)$ (i.e., using a thermal bath at the same temperature T_G as before). This particular work cost is an example of $W_{\text{diss}}^1(S)$ [see Eq. (3)], as has been discussed extensively throughout the article and appendices.

Thus, the memory is now perfectly correlated with the patterned tape emitted, and ready to accept the next system on the tape and continue generating the pattern. The generation stage is hence complete, requiring a total work cost of $\frac{1}{\beta_G} [H(X_{\text{dft}}) - h(p) + h(p)] = k_B T_G H(X_{\text{dft}})$.

Now consider the behavior of the extractor, which we assume also has access to a heat bath at (possibly) different temperature T_E . Initially, the system on the incoming tape is X^{t+1} is, and the internal memory is in state S' . The first stage is to reversibly swap the state of the tape system with the memory state. Since the perturbed coin process is Markovian and $X^{t+1} = S^{t+1}$, this alone will ensure that the memory is in the correct state to anticipate future the extraction of upcoming parts of the pattern. Now, we can apply a (reversible) CNOT on the system on the tape (currently in state S') controlled by the system in memory S^{t+1} , noting that as above there was a probability $1 - p$ that the systems are the same and of p that they are different. The system on the tape is now in the state $(1 - p, p)$, and is uncorrelated from the state in the extractor's memory. Thus far, no exchange of heat or work has been required.

For the final stage of extraction, however, we must again employ the protocol in Lemma 12, and in conjunction with heat bath at temperature T_E quasistatically reset the state on the tape from $(1 - p, p)$ to X_{dft} . This protocol required a work exchange of $k_B T_E [h(p) - H(X_{\text{dft}})]$, corresponding to $-W_{\text{out}}^1$ from Eq. (5), and concludes the extraction.

In summary, if we now consider a cycle of generation followed by extraction (as above) and set $T_G = T_E$, then the total of all work exchange terms is a net dissipation over the entire cycle is given by $k_B T h(p)$ —the value of $\mathcal{W}_{\text{diss}}^1(S)$ as predicted in Eq. (6). We have thus established constructive

protocol in the trajectory formalism for manipulating the pattern at a work cost that saturates the bounds given in the top row of table Table I.

Alternatively, if we choose $p < q < \frac{1}{2}$ [where p is the parameter from the perturbed coin process, and $X_{\text{dft}} = (q, 1 - q)$] and choose $T_E > T_G$, the above cycle can now function as a heat engine, as drawn in Fig. 1. The efficiency of this engine is directly calculated

$$\begin{aligned} \eta &= \frac{k_B T_E [H(X_{\text{dft}}) - h(p)] - k_B T_G H(X_{\text{dft}})}{k_B T_E [H(X_{\text{dft}}) - h(p)]} \\ &= 1 - \frac{T_G}{T_E} - \frac{T_G h(p)}{T_E [H(X_{\text{dft}}) - h(p)]}, \end{aligned} \quad (\text{A94})$$

and exactly matches the bound in Eq. (8) previously derived using information theory. For case of the perturbed coin, we thus conclude that only the simplest pattern (where $p = 0$ such that all states are the same, or $p = 1$ such that all states perfectly alternate) will achieve the Carnot efficiency.

An emergent proof (one that does not already accept the second law as true) that the protocols detailed in this example are optimal would well beyond the scope of this illustrative example,⁶ as it requires a tricky optimization over *all* possible operations that could be done on the joint tape-memory system (including allowing for an arbitrary amount of ancillary memory, etc.). On the other hand, this is where the power of the information-theoretic results derived in our article can be demonstrated: we may assert with confidence that the above protocol *is* optimal, since it saturates our bounds. If there were a protocol using prescient memory that is more work-efficient than this (in the trajectory formalism, or indeed any other framework), then Landauer's principle could not hold, and this would have drastic impact on our understanding of the second law—at least as to how it applies to the particular physical framework employed. Thus, if we have faith in the second law and how it has been applied within the physical mechanism, then we can be content to halt our search for a better mechanism here.

⁶If one takes a prescriptive “bottom-up” approach (e.g., Ref. [13]) rather than a proscriptive “top-down” approach (e.g., this article) to pattern thermodynamics, one can specify a particular process and consider what it does, rather than consider a particular task (here, a pattern manipulation defined only by its inputs and outputs) and work out how best to do it. The prescriptive approach typically avoids the question of optimality, as it tends to provide upper bounds (“there is a process at least this efficient”) rather than lower bounds (“no process can be more efficient than this”).

- [1] J. P. Crutchfield and K. Young, Inferring Statistical Complexity, *Phys. Rev. Lett.* **63**, 105 (1989).
- [2] C. R. Shalizi and J. P. Crutchfield, Computational mechanics: Pattern and prediction, structure and simplicity, *J. Stat. Phys.* **104**, 817 (2001).

- [3] J. P. Crutchfield, Between order and chaos, *Nat. Phys.* **8**, 17 (2011).
- [4] R. Haslinger, K. L. Klinkner, and C. R. Shalizi, The computational structure of spike trains, *Neural Comput.* **22**, 121 (2010).

- [5] J. B. Park, J. Won Lee, J.-S. Yang, H.-H. Jo, and H.-T. Moon, Complexity analysis of the stock market, *Physica A: Stat. Mech. Appl.* **379**, 179 (2007).
- [6] C. Lu and R. R. Brooks, P2P hierarchical botnet traffic detection using hidden Markov models, in *Proceedings of the 2012 Workshop on Learning from Authoritative Security Experiment Results—LASER '12* (ACM Press, New York, NY, 2012), pp. 41–46.
- [7] F. G. S. L. Brandão, M. Horodecki, J. Oppenheim, J. M. Renes, and R. W. Spekkens, Resource Theory of Quantum States Out of Thermal Equilibrium, *Phys. Rev. Lett.* **111**, 250404 (2013).
- [8] M. Horodecki and J. Oppenheim, Fundamental limitations for quantum and nanoscale thermodynamics, *Nat. Commun.* **4**, 2059 (2013).
- [9] C. H. Bennett, The thermodynamics of computation: a review, *Int. J. Theor. Phys.* **21**, 905 (1982).
- [10] R. Landauer, The physical nature of information, *Phys. Lett. A* **217**, 188 (1996).
- [11] H. Leff and A. F. Rex, *Maxwell's Demon 2: Entropy, Classical and Quantum Information, Computing* (Taylor & Francis, London, 2002).
- [12] J. M. R. Parrondo, J. M. Horowitz, and T. Sagawa, Thermodynamics of information, *Nat. Phys.* **11**, 131 (2015).
- [13] D. Mandal and C. Jarzynski, Work and information processing in a solvable model of Maxwell's demon, *Proc. Natl. Acad. Sci. USA* **109**, 11641 (2012).
- [14] S. Deffner and C. Jarzynski, Information Processing and the Second Law of Thermodynamics: An Inclusive, Hamiltonian Approach, *Phys. Rev. X* **3**, 041003 (2013).
- [15] A. C. Barato and U. Seifert, Stochastic thermodynamics with information reservoirs, *Phys. Rev. E* **90**, 042150 (2014).
- [16] D. H. Wolpert, Extending Landauer's bound from bit erasure to arbitrary computation, [arXiv:1508.05319](https://arxiv.org/abs/1508.05319).
- [17] A. B. Boyd, D. Mandal, and J. P. Crutchfield, Identifying functional thermodynamics in autonomous Maxwellian ratchets, *New J. Phys.* **18**, 023049 (2016).
- [18] A. B. Boyd, D. Mandal, and J. P. Crutchfield, Correlation-powered information engines and the thermodynamics of self-correction, *Phys. Rev. E* **95**, 012152 (2017).
- [19] A. B. Boyd, D. Mandal, and J. P. Crutchfield, Leveraging environmental correlations: The thermodynamics of requisite variety, [arXiv:1609.05353](https://arxiv.org/abs/1609.05353).
- [20] J. P. Crutchfield, C. J. Ellison, and J. R. Mahoney, Time's Barbed Arrow: Irreversibility, Crypticity, and Stored information, *Phys. Rev. Lett.* **103**, 094101 (2009).
- [21] R. Alicki, The quantum open system as a model of the heat engine, *J. Phys. A: Math. Gen.* **12**, L103 (1979).
- [22] R. Alicki, M. Horodecki, P. Horodecki, and R. Horodecki, Thermodynamics of quantum informational systems—Hamiltonian description, *Open Syst. Inf. Dyn.* **11**, 205 (2004).
- [23] T. D. Kieu, Quantum heat engines, the second law and Maxwell's daemon, *Eur. Phys. J. D* **39**, 115 (2006).
- [24] H. T. Quan and H. Dong, Quantum Crooks fluctuation theorem and quantum Jarzynski equality in the presence of a reservoir, [arXiv:0812.4955](https://arxiv.org/abs/0812.4955).
- [25] S. W. Kim, T. Sagawa, S. De Liberato, and M. Ueda, Quantum Szilard Engine, *Phys. Rev. Lett.* **106**, 070401 (2011).
- [26] C. Browne, A. J. P. Garner, O. C. O. Dahlsten, and V. Vedral, Guaranteed Energy-Efficient Bit Reset in Finite Time, *Phys. Rev. Lett.* **113**, 100603 (2014).
- [27] J. Åberg, Truly work-like work extraction via a single-shot analysis, *Nat. Commun.* **4**, 1925 (2013).
- [28] N. Yunger Halpern, A. J. P. Garner, O. C. O. Dahlsten, and V. Vedral, Introducing one-shot work into fluctuation theorems, *New J. Phys.* **17**, 095003 (2015).
- [29] D. Egloff, O. C. O. Dahlsten, R. Renner, and V. Vedral, A measure of majorization emerging from single-shot statistical mechanics, *New J. Phys.* **17**, 073001 (2015).
- [30] W. Thomson, On the dynamical theory of heat, with numerical results deduced from Mr Joule's equivalent of a thermal unit, and M. Regnault's observations on steam, *Trans. R. Soc. Edinburgh* **20**, 261 (1851).
- [31] J. P. Crutchfield, C. J. Ellison, R. G. James, and J. R. Mahoney, Synchronization and control in intrinsic and designed computation: An information-theoretic analysis of competing models of stochastic computation, *Chaos: Interdisc. J. Nonlin. Sci.* **20**, 037105 (2010).
- [32] C. R. Shalizi, Causal architecture, complexity and self-organization in time series and cellular automata, Ph.D. thesis, University of Wisconsin at Madison, 2001.
- [33] R. G. James, J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield, Many roads to synchrony: Natural time scales and their algorithms, *Phys. Rev. E* **89**, 042135 (2014).
- [34] J. R. Mahoney, C. J. Ellison, R. G. James, and J. P. Crutchfield, How hidden are hidden processes? A primer on crypticity and entropy convergence, *Chaos* **21**, 037112 (2011).
- [35] K. Wiesner, M. Gu, E. Rieper, and V. Vedral, Information-theoretic lower bound on energy cost of stochastic computation, *Proc. R. Soc. London A* **468**, 4058 (2012).
- [36] S. Still, D. A. Sivak, A. J. Bell, and G. E. Crooks, Thermodynamics of Prediction, *Phys. Rev. Lett.* **109**, 120604 (2012).
- [37] J. R. Mahoney, C. J. Ellison, and J. P. Crutchfield, Information accessibility and cryptic processes, *J. Phys. A: Math. Theor.* **42**, 362002 (2009).
- [38] S. E. Marzen and J. P. Crutchfield, Informational and causal architecture of continuous-time renewal and hidden semi-Markov processes, [arXiv:1611.01099](https://arxiv.org/abs/1611.01099).
- [39] M. Gu, K. Wiesner, E. Rieper, and V. Vedral, Quantum mechanics can reduce the complexity of classical models, *Nat. Commun.* **3**, 762 (2012).
- [40] W. Y. Suen, J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, The classical-quantum divergence of complexity in the Ising spin chain, [arXiv:1511.05738](https://arxiv.org/abs/1511.05738).
- [41] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, Occam's quantum Strop: Synchronizing and compressing classical cryptic processes via a quantum channel, *Sci. Rep.* **6**, 20495 (2016).
- [42] J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, Using quantum theory to simplify input-output processes, *npj Quantum Inf.* **3**, 6 (2017).
- [43] C. Aghamohammadi, J. R. Mahoney, and J. P. Crutchfield, The ambiguity of simplicity in quantum and classical simulation, *Phys. Lett. A* **381**, 1223 (2017).
- [44] P. M. Riechers, J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, Minimized state complexity of quantum-encoded cryptic processes, *Phys. Rev. A* **93**, 052317 (2016).

- [45] J. Pearl, *Causality*, 2nd ed. (Cambridge University Press, Cambridge, 2000).
- [46] C. Jarzynski, Nonequilibrium Equality for Free Energy Differences, *Phys. Rev. Lett.* **78**, 2690 (1997).
- [47] G. E. Crooks, Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences, *Phys. Rev. E* **60**, 2721 (1999).
- [48] H. Weimer, M. J. Henrich, F. Rempp, H. Schröder, and G. Mahler, Local effective dynamics of quantum systems: A generalized approach to work and heat, *Europhys. Lett.* **83**, 30008 (2008).
- [49] S. Vinjanampathy and J. Anders, Quantum thermodynamics, *Contemp. Phys.* **57**, 545 (2016).
- [50] J. Goold, M. Huber, A. Riera, L. del Rio, and P. Skrzypczyk, The role of quantum information in thermodynamicsa topical review, *J. Phys. A: Math. Theor.* **49**, 143001 (2016).