

UC Davis

UC Davis Previously Published Works

Title

Leveraging Environmental Correlations: The Thermodynamics of Requisite Variety

Permalink

<https://escholarship.org/uc/item/0fc0443j>

Journal

Journal of Statistical Physics, 167(6)

ISSN

0022-4715

Authors

Boyd, Alexander B
Mandal, Dibyendu
Crutchfield, James P

Publication Date

2017-06-01

DOI

10.1007/s10955-017-1776-0

Peer reviewed

Leveraging Environmental Correlations: The Thermodynamics of Requisite Variety

Alexander B. Boyd,^{1,*} Dibyendu Mandal,^{2,†} and James P. Crutchfield^{1,‡}

¹*Complexity Sciences Center and Physics Department,*

University of California at Davis, One Shields Avenue, Davis, CA 95616

²*Department of Physics, University of California, Berkeley, CA 94720, U.S.A.*

(Dated: March 11, 2017)

Key to biological success, the *requisite variety* that confronts an adaptive organism is the set of detectable, accessible, and controllable states in its environment. We analyze its role in the thermodynamic functioning of information ratchets—a form of autonomous Maxwellian Demon capable of exploiting fluctuations in an external information reservoir to harvest useful work from a thermal bath. This establishes a quantitative paradigm for understanding how adaptive agents leverage structured thermal environments for their own thermodynamic benefit. General ratchets behave as memoryful communication channels, interacting with their environment sequentially and storing results to an output. The bulk of thermal ratchets analyzed to date, however, assume memoryless environments that generate input signals without temporal correlations. Employing computational mechanics and a new information-processing Second Law of Thermodynamics (IPSL) we remove these restrictions, analyzing general finite-state ratchets interacting with structured environments that generate correlated input signals. On the one hand, we demonstrate that a ratchet need not have memory to exploit an uncorrelated environment. On the other, and more appropriate to biological adaptation, we show that a ratchet must have memory to most effectively leverage structure and correlation in its environment. The lesson is that to optimally harvest work a ratchet’s memory must reflect the input generator’s memory. Finally, we investigate achieving the IPSL bounds on the amount of work a ratchet can extract from its environment, discovering that finite-state, optimal ratchets are unable to reach these bounds. In contrast, we show that infinite-state ratchets can go well beyond these bounds by utilizing their own infinite “negentropy”. We conclude with an outline of the collective thermodynamics of information-ratchet swarms.

PACS numbers: 05.70.Ln 89.70.-a 05.20.-y 05.45.-a

Keywords: Maxwell’s Demon, cybernetics, detailed balance, entropy rate, Second Law of Thermodynamics, transducer, adaptation

I. INTRODUCTION

The mid-twentieth century witnessed an efflorescence in information and control and, in particular, the roles they play in biological adaptation [1]. Norbert Wiener’s linear prediction theory [2, 3] and Claude Shannon’s mathematical theory of communication [4–7] stood out as the technical underpinnings. It was Wiener, though, who advocated most directly for a broad development of a new calculus of control and adaptation, coining the term “cybernetics” [8, 9]. The overall vision and new methods of information theory and linear stochastic processes stimulated a tremendous enthusiasm and creativity during this period.

It must be said that, despite substantial efforts throughout the 1950s and 1960s to develop “general systems” theories and the like [10, 11], at best, only modest successes transpired which addressed Wiener’s challenges for cybernetics [12]. Historians of science claimed, in fact,

that progress was inhibited by the political tensions between the West and East during the Cold War [13]. More practically, one cause was the immodest complicatedness of the systems targeted—weather control, the brain, and social design. In short, there simply were not the powerful computational and mathematical tools required to understand such large-scale, complex systems. This all said, we must not forget that the intellectual fallouts from this period—the development of communication, coding, computation, and control theories—substantially changed the landscape of the engineering disciplines and irrevocably modified modern society.

Now, at the beginning of the 21st century, it seems time to revisit the broad and ambitious goals these early pioneers laid out. For, indeed, the challenges they introduced are still with us and are evidently starting to reveal dire consequences of our failure to understand the dynamics and emergent properties of large-scale complex systems, both natural and man-made. Optimistically, very recent developments in nonlinear dynamics [14] and nonequilibrium thermodynamics [15] give hope to finally achieving several of their goals, including reframing them in ways that will facilitate physical implementation. Here, we elucidate cybernetics’ Law of Requisite Variety

* aboyd@ucdavis.edu

† dibyendu.mandal@berkeley.edu

‡ chaos@ucdavis.edu

in light of these recent advances.

W. Ross Ashby was one of cybernetics’s best expositors [16], having an impact that rivaled Wiener’s advocacy. Principle to Ashby’s approach was his concept of *requisite variety*. The requisite variety that confronts an adaptive system is the set of accessible, detectable, and controllable states in its environment. In its most elementary form, Ashby re-interpreted Shannon’s notion of information-as-surprise, retooling it for broader application to biological and cognitive systems [11]. In this, though, he was anticipated by 30 years by Leo Szilard’s successful purging of Maxwell Demon [17, 18]: “... a simple inanimate device can achieve the same essential result as would be achieved by the intervention of intelligent beings. We have examined the ‘biological phenomena’ of a nonliving device and have seen that it generates exactly that quantity of entropy which is required by thermodynamics”. In laying out the thermodynamic costs of measurement, and so showing any demon is consistent with the Second Law of Thermodynamics, Szilard not only anticipates by two decades Shannon’s quantitative measure of information but also Wiener’s conception of cybernetics in which stored information plays a *functional role*.

The conceptual innovation in Szilard’s analysis, still largely underappreciated, is his identifying two distinct kinds of information. On the one hand, there is surprise; Shannon’s notion that later on lead to an algorithmic foundation for randomness and probability [19–22]. Its parallel in physics is a system’s thermodynamic entropy [23]. The Demon monitors statistical fluctuations in its heat-bath environment. On the other hand, there is information stored as historical contingency and memory. It is this latter kind that explains the thermodynamic functionality of Maxwell’s Demon, as it uses stored information about the thermal fluctuations to convert them to useful work [24]. This recognition handily resolves Maxwell’s Second Law paradox. This information dichotomy was recently laid bare by mapping Szilard’s single-molecule engine to chaotic dynamical system; a mapping so simple that all questions can be analytically addressed [25]. The role of both informative measurement and its use, when stored, for control illustrates the complementary role and functional consequences of both kinds of information in an adaptive system.

In this way, the now-familiar physical setting of Maxwell’s paradox highlights how the distinction between *information-as-surprise* and *stored actionable-information* motivated Ashby’s emphasizing requisite variety in adaptation. Detecting environmental fluctuations and acting on their structure (such as temporal correlations) are critical to the Demon’s functioning. Appealing to new results in nonlinear dynamics and nonequilibrium thermodynamics, the distinction simi-

larly motivates our re-addressing this central concern in cybernetics, so basic to the operation of adaptive systems, but in a fully thermodynamic setting: What requisite variety (range of historical contexts) must an adaptive agent recognize in its environment to realize thermodynamic benefits?

In the following, we first give an overview of our contributions (Sec. II). We mention how Ashby’s law of requisite variety is faithfully reflected in the behavior of *information engines*—autonomous versions of Maxwell’s Demon. This close connection follows from the bounds set by the Second Law of Thermodynamics for information processing [26–28]. Important for engineered and biological implementations, we note that these bounds, and so those specified by Landauer’s Principle [29, 30], are not generally achievable. The subsequent sections form the technical components of the development, key to which is representing an information reservoir in terms of the outcomes of a hidden Markov process.

Section III considers (i) the meaning of memory for the input processes of information engines and for the engines themselves, (ii) their energetics, and (iii) the role of memory in information thermodynamics more generally [31, 32]. It is the thermodynamics of memory that establishes the correspondence between Ashby’s law and the behavior of information engines. Section IV addresses the limits on information engines achieving the informational Second Law bounds. We see that the bounds are not saturated even by optimal, finite-state engines. We also mention the curious case of infinite-memory information engines that can achieve and then go beyond these bounds, essentially by leveraging their internal infinite “negentropy” to generate work [33]. These results bear directly on the description of Maxwell’s original Demon and, more contemporarily, stochastic universal Turing machines built out of information engines. Finally, we conclude with a summary of our results and their implications for biological physics and engineering.

II. SYNOPSIS OF MAIN RESULTS

Szilard’s Engine and related Maxwellian Demons are instances of thermal agents processing environmental information in order to convert thermal energy into work. Turning disordered thermal energy into work (ordered energy) was long thought to violate the Second Law of Thermodynamics [34]. However, the past century resolved the apparent violation by recognizing that information processing has unavoidable energy costs. Rolf Landauer was one of the first to set bounds on information processing—specifically, erasing a bit—such that the work production over a thermodynamic cycle cannot be

positive, satisfying the Second Law of thermodynamics [29, 30].

However, if the Demon accesses an information reservoir in its environment, it can use the reservoir’s statistics as a resource to convert thermal energy into work. This view of a Demon taking advantage of a structured environment connects back to cybernetics. Just as Ashby asked how a controller’s variety should match that of its inputs, we ask how the Demon’s internal structure should match the structure of an input process, which characterizes the information reservoir, in order to generate work. In contrast to cybernetics, though, we consider the variety inherent in “information ratchets” viewed as thermodynamic systems and, by implication, the variety they can detect and then leverage in their environments.

An information ratchet is an explicit construction of an autonomous Maxwellian Demon that uses an input symbol sequence to turn thermal energy into work energy [26, 35]. The ratchet steadily transduces the input symbols into an output sequence, processing the input information into an output while effecting thermodynamic transformations—implementing a physically embedded, real-time computation. This is accomplished by driving the ratchet along the input symbol sequence unidirectionally, so that the ratchet (with states in set \mathcal{X}) interacts once with each symbol (with values in alphabet \mathcal{Y}). During the interaction, the ratchet and current symbol make a thermally activated joint transition from $x \otimes y \in \mathcal{X} \otimes \mathcal{Y}$ to $x' \otimes y'$ with probability [27]:

$$M_{x \otimes y \rightarrow x' \otimes y'} = \Pr(X_{N+1} = x', Y'_N = y' | X_N = x, Y_N = y) .$$

Because thermal dynamics are reversible, the Markov chain described by the transition matrix M must have detailed balance. We define a physical ratchet as one which can be achieved through such reversible dynamics, and thus have calculable energetics. (The requirement of detailed balance comes from the fact that the ratchets correspond to thermal physical systems performing computation and, thereby, must satisfy this condition in the absence of external, nonconservative forces, which we assume to be the case.) The transition matrix M determines the energetics as well as the ratchet’s information processing capacity.

Recent work introduces a general computational mechanics [14, 36] framework for analyzing thermodynamic devices that transduce an input process into an output process [27, 36]. Figure 1 depicts the relative roles of the input process specified by a finite-state hidden Markov model (HMM), the ratchet as transducer operating on the input process, and the resulting output process, also given by an HMM.




FIG. 1. Computational mechanics view of an information ratchet: The input signal (environment) is described by a hidden Markov model (HMM) that generates the input symbol sequence. The ratchet itself acts as a transducer, using its internal states or memory to map input symbols to output symbols. The resulting output sequence is described by an HMM that results from composing the transducer with the input HMM. The current internal state of the input HMM, transducer, and output HMM are each highlighted by a dashed red circle. These are the states achieved after the last output symbol (highlighted by a red box) of each machine. We see that the internal state of the output HMM is the direct product of the internal state of the transducer and the input HMM.

The tools of computational mechanics were developed to quantitatively analyze how a ratchet’s structure should match that of its input for maximum efficacy, since they use a consistent notion of structure for general processes and transformations. In particular, using them we recently established a general information processing Second Law (IPSL) for thermodynamically embedded information processing by finite ratchets that bounds the asymptotic work per cycle $\langle W \rangle$ in terms of the difference in entropy rates of the input and output processes, h_μ and h'_μ , respectively [27] (see App. B for a proof):

$$\langle W \rangle \leq k_B T \ln 2 (h'_\mu - h_\mu) . \quad (1)$$

(Definitions are given shortly in Sec. III.) Employing entropy rates—the Shannon entropy rate of the symbol sequence or, equivalently here, the Kolmogorov-Sinai entropy of its generative dynamical system—the bound accounts for all temporal correlations in the input and output processes as well as the single-symbol biases. While this bound appears similar to that $\langle W \rangle \leq \langle I \rangle - \Delta F$ [37] on work production in a system with feedback control, $\langle I \rangle$ quantifies correlations between the controller and en-

vironment rather than temporal correlations induced in the environment.

Two uses of Eq. (1)’s IPSL suggest themselves. First, it sets an informational upper bound on the maximum average work production $\langle W \rangle$ per thermodynamic cycle. Here, W is the flow of work from the ratchet to an external driver. Second, and complementarily, it places an energetic lower bound on the minimal work $\langle W_d \rangle$ required to drive a given amount (Δh_μ) of computation forward. Here, $W_d = -W$ is the flow of work from the driver into the ratchet. In this second use, the IPSL is a substantial extension of Landauer’s Principle. The latter says that erasing a bit of information requires a minimum energy expenditure of $k_B T \ln 2$ while the IPSL applies to any kind of computational processing that transforms an input process to an output process, not simply erasure. The first use appears, in this light, as a rough converse to Landauer’s limit: There is a potential thermodynamic benefit of “destroying variety” in the form of work [29, 30].

Practically, computational mechanics gives a means to partition the ratchet and input process into different cases: memoryful and memoryless. Whether or not the input process or ratchet have memory substantially changes the bound on work production. And so, we can examine how environmental and demon varieties interact. For example, in the case in which temporal correlations (varieties) vanish, the difference between the input’s single-symbol entropy H_1 and the output’s H'_1 gives an analogous bound [38]:

$$\langle W \rangle \leq k_B T \ln 2 (H'_1 - H_1) , \quad (2)$$

Using the single-symbol approximation H_1 of the true entropy rate h_μ can be quite convenient since H_1 is much easier to calculate than h_μ , as the latter requires asymptotic (long-range) sequence statistics. (Again, definitions are given shortly in Sec. III.) Likely, this is why the H_1 -bound has appeared frequently to describe ratchet information processing [26, 32, 39–41]. Also, Eq. (2) is a rather direct generalization of the Landauer limit, since the input entropy $H_1 = 1$ bit and the output $H'_1 = 0$ bits saturate the bound on the work required to drive erasing a binary symbol. However, a key difference is that Eq. (1)’s entropy rates are dynamical invariants; unchanged by smooth transformations [42, 43]. The single-symbol Shannon entropies are not dynamical invariants. In addition, the single-symbol bound does not properly account for the temporal correlations in the input process or those created by the ratchet in the output process and so leads to several kinds of error in thermodynamic analysis. Let us explore these.

First, the average total temporal correlation in a pro-

cess can be quantified by the difference between the single-symbol entropy and the entropy rate, known as a process’ length-1 *redundancy* [44]:

$$H_1 - h_\mu \geq 0 . \quad (3)$$

This is the extent to which single-symbol entropy-rate estimates (H_1) exceed the actual per-symbol uncertainty (h_μ); and it is always nonnegative. This measure describes a type of structure distinct from statistical auto-correlations. Unless stated otherwise, going forward, the informational temporal correlations quantified in Eq. (3) are what we mean by *correlations*.

How inputs or ratchets create or destroy these correlations determines the relative strength and validity of the Eq. (1) and Eq. (2) work bounds. These bounds, in turn, suggest that memoryless ratchets are best for leveraging memoryless inputs and memoryful ratchets are best for leveraging memoryful inputs and generating work. However, it is not clear if and when the bounds are achievable. So, more effort is required to establish this thermodynamic version of Ashby’s Law of Requisite Variety.

To address achievability, we turn to a general energetic framework for calculating ratchet work production [28]. There it was shown that memoryful ratchets can leverage temporal correlations which memoryless ratchets cannot. In short, memoryful ratchets are indeed best for leveraging memoryful inputs. This gives an explicit violation of Eq. (2). However, for memoryless ratchets both Eqs. (2) and (1) are valid bounds [38]. We show, with proof given in App. A, that memoryless ratchets are the best among all finite ratchets at leveraging statistical biases in memoryless inputs to produce work. Notably, these ratchets do not achieve the derived upper bounds on work production, demonstrating fundamental inefficiencies in the information-to-work conversion in this class of an autonomous Maxwellian Demon.

To approach the bounds described by Eqs. (1) and (2) it is necessary to go beyond the information processing paradigm of a single finite-memory ratchet that interacts with a single symbol at a time. For instance, consider a “swarm” of finely tuned ratchets that work in a sequence, the output of one acting as the input of the next, and each ratchet being optimized with respect to its own input. This stepwise, sequential processing of the information reservoir is more efficient than the single-ratchet paradigm and is able to approach the upper bounds on information processing as the number of ratchets in the army grows. (This is reminiscent of the higher efficiency of quasistatic thermodynamic processes compared to finite-time, irreversible processes.) We reserve the detailed analysis of this phenomenon for a later work since the framework for collective thermodynamics

is less developed than the single-ratchet setting we focus on here.

While the IPSL and related bounds on work are suggestive of how the structure of the input matches the output, the fact that they are unachievable for single-information ratchets means we must reach further to solidify the relationship between input statistics and ratchet thermodynamics. Exact calculations here for the work production verify the intuition that the memory of an optimal ratchet must match the memory of the input. This leads to a variation on Ashby’s Law of Requisite Variety: “memory leverages memory”.

In this way, the transducer framework for information ratchets gives insight into how adaptive agents leverage structure. Its importance extends far beyond, however, to general computation. On the one hand, transducers describe mappings from input sequences to distributions over output sequences [36, 45] and do so in real time. Turing machines, on the other, map individual input sequences to individual output sequences with no particular reference to physical time. In this sense, Turing machines are a subclass of transducers, emphasizing that transducers are a general model for physical computation and information processing. However, to do universal computation, as properly configured Turing machines can, requires infinitely many states [45]. And, this suggests examining the thermodynamics of infinite-memory ratchets.

It turns out that infinite ratchets with states having finite energy differences are pathological in that they violate both the IPSL and its single-symbol sister bounds on work production—Eqs. (1) and (2), respectively. The proof of Eq. (1) assumes a stationary distribution over the ratchet state and input symbol. This need not exist for infinite ratchets [27]. In this case structure in the ratchet’s memory, rather than structure in the information reservoir, can be used as an additional thermodynamic resource to produce work. And, this means that a framework for general computation requires more detailed analysis to set bounds on work production that account for the ratchet’s memory. While we leave this for upcoming work, it does call into question any discussion of the thermodynamics of universal computation.

Following are the major contributions:

1. To address the role of memory in ratchets, we introduce thermodynamically predictive definitions of memory for *both* the input string and the ratchet which performs a computation on the input.
2. The validity of IPSL bounds changes depending on whether or not the input or ratchet are separately or together memoryful or memoryless.
3. The memory dependence of IPSLs implies that,

if the IPSL derived in our previous work [27] is achievable, then we arrive at a thermodynamic Law of Requisite Variety [11].

4. However, our exact analysis of memoryless ratchets driven by memoryless inputs tells us that to achieve IPSL bounds we must go beyond the class of individual autonomous ratchets that operate on a single bit at a time. The Law of Requisite Variety may not hold for such generalized ratchets.
5. Fortunately, we also complete a suite of results about information ratchets that show memoryful ratchets are best for leveraging memoryful inputs and memoryless ratchets are best for leveraging memoryless inputs. This confirms the Law of Requisite Variety from a dynamical perspective, independent of IPSLs.
6. While the results hold for finite ratchets, for the potentially unphysical case of infinite ratchets, both IPSLs and the Law of Requisite Variety fail: an infinite ratchet can violate IPSL bounds, extracting more energy than a memoryless input allows.

With this overview laid out, with the goals and strategy stated, we now are ready to delve into memory’s role in information-engine thermodynamics and the achievability of the IPSL and its related bounds.

III. MEMORY

To explore how a ratchet’s structure “matches” (or not) that of an environmental signal requires quantifying what is meant by structure. In terms of their structure, both ratchets and environmental inputs can be either memoryless or memoryful and this distinction delineates a ratchet’s thermodynamic functioning via the IPSL. This section introduces what we mean by the distinction, describes how it affects identifying temporal correlations, and shows how it determines bounds on work production and functionality. The results, though, can be concisely summarized. Figure 2 presents a tableau of memoryless and memoryful ratchets and inputs in terms of example HMM state-transition diagrams. Figure 3 then summarizes IPSL bounds for the possible cases.

A. Process memory

The amount of memory in the input or output processes is determined by the number of states in the minimal representative dynamics that generates the associated sequence probability distributions. We make this

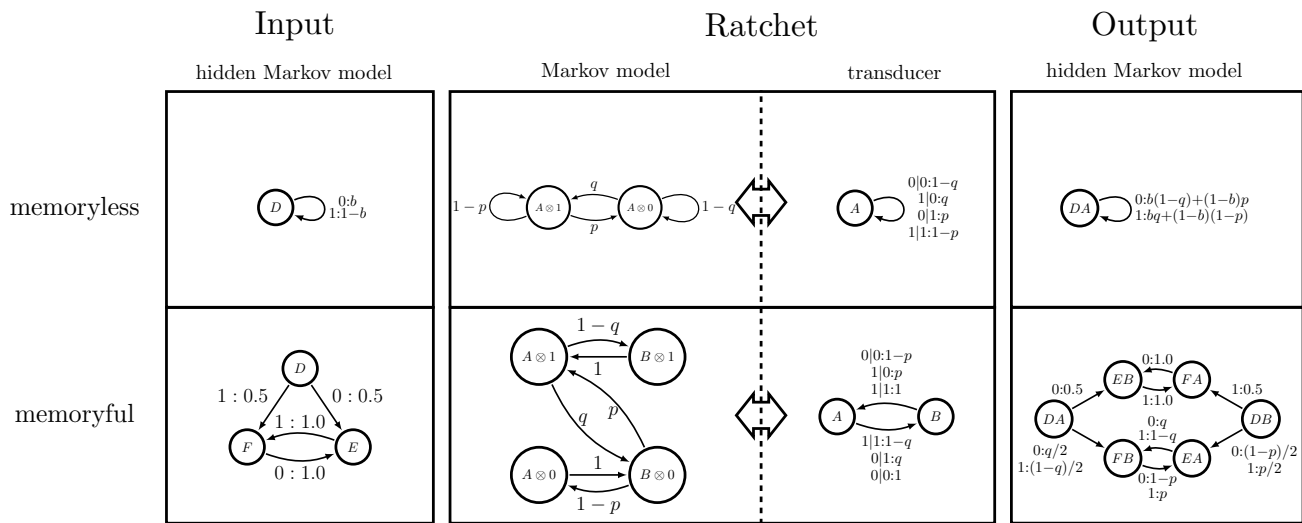


FIG. 2. Ratchets and input and output signals can be either memoryful or memoryless. For the input or output signal to be memoryless, the generating (minimal) HMM must have more than one internal state. The action of a ratchet can be represented in two different ways: either by a detailed Markov model involving the joint state space of the ratchet and an input symbol or by a symbol-labeled Markov dynamic on the ratchet’s state space. We call the latter the *transducer representation* [36]. Similar to the input and output signals, if the (minimal) transducer has more than one internal state, then the ratchet is memoryful.

definition clear in the following section and show that it is relevant to predicting thermodynamic bounds.

While there are many ways to generate a process, HMMs are a particularly useful representation of generating mechanisms. For example, they describe a broader class of processes than finite-order Markov models, since they can generate infinite Markov-order processes using only a finite number of hidden states [44].

Here, we use the Mealy representation of HMMs [46–49], which consists of a set \mathcal{S} of internal states and an alphabet \mathcal{A} of symbols that are emitted. As with a Markov chain, transitions between hidden states in \mathcal{S} are made according to conditional probabilities. However, the generated symbols in \mathcal{Y} are emitted during transitions between hidden states, rather than when entering states [50]. The Mealy HMM dynamic is specified by a set of symbol-labeled transition matrices:

$$T_{s_N \rightarrow s_{N+1}}^{(y_N)} = \Pr(Y_N = y_N, S_{N+1} = s_{N+1} | S_N = s_N),$$

which give the joint probability of emitting y_N and transitioning to hidden state s_{N+1} given that the current hidden state is s_N . For the special class of *unifilar* HMMs the current hidden state s and emitted symbol y uniquely determine the next hidden state $s'(s, y)$. Helpfully, for unifilar HMMs the generated process’ entropy rate h_μ is exactly given by the state-averaged uncertainty in the

emitted symbols given the current state [44]:

$$\begin{aligned} h_\mu &= \lim_{N \rightarrow \infty} \mathbb{H}[Y_N | Y_{0:N}] \\ &= \lim_{N \rightarrow \infty} \mathbb{H}[Y_N | S_N] \\ &= \sum_{s \in \mathcal{S}} \pi_s \lim_{N \rightarrow \infty} \mathbb{H}[Y_N | S_N = s] \\ &= - \sum_{s \in \mathcal{S}} \pi_s \sum_{y \in \mathcal{Y}} T_{s \rightarrow s'(s,y)}^{(y)} \log_2 T_{s \rightarrow s'(s,y)}^{(y)}, \end{aligned}$$

where π_s is the steady-state distribution over the hidden states. A process’ ϵ -machine is its minimal unifilar HMM generator, where minimality is determined by having the smallest internal-state Shannon entropy [14]:

$$\begin{aligned} \lim_{N \rightarrow \infty} \mathbb{H}[S_N] &= - \lim_{N \rightarrow \infty} \sum_{s \in \mathcal{S}} \Pr(S_N = s) \log_2 \Pr(S_N = s) \\ &= - \sum_{s \in \mathcal{S}} \pi_s \log_2 \pi_s \\ &\equiv C_\mu. \end{aligned}$$

where in the last line we defined the process’ *statistical complexity* C_μ . Since h_μ gives an exact expression for process entropy rate and C_μ a unique definition of process memory [51], throughout we represent processes by their ϵ -machines. An ϵ -machine’s internal states are called *causal states*.

Broadly, the memory of an ϵ -machine refers to its hidden states. As shown in Fig. 2, memoryless input processes have ϵ -machines with a single state: $|\mathcal{S}| = 1$. The

sequence distributions for such processes are given by a product of single-symbol marginal distributions. For a stationary process, the single-symbol marginal entropy H_1 is the same for every symbol:

$$H_1 \equiv H[Y_N] \text{ for all } N \in \mathbb{N}. \quad (4)$$

For memoryless processes, the entropy rate is the same as the single-symbol entropy:

$$\begin{aligned} h_\mu &= \lim_{N \rightarrow \infty} H[Y_N | Y_{0:N}] \\ &= \lim_{N \rightarrow \infty} H[Y_N] \\ &= H_1 . \end{aligned}$$

This means that their difference vanishes:

$$H_1 - h_\mu = 0 . \quad (5)$$

Thus, there are no temporal correlations in the symbol string, since $H_1 - h_\mu$ quantifies the informational correlation of individual input symbols with past inputs:

$$\begin{aligned} H_1 - h_\mu &= \lim_{N \rightarrow \infty} (H[Y_N] - H[Y_N | Y_{0:N}]) \\ &= \lim_{N \rightarrow \infty} I[Y_N : Y_{0:N}] , \end{aligned} \quad (6)$$

where $I[W : Z]$ is the mutual information of random variables W and Z [52].

For memoryful input processes, as shown in Fig. 2, there are multiple causal states for the ϵ -machine: $|\mathcal{S}| > 1$. In other words, sequence probabilities cannot be broken into a product of marginals. And so, in general, we have:

$$H_1 > h_\mu .$$

Thus, there are temporal correlations in the input process:

$$H_1 - h_\mu > 0 . \quad (7)$$

This means that individual symbols of the input sequence share information with past inputs. In the maximally correlated case, every symbol is exactly predictable from its past. As a result the entropy rate vanishes and the temporal correlation measure in Eq. (6) is equal to the single-symbol entropy.

To summarize, memoryless input signals have a single causal state and, thus, do not exhibit temporal correlations, since they have no way to store information from the past. Meanwhile, memoryful inputs have multiple hidden states that are used to transmit information from the past to the present and so express temporal correla-

tions.

B. Ratchet memory

From the perspective of information processing, the ratchet is a transducer that interacts with each symbol in the input sequence in turn, converting it into a output symbol stored in the output sequence [27, 36]. The ratchet is a form of communication channel [52]. One that is determined by a detailed-balanced Markov dynamic:

$$\begin{aligned} M_{x_N \otimes y_N \rightarrow x_{N+1} \otimes y'_N} \\ = \Pr(Y'_N = y'_N, X_{N+1} = x_{N+1} | X_N = x_N, Y_N = y_N) \end{aligned}$$

over the ratchet's state space \mathcal{X} and a symbol alphabet \mathcal{Y} . This is the probability that the ratchet ends in state x_{N+1} and writes a symbol y'_N to the output sequence, given that the input symbol was y_N and the ratchet's state was x_N before the symbol-state interaction interval.

The Markovian dynamic describes the behavior of the joint event (ratchet-state \otimes symbol-value) during the interaction transition and leads to the transducer representation of the ratchet's functionality, illustrated in Fig. 2. As we use the terms, the *ratchet* refers to the physical device implementing the Markovian dynamic, whereas *transducer* refers to the computational mechanics state-transition machine (ϵ -transducer) that captures its information-theoretic functionalities in a compact way [36]. The form of the transducer is:

$$M_{x_N \rightarrow x_{N+1}}^{(y'_N | y_N)} = M_{x_N \otimes y_N \rightarrow x_{N+1} \otimes y'_N} . \quad (8)$$

The distinction between the Markov dynamic and the transducer representation is best illustrated graphically, as in the second column of Fig. 2.

The definition of a ratchet's memory involves its ϵ -transducer representation. In other words, memory is related to the size of the ratchet's causal state space $|\mathcal{X}|$ in its ϵ -transducer representation. (The very definition of ϵ -machines and ϵ -transducers entails that they have the minimal set of states for a given input, output, or input-output process.) As seen in the top middle of Fig. 2, memoryless ratchets have only a single internal (hidden) state: $|\mathcal{X}| = 1$. Thus, the ratchet behaves as a memoryless channel from input to output [52]. And, in this, it reduces temporal correlations in the input signal:

$$H'_1 - h'_\mu \leq H_1 - h_\mu , \quad (9)$$

according to the Data Processing Inequality [52]. As shown by Eq. (6) the difference between the single sym-

bol entropy and the entropy rate is the mutual information between the infinite past symbols and the current symbol. If the channel is memoryless, then $Y_N \rightarrow Y'_N$ and $Y_{0:N} \rightarrow Y'_{0:N}$, so by applying the Data Processing Inequality twice, we find $I[Y'_N; Y'_{0:N}] \leq I[Y_N; Y_{0:N}]$. Thus, the change in single-symbol entropy is a lower bound for the change in entropy rates [38]. In contrast, a memoryful ratchet has more than one state, $|\mathcal{X}| > 1$, and behaves as a memoryful channel [36]; bottom right of Fig. 2.

How the ratchet transduces the current input to the current output depends on in which state it is. As a result, the ratchet can create correlations in the output such that, regardless the input process:

$$H'_1 - h'_\mu \geq 0 . \quad (10)$$

Several explicit constructions of the output process based on given input and ratchet are shown in the last column of Fig. 2.

C. Thermodynamics of memory

This section considers the role of memory in the thermodynamic efficacy of information engines. In particular, we consider the average work production per cycle $\langle W \rangle$. The role can be explored in two complementary ways: either following the IPSL and related bounds, Eqs. (1) and (2), or from the exact expression of $\langle W \rangle$.

1. Information Processing Second Law bounds

The thermodynamics of memory is summarized in Fig. 3's table, where each row considers a different combination of input process and ratchet. This section addresses each cell in the table individually.

Consider the case of memoryless input and a memoryless ratchet. In Eq. (9), we saw that the temporal correlations in the input signal cannot be increased by such ratchets. Since the input signal is memoryless, the output signal must also be memoryless. For memoryless signals, however, we saw via Eq. (5) that the entropy rate h_μ is the same as the single-symbol entropy H_1 . We conclude that the temporal correlations vanish in both the input and output and, thus, that the single-symbol entropy input-to-output difference is the same as the entropy-rate difference:

$$H'_1 - h'_\mu = H_1 - h_\mu .$$

As a result both Eqs. (1) and (2) give the same bound

on the the average rate of work production:

$$\langle W \rangle \leq k_B T \ln 2 \Delta H_1 \quad (11)$$

$$= k_B T \ln 2 \Delta h_\mu , \quad (12)$$

This is noted at the right column in the table's first row.

Consider now the case of memoryful input with, again, a memoryless ratchet. A memoryful input contains temporal correlations that are decreased by the memoryless ratchet, from Eq. (9). The same equation implies that the single-symbol entropy difference is an upper bound on the entropy-rate difference. As a result, Eq. (2) provides a quantitatively tighter bound on the work production compared to the IPSL of Eq. (1) [38]:

$$\langle W \rangle \leq k_B T \ln 2 \Delta H_1$$

$$\leq k_B T \ln 2 \Delta h_\mu ,$$

These observations suggest that memoryless ratchets cannot leverage temporal correlations, since the stricter bound (single symbol) on work production stays fixed as we hold the single-symbol entropy fixed but vary the temporal correlations in the input. It appears that to leverage temporal correlations, one must use a memoryful ratchet.

We now address the case of memoryful ratchets. First, consider the case of memoryless inputs (no temporal correlations: $h_\mu = H_1$). From Eq. (10), we know that memoryful ratchets can create correlations in the output. In other words, the output signal is generally memoryful, implying $H'_1 - h'_\mu \geq 0$. As a result, $H'_1 - h'_\mu \geq H_1 - h'_\mu$, which implies $\Delta h_\mu \leq \Delta H_1$. And so, the change in entropy rate is a stricter bound than the change in single-symbol entropy:

$$\begin{aligned} \langle W \rangle &\leq k_B T \ln 2 \Delta h_\mu \\ &\leq k_B T \ln 2 \Delta H_1 , \end{aligned}$$

as seen in Table 3's second row. We explored this in some detail previously [27]. By calculating Δh_μ , we found a novel type of functionality in which the ratchet used stored work energy to increase temporal correlations in the input while simultaneously *increasing* the single-symbol uncertainty. The above relations also imply that memoryless ratchets may be best suited for leveraging memoryless input processes, since the bounds on work production for memoryless ratchets are higher than the bounds for memoryful ratchets.

Consider now a memoryful input driving a memoryful ratchet. In this case, memory in the ratchet is useful for work production. A companion work [28] considers a maximally correlated, period-2 input process, that has no single-symbol negentropy to leverage ($H_1 = 1$ bit of in-

Input Process	Ratchet Transducer	Output Process	Thermal Relations
memoryless input memoryless ratchet			 $0 = H_1 - h_\mu = H'_1 - h'_\mu$ $\langle W \rangle \leq \Delta h_\mu = \Delta H_1$
memoryless input memoryful ratchet			 $0 = H_1 - h_\mu \leq H'_1 - h'_\mu$ $\langle W \rangle \leq \Delta h_\mu \leq \Delta H_1$
memoryful input memoryless ratchet			 $0 \leq H'_1 - h'_\mu \leq H_1 - h_\mu$ $\langle W \rangle \leq \Delta H_1 \leq \Delta h_\mu$
memoryful input memoryful ratchet			 $H'_1 - h'_\mu \stackrel{?}{=} H_1 - h_\mu$ $\langle W \rangle \leq \Delta h_\mu$ $\langle W \rangle \leq \Delta H_1$

FIG. 3. The informational (IPSL) bounds on work that use Δh_μ or $\Delta H(1)$ depend critically on input signal and ratchet memory. In all finite memory cases, Δh_μ is a valid bound on $\langle W \rangle / k_B T \ln 2$, but the same is not true of ΔH_1 , as indicated in the far right column on thermal relations. The bounds shown in column have $k_B T \ln 2$ set to unity, so that the relations can be shown in compact form. If the ratchet is memoryless, then ΔH_1 is a valid and stronger bound than Δh_μ , because these channels decrease the temporal correlations in transducing input to output. For a memoryless input with memoryful ratchet, ΔH_1 is still a valid bound, but it is a weaker bound than Δh_μ , because a memoryful ratchet typically creates temporal correlations in its output. However, in the case where both input and output are memoryful, the ΔH_1 bound is invalid. It is violated by systems that turn temporal correlations into work by using ratchet memory to synchronize to the input memory.

formation), but that has maximal temporal correlations ($H_1 - h_\mu = 1$ bit). Notably, the single-symbol bound indicates that no work can be produced, since $\Delta H_1 \leq 0$ regardless of the output. Critically, though, the IPSL bound indicates that work production is possible, since $h'_\mu - h_\mu > 0$ as long as the output has some uncertainty in each sequential symbol. Indeed, Ref. [28] constructs a ratchet that produces positive work: $\langle W \rangle = k_B T \frac{1-\delta}{e}$, where $\delta \in (0, 1)$. Thus, the single-symbol bound is violated, but the IPSL bound is satisfied, as shown in Fig. 3's last row.

The final case to consider, in fact, is left out of Fig. 3: infinite-memory ratchets. This is because infinite memory ratchets do not necessarily have a steady state, so the IPSL bound in Ref. [27] does not hold. There are, as yet, no predictions for infinite-memory ratchets based on the information measures of the input or output processes. However, this is an intriguing case. And so, we turn to infinite ratchets in Sec. IV C.

Stepping back, Fig. 3's table details a constructive thermodynamic parallel to Ashby's Law of Requisite Variety: *Memory can leverage memory*. However, the bounds do not constitute existence proofs, since it is

not yet known if the specified bounds are achievable. Though, we constructed an example of a temporally correlated process that is best leveraged by memoryful ratchets, it is possible that there is an alternative temporally correlated input process that is best leveraged by a memoryless ratchet. Similarly, we see that the bounds on memoryless inputs are stricter for memoryful ratchets than for memoryless ratchets. If these bounds are not achievable, however, then this does not translate into a statement about the ratchet's actual efficiency in producing work.

Before addressing the puzzle of achievability, we need to determine the work production.

2. Exact work production

An exact expression for the average work production rate was introduced in Ref. [28]:

$$\langle W \rangle = k_B T \sum_{\substack{x, x' \in \mathcal{X} \\ y, y' \in \mathcal{Y}}} \pi_{x \otimes y} M_{x \otimes y \rightarrow x' \otimes y'} \ln \frac{M_{x' \otimes y' \rightarrow x \otimes y}}{M_{x \otimes y \rightarrow x' \otimes y'}}, \quad (13)$$

where $\{\pi_{x\otimes y}\}$ is the steady-state joint probability distribution of the ratchet and the input symbol before interaction. Heuristically, the formula can be understood in the following way. At the beginning of the interaction interval, the ratchet and the incoming bit have probability $\pi_{x\otimes y}$ to be in state $x\otimes y$. Thus, the joint system has the probability $\pi_{x\otimes y}M_{x\otimes y\rightarrow x'\otimes y'}$ to make the transition $x\otimes y\rightarrow x'\otimes y'$. Since M specifies a detailed-balanced thermal dynamic, the amount of energy extracted from the reservoir in each transition is given by the log-ratio $\ln(M_{x'\otimes y'\rightarrow x\otimes y}/M_{x\otimes y\rightarrow x'\otimes y'})$. The right-hand side of Eq. (13) therefore gives the average energy extracted from the heat reservoir every thermodynamic cycle. From the First Law of Thermodynamics, this must be the ratchet's average work production, since its energy is fixed in the steady state. Not only does the expression confirm our physical law of requisite memory, it also expands our understanding of the validity of IPSL-like bounds, as we see below.

Irrespective of the nature of the input, consider the case of memoryless ratchets for which we have:

$$\begin{aligned}\pi_{x\otimes y} &= \lim_{N\rightarrow\infty} \Pr(X_N = x, Y_N = y) \\ &= \lim_{N\rightarrow\infty} \Pr(Y_N = y) \\ &= \Pr(Y_N = y),\end{aligned}$$

simply the single-symbol probabilities of the input process. This follows since there is only a single ratchet state x . Thus, from Eq. (13), the only dependence the work has on the input process is on the latter's single-symbol distribution. In short, memoryless ratchets are insensitive to correlations in the inputs. To leverage correlations beyond single symbols in the input process it is necessary to add memory to the ratchet, as discussed in the previous section and in our companion work [28].

Conversely, as App. A establishes, if the input process is memoryless, there is no energetic advantage of using finite memoryful ratchets for binary input processes. For any finite memoryful ratchet that extracts work using the input process, there exists a memoryless ratchet that extracts at least as much work.

These two results confirm the intuition that to be thermodynamically optimal a ratchet's memory must match that of the input: Memoryful ratchets best leverage memoryful inputs and memoryless ratchets best leverage memoryless inputs.

IV. ACHIEVABILITY OF BOUNDS

The IPSL bound on average work production rate was derived based on the Second Law of Thermodynamics

applied to the joint evolution of the ratchet, the input-output symbol sequence, and the heat reservoir. Since the Second Law is merely an inequality, it does not guarantee that the bounds are actually achievable, at least for the class of information engines considered here. In point of fact, we saw that the bound cannot be saturated by memoryless ratchets. A somewhat opposite picture is presented by infinite-memory ratchets. And, understanding these is a necessity if we wish to build a thermodynamics of general computation; that is, of physically embedded universal Turing machines. As we will show shortly, infinite-memory ratchets can violate the IPSL bound since they can leverage the steady, indefinite increase in their own entropy to reduce the entropy of the heat reservoir, in addition to the contributions from an input signal. The following analyzes these cases individually.

A. Memoryless ratchets

This section applies the work expression of Eq. (13) to find optimal memoryless ratchets and then compares their optimal work production to the preceding information thermodynamics bounds to determine their achievability. Understanding the relationships between the memory of the ratchet and that of the input process, as discussed above, deepens the interpretation of the analysis. Since memoryless ratchets are insensitive to correlations, our calculated work productions are not only the work productions for memoryless inputs, but the work productions for all inputs with the same single-symbol statistical biases.

A memoryless ratchet's memory consists of a single state. As a result, the Markovian dynamic M acts only on individual input symbols. Thus, the work for any input process is a function only of its single-symbol distribution $\pi_y = \Pr(Y_N = y)$ (given M):

$$\langle W \rangle = k_B T \sum_{y, y' \in \mathcal{Y}} \pi_y M_{y \rightarrow y'} \ln \frac{M_{y' \rightarrow y}}{M_{y \rightarrow y'}}.$$

Here, we discuss in detail the particular case of a memoryless ratchet driven by binary inputs. The relevant class of transducers comprises all two-state HMMs over the state space $\{A\} \otimes \{0, 1\}$, where A is the ratchet's sole state. Since the transducers' state space is two-dimensional, the Markovian dynamic M is guaranteed to be detailed balanced. Moreover, we can parametrize this class by two transition probabilities p and q , as shown in Fig. 4. This, then, allows us to optimize over p and q to maximize work production.

For the ratchet shown in Fig. 4 driven by a process




FIG. 4. All possible memoryless ratchets that operate on a binary input, parametrized by transition probabilities p and q .

with single-symbol probabilities $\Pr(Y_N = 0) = b$ and $\Pr(Y_N = 1) = 1 - b$, the average work done is a function of b , p , and q :

$$\langle W \rangle(b, p, q) = k_B T (b - b') \ln \frac{p}{q}, \quad (14)$$

where $b' = b'(b, p, q) = (1 - q)b + (1 - b)p$ is the probability $\Pr(Y'_N = 0)$ of symbol 0 in the output. The expression for b' follows from the dynamic depicted in Fig. 4, whereas Eq. (14) follows from the fact that work $\ln(p/q)$ is gained for each transformation $0 \rightarrow 1$. For a given input bias b , optimization of the ratchet's transducer dynamic to produce maximal work yields ratchet parameters $p_{\max}(b)$ and $q_{\max}(b)$:

$$[q_{\max}(b), p_{\max}(b)] = \begin{cases} [\frac{1-b}{b \Omega(e^{(1-b)/b})}, 1], & 1/2 \leq b \leq 1 \\ [1, \frac{b}{(1-b) \Omega(e^{b/(1-b)})}], & 0 \leq b < 1/2 \end{cases},$$

where the function $\Omega(\cdot)$ is defined implicitly as $\Omega(ze^z) = z$. To confirm that these are indeed the maximal parameters for a given input $b \in [0, 1]$, note that $\langle W \rangle(b, p, q)$ is concave down over the physical parameter range $p, q \in [0, 1]$. For $b < 1/2$, plugging $q = 1$ and $p = b/(1 - b)\Omega(e^{b/(1-b)})$ into partial derivatives of the work yields $\partial \langle W \rangle / \partial q \geq 0$ and $\partial \langle W \rangle / \partial p = 0$. However, since q is already at its maximum, this a local maximum for the allowed parameter range. And, since the work is concave down, we know that this local maximum is the global maximum. The same can be shown for $b \geq 1/2$ by using the symmetry with respect the simultaneous exchanges $\{p \leftrightarrow q, b \leftrightarrow 1 - b\}$. Figure 5 shows how the optimal parameters depend on input bias $\Pr(Y_N = 0) = b$.

Substituting q_{\max} and p_{\max} into the work production expression, we find the maximum:

$$\langle W \rangle_{\max}(b) = \langle W \rangle(b, p_{\max}(b), q_{\max}(b)),$$

yielding the solid (blue) curve in Fig. 6. The curve is the maximum work production $\langle W \rangle_{\max}(b)$ of a memoryless ratchet for an input with bias b . This may seem like a limited result at first, since it was calculated by driving a memoryless ratchet with memoryless inputs. However, memoryless ratchets are insensitive to temporal corre-




FIG. 5. Optimal ratchet parameters $p_{\max}(b)$ (solid orange line) and $q_{\max}(b)$ (dashed blue line) are mirror images about $b = 1/2$. For $b < 1/2$, we set $p_{\max}(b) < 1$ and $q_{\max} = 1$ so that the interaction transition $1 \rightarrow 0$ has a positive energy change $\Delta E_{1 \rightarrow 0} = k_B T \ln(q/p)$ and, thus, absorbs heat from the thermal reservoir. The same reasoning applies to $b > 1/2$, where $p_{\max}(b) = 1$ and $q_{\max} < 1$. In the unique case where the input is all 1s, the most effective ratchet for generating work has $p_{\max} = 1/e$. Both functions realize a minimum value of $1/e$, as shown.

lations, and finite memory ratchets are no better than memoryless ratchets when driven by memoryless inputs, as discussed in the next section. Thus, the blue curve also represents the maximum work production of memoryless ratchets with memoryful inputs as well as of memoryful ratchets with memoryless inputs, as long as the probability of an input 0 is b .

To compare work production directly with the IPSL and related bounds, Eqs. (1) and (2), we need to calculate the changes in single-symbol entropy difference ΔH_1 and entropy-rate difference Δh_μ . Reminding ourselves that the ratchet is memoryless, these differences are the same if we assume the input to be memoryless. We find:

$$\begin{aligned} \Delta H_1 &= \Delta h_\mu(b, p, q) \\ &= H_B(b') - H_B(b), \end{aligned}$$

with $H_B(z) = H(\{z, 1 - z\})$ for $z \in [0, 1]$, the binary entropy function [52]. We obtain the bounds for an optimal ratchet, for a given input bias b , by substituting p_{\max} and q_{\max} for p and q , respectively. We plot this optimal bound as the dashed line (orange) in Fig. 6. Even though we maximized over the memoryless ratchet's parameters (p and q), the output work $\langle W \rangle_{\max}(b)$ falls far short of the bounds set on it, as the solid (blue) curve lies below the dashed (orange) curve except exactly at $b = 1/2$, where there is zero work production. This demonstrates that there are inherent inefficiencies in memoryless information ratchets.

There is a second source of inefficiency for memoryless




FIG. 6. Maximum work production $\langle W \rangle_{\max}$ for any input bias b is $k_B T/e$ (horizontal dashed line) and so ratchets do not achieve the IPSL upper bound $\langle W \rangle \leq k_B T \Delta h_{\mu}(b, p_{\max}, q_{\max})$ that derives from pure informational properties of the input and output processes. Also, $\Delta h_{\mu}(b, p_{\max}, q_{\max})$ itself is slightly less than the absolute maximum possible change in entropy $\Delta h_{\max}(b)$ given an input bias b . This means that a memoryless ratchet does not leverage all of the single-symbol statistical order in the input. This is also true of memoryful ratchets when driven by memoryless processes of the same input bias. Appendix A tells us that the blue curve is also the maximal work production for these ratchets with memoryless inputs.

ratchets. The maximum possible bound for the generated work comes from the case where there are no statistical biases and no correlations left in the output sequence, so that the output has maximum Shannon entropy. In this case we have $b' = 1/2$, the maximal entropy change being:

$$\Delta h_{\max}(b) = 1 - H_B(b) .$$

Figure 6 plots the corresponding bound as a dashed line (green), showing that it lies above the actual change in entropy for an optimal ratchet. Thus, not all of the order in the input sequence is being leveraged to generate work. In fact, the output bias $b'(b, p_{\max}, q_{\max})$ for an optimal ratchet is generally not equal to $1/2$.

B. Optimal memoryless ratchets versus memoryful ratchets

At this point we may ask: Is it possible to surpass the optimal memoryless ratchet in terms of work production with a memoryful ratchet? The answer seems to be negative for memoryless inputs. More to the point, Appendix A proves the following statement:

For memoryless, binary inputs work production by the optimal memoryless ratchet cannot be surpassed by any memoryful ratchet.

Thus, by optimizing over memoryless ratchets, we can actually determine the optimum work production over all finite memoryful ratchets. Appendix A proves that for sequences of binary symbols, memoryless ratchets are optimal for producing work.

This has a number of implications. First of all, it means that the dashed (blue) curve in Fig. 6 is not only a bound on the work production of a memoryless ratchet for any input with bias b , but it is also a bound on the work production of *any* finite memory ratchet with a memoryless input with the same bias. Second, in particular, the work production is at most $k_B T/e$, as shown by the dashed (red) horizontal line.

Third, importantly, this line is less than the conventional Landauer bound of $k_B T \ln 2$. It may seem counterintuitive that no single ratchet can autonomously achieve the Landauer bound, but this is a natural result of dynamics of this class of autonomous Maxwellian Demon. Between each interaction interval, as the ratchet switches between inputs, the Hamiltonian changes instantaneously and discontinuously. As a result, the ratchet and bit exist in a nonequilibrium distribution, which dissipates unrecoverable heat as it relaxes towards the Boltzmann distribution.

Finally, to achieve entropic bounds, the joint state of the ratchet and bit should follow the Boltzmann distribution during the interaction. However, to do this while also performing meaningful computation, it is necessary to implement some form of adiabatic protocol [53]. The latter can be implemented either by dynamically controlling the ratchet's energies over the interaction interval or by stringing together a series of ratchets, each one gradually updating the distribution of its input infinitesimally.

Appendix A's observation also suggests that multiple ratchets in series—the output sequence of one is input to the next—cannot be represented as a single finite-memory ratchet that interacts with one bit at a time and only once. This is because we can surpass the work production of an optimal memoryless ratchet with multiple ratchets interacting with multiple symbols at a time, as we noted already. Ratchets composed in series form a fundamentally different construction than a single memoryful ratchet; a topic of some biological importance to which we will return elsewhere.

C. Infinite-memory ratchets

We emphasized that the very general IPSL bound on information processing based on input-output entropy rate change holds for finite-state ratchets. What happens if infinite memory is available to a ratchet? This section constructs infinite-memory ratchets that can vi-

olate both Eqs. (1) and (2) and, by implication, Landauer’s bound. The intuition behind this is that, due to the infinite memory, the ratchet can continue indefinitely to store information that need not be written to the output. In effect, an apparent violation of the IPSL bound arises since the hidden degrees of freedom of the ratchet’s memory are not accounted for.

Nonetheless, infinite-memory ratchets offer intriguing possibilities for thermodynamic and computational functionality. While finite-memory ratchets can do meaningful computations and can even be appropriate models for, say, biological organisms that have finite information-processing capacities and response times, they cannot be computation universal in the current architecture [54, 55]. More precisely, a one-way universal Turing machine (UTM), like our ratchet, that reads its input once and never again, requires an “internal” infinite work tape to read from and write on. So, an infinite-state ratchet of our type is needed to emulate the infinite bidirectional read-write tape of the usual UTM [45].

Appendix A shows that memoryless ratchets are able to extract the most work from memoryless binary input processes, under the assumption that the ratchet’s memory is finite. Without finiteness the proof breaks down, since an asymptotic state distribution may not exist over infinite states [56]. In addition, the proof of Eq. (1) fails for the same reason. Thus, we turn to other tools for understanding the behavior in this case. The expression for work production still holds, so despite not having general informational bounds on work production, we can still calculate the exact work production for a prototype infinite ratchet.

Here, we present an infinite-state ratchet with finite energy-differences between all states. Our main result is that it produces more work than any finite memory ratchet for a given input. More to the point, it violates both the bounds in Eqs. (1) and (2). This demonstrates the need for the finite-memory assumption in developing Landauer and IPSL bounds. Consider, for example, an input process of all 1s. According to Sec. IV A, the maximum amount of work that can be extracted from this input by a memoryless ratchet is given by:

$$\langle W \rangle_{\max} = \frac{k_B T}{e} .$$

The discussion in App. A indicates that this should be the maximum amount of work that can be extracted by any finite-memory ratchet (for the same input). However, the infinite-state ratchet shown in Fig. 7 produces twice as much work, as we now show:

$$\langle W \rangle^{\infty} = \frac{2k_B T}{e} .$$

The infinite-state ratchet also violates both of the IPSL and single-symbol bounds, Eqs. (1) and (2), since $k_B T \ln 2$ is an upper bound for the work generation in all binary input processes according to these bounds, whereas $2/e > \ln 2$.

Let’s describe the structure and dynamics of the infinite-state ratchet in Fig. 7 in detail. This ratchet has a countably infinite number of states A_i , with $i \in \{0, 1, 2, \dots\}$. In other words, the ratchet state space is $\mathcal{X} = \{A_0, A_1, A_2, \dots\}$. The joint dynamic of the ratchet and the interacting symbol is shown in Fig. 7, where the arrows indicate allowed transitions and the number along the arrow, the associated transition probabilities. Apart from the case $i = 0$, only the following transitions are allowed: $A_i \otimes 1 \rightarrow \{A_{i\pm 1} \otimes 1, A_{i+1} \otimes 0\}$ and $A_i \otimes 0 \rightarrow A_j \otimes 1$ with $j = i/2$ for even i and $(i-1)/2$ for odd i . If the incoming symbol is 0, the only transition allowed involves a simultaneous change in the ratchet state and symbol, switching over to state $A_{j(i)}$ if it started in state A_i and the symbol switching to 1. The only exception is the case $i = 0$ in which the ratchet stays in the same state, while the symbol switches to 1. If the incoming symbol is 1, there are generally three possible transitions: $A_i \otimes 1 \rightarrow A_{i\pm 1} \otimes 1$ and $A_i \otimes 1 \rightarrow A_{i+1} \otimes 0$. The first two transitions occur with equal probabilities $1/2 - 1/e$, while the third transition occurs with probability $1/e$. For $i = 0$, there are four transitions possible: $A_0 \otimes 1 \rightarrow \{A_0 \otimes 1$ (self-loop), $A_1 \otimes 1, A_0 \otimes 0, A_1 \otimes 1\}$. The transition probabilities are shown in the figure.

We can assign relative energy levels for the joint states $A_i \otimes \{0, 1\}$ based on the transition probabilities. Since the (horizontal) transitions $A_i \otimes 1 \leftrightarrow A_{i+1} \otimes 1$ have equal forward and reverse transition probabilities, all the joint states $A_i \otimes 1$ have the same energy. Any state $A_i \otimes 0$ is higher than the state $A_{j(i) \otimes 1}$ by an energy:

$$\begin{aligned} \Delta E_{A_i \otimes 1 \rightarrow A_j \otimes 0} &= k_B T \ln \frac{1}{1/e} \\ &= k_B T . \end{aligned}$$

As a result, all states $A_i \otimes 0$ have the same energy, higher than that of the states $A_i \otimes 1$ by $k_B T$. This energy difference is responsible for producing the work. When the ratchet is driven by the all-1s process, if it is in an $A_i \otimes 0$ state after the previous interaction transition, then the switching transition changes the state to $A_i \otimes 1$ gaining $\Delta E_{A_i \otimes 0 \rightarrow A_i \otimes 1} = k_B T$ in work. The probability of being in a $Y_N = 0$ state after an interaction interval is $2/e$, so the work production is $\langle W \rangle = 2k_B T/e$, as stated above.

The reason this infinite-state ratchet violates the information-theoretic bounds is that those bounds ignore the asymptotic entropy production in the ratchet’s internal state space. There is no steady state over the infinite




FIG. 7. Infinite-state ratchet that violates the IP SL and single-symbol bounds, Eqs. (1) and (2), respectively. The ratchet state-space is $\mathcal{X} = \{A_0, A_1, A_2, \dots\}$; all states effectively have the same energy. The symbol values $\mathcal{Y} = \{0, 1\}$ differ by energy $\Delta E = k_B T$, with 0 having higher energy. The black arrows indicate the possible *interaction transitions* among the shown joint states of the ratchet and symbol during the interaction interval. For example, transitions $A_0 \otimes 1 \leftrightarrow A_1 \otimes 1$ are allowed whereas transitions $A_0 \otimes 0 \leftrightarrow A_1 \otimes 0$ are not. The dashed black lines show interaction transitions between the shown joint states and joint states that could not be shown. Briefly, for $i \geq 1$, there can be only the following interaction transitions: $A_i \otimes 1 \rightarrow \{A_{i\pm 1} \otimes 1, A_{2i} \otimes 0, A_{2i+1} \otimes 0\}$ and $A_i \otimes 0 \rightarrow A_{j(i)} \otimes 1$ with $j(i) = i/2$ for even i and $(i-1)/2$ for odd i . For the $i = 0$ transitions, see the diagram. Every interaction transition is followed by a switching transition and vice versa. The red dotted lines are the possible paths for driven *switching transitions* between the joint states, which correspond to the production or dissipation of work. During the switching interval, the only allowed transitions are the vertical transitions between energy levels $A_i \otimes 0 \leftrightarrow A_i \otimes 1$. The probability of these transitions depends on the input bias.

set of states and this leads to continual entropy production within the ratchet's state space \mathcal{X} . For the specific case of the all-1s input process note that, before the interaction interval, the joint state-space distribution of the ratchet and the incoming symbol must be positioned over only $A_i \otimes 1$ states. This is due to the fact that the switching transition always changes the symbol value to 1. From a distribution $\{\Pr(X_N = A_i, Y_N = 1)\}_{i \in \{0,1,\dots\}}$ over the $A_i \otimes 1$ states at time N , the interaction interval spreads the joint distribution to both $A_i \otimes 0$ and $A_i \otimes 1$ states. However, they are reset to a new distribution over the $A_i \otimes 1$ states $\{\Pr(X_{N+1} = A_i, Y_{N+1} = 1)\}_{i \in \{0,1,\dots\}}$ after the following switching transition. This leads to a spreading of the probability distribution—and, therefore, to an increase in entropy—in the ratchet space \mathcal{X} after each time step.

Figure 8 demonstrates the spreading by setting the initial joint ratchet-symbol state $X_0 \otimes Y_0$ to $A_0 \otimes 0$ and letting the distribution evolve for $N = 15$ time steps over the ratchet states. The ratchet states are indexed by i and the time steps are indexed by N , going from 1 to 15. The curves show the probabilities $\Pr(X_N = A_i)$ of the ratchet at time step N being in the i th ratchet state. By filling the area under each distribution curve and plotting the ratchet-state index in logarithm base 2, we see that the distribution's support doubles in size after every time step. This indicates an increase in the ratchet's internal entropy at each time step. This increase in internal en-




FIG. 8. Evolution of infinite-ratchet state distribution starting from an initial distribution peaked over the set \mathcal{X} , whose states are indexed A_i . The state distribution curves are plotted over 15 time steps, starting in red at time step $N = 1$ and slowly turning to blue at time step $N = 15$. With each sequential step, the support of the ratchet state distribution doubles in size, leading to increasing uncertainty in the ratchet state space and so increasing state entropy.

ropy is responsible for the violation of the IP SL bounds in Eqs. (1) and (2).

We have yet to discover a functional form for a steady state that is invariant—that maps to itself under one time-step. We made numerical estimates of the ratchet's entropy production, though. From the distributions




FIG. 9. The dashed (orange) line indicates average work production $\langle W \rangle$ per time step. It lies above the dotted (green) curve that indicates the IPSL entropy-rate bound on $\langle W \rangle$ (Eq. (1)), indicating a violation of the latter. The interpretation of the violation comes from the solid (blue) curve that indicates the joint entropy production of the input process and the ratchet together. We see a violation of the entropy-rate bound since there is continuous entropy production in the ratchet’s (infinite) state space.

shown in Fig. 8, we calculated the ratchet’s state entropies at each time step N . The entropy production $\Delta H[X_N] = H[X_{N+1}] - H[X_N]$ at the N th step is shown in Fig. 9. We see that the sum $\Delta H[X_N] + \Delta h_\mu$ of the changes in ratchet entropy and symbol entropy upper bounds the work production. Note that only the Δh_μ curve lies below the work production. Thus, while this infinite ratchet violates the IPSL bounds of Eqs. (1) and (2), it still satisfies a more general version of the Second Law of Thermodynamics for information ratchets—Eq. (A7) of Ref. [27]:

$$\langle W_N \rangle \leq k_B T \ln 2 (H_{N+1} - H_N) , \quad (15)$$

where W_N is the work gain at the N th time step and $H_N = H[X_N, Y_{N:\infty}, Y'_{0:N}]$ is the joint Shannon entropy of the ratchet and the input and output symbol sequences $Y_{N:\infty}$ and $Y'_{0:N}$, respectively, at time $t = N$. As we can see, this bound is based on not only the input and output process statistics, but also the ratchet memory.

CONCLUSION

How an agent interacts with and leverages it’s environment is a topic of broad interest, from engineering and cybernetics to biology and now physics [16, 57]. General principles for how the structure of an agent must match that of its environment will become essential tools for understanding how to take thermodynamic advantage of correlations in structured environments, whether the

correlations are temporal or spatial. Ashby’s Law of Requisite Variety—a controller must have at least the same variety as its input so that the whole system can adapt to and compensate that variety and achieve homeostasis [16]—was an early attempt at such a general principle of regulation and control. In essence, a controller’s variety should match that of its environment. Above, paralleling this, we showed that a near-optimal thermal agent (information engine) interacting with a structured input (information reservoir) obeys a similar variety-matching principle.

For an efficient finite-state information ratchet, the ratchet memory should reflect the memory of the input process. More precisely, memoryless ratchets are optimal for leveraging memoryless inputs, while memoryful ratchets are optimal for leveraging memoryful inputs. This can be appreciated in a two different ways.

On the one hand, the first comes from information processing properties of the ratchet and input and the associated IPSL bounds on work. The operation of memoryless ratchets can only destroy temporal correlations. These ratchets’ work production is still bounded by single-symbol entropy changes, as in Eq. (2). And, since memoryless input processes only produce single-symbol correlations (statistical biases), the memoryless ratchet bound of Eq. (2) allows for maximal work production. Thus, according to their bounds, memoryless ratchets and inputs produce the most work *when paired*.

On the other hand, in the second view memoryful input processes exhibit multiple-symbol temporal correlations. And, the entropy rate bound of Eq. (1) suggests that the memoryful input processes can be used to produce work in a memoryful ratchet, but not a memoryless one. More precisely, we can conceive of memoryful input processes whose single-symbol statistics are unbiased (equal proportions of 0s and 1s, in case of binary alphabet) but the entropy rate is smaller than the single-symbol entropy: $h_\mu < H_1 (= \ln 2 \text{ for a binary alphabet})$. In this case, since the single-symbol entropy is already at its maximum possible value, memoryless ratchets are unable to extract any work. Since the memoryful ratchets satisfy the IPSL bound of Eq. (1), however, they can extract work from such memoryful processes. One such example is studied in detail by Ref. [28]. (For a quantum-mechanical ratchet, compare Ref. [58].) Thus, memoryful ratchets are best paired with memoryful inputs. This and its complement result—memoryless inputs are optimally used by memoryless ratchets—is biologically suggestive. If one observes memory (temporal correlations) in the transduction implemented by a biomolecular assembly, for example, then it has adapted to some structured environment.

We summarized the role of memory in thermodynamic

processes in Fig. 3 which considers each of the four possible combinations of memoryful or memoryless ratchets with memoryful or memoryless input.

While the Second Law of Thermodynamics determines the IPSL and related bounds discussed here, it does not follow that the bounds are achievable for the class of information ratchets considered. Based on an exact method for calculating the average work production [28], we saw that there are indeed situations where the bounds are not achievable (Fig. 6). In Sec. IV A, we saw that memoryless ratchets cannot generally saturate their bound (Eq. (2)). Furthermore, based on the results of App. A we could prove that finite memoryful ratchets fare no better than memoryless ratchets at leveraging memoryless inputs. Thus, not even memoryful ratchets can extract the maximum amount of work possible from a memoryless input. There are some hints, though, as to what the architecture of information engines should be to extract the maximum possible work allowed by the Second Law. We alluded to one such situation in Sec. IV A involving a “swarm” of memoryless, optimized ratchets.

The unattainability of the IPSL bound observed above pertains to the architecture of information engines where there is only a single ratchet that interacts with one environmental signal value at a time. This leads one to speculate that multiple ratchets interacting with different signals—say, chained together so that the output of one is the input of another—will lead to a closer approach to the bound. Simply having multiple copies of the optimal memoryless ratchets one after another, however, will not necessarily address unattainability. Interestingly, depending on input bias b , there may be oscillations in the amount of work that is gained per cycle. And, even with infinitely many ratchets chained together sequentially, we may still be far from the IPSL bound. Based on our intuition about thermodynamically reversible processes, we postulate that to approach the bound more closely we need increasingly many memoryless ratchets, each optimized with respect to *its own input*. We leave the verification of this intuition for a future investigation. This does suggest, though, architectural trade-offs that should manifest themselves in evolved biological thermodynamic processes.

To complete our exploration of the role of memory in thermodynamic processes, we considered infinite-state ratchets, which are necessary if we wish to physically implement universal Turing machines with the unidirectional information ratchets. Infinite ratchets, however, pose a fundamental challenge since the IPSL entropy-rate bound on work production does not apply to them. The proof of the bound (Eq. (1) [28]) is based on the assumption that the ratchet reaches a steady state after interacting with sufficiently many input symbols. This

need not be the case for infinite-state ratchets. In fact, the numerical investigations of Sec. IV C indicate that the probability distribution in the state space of an infinite ratchet can continue to spread indefinitely, without any sign of relaxing to a steady state; recall Fig. 8. By calculating both the average work production per time step and the amount of change in the entropy rate, Fig. 9 showed that there is a violation of the IPSL and related bounds. This necessitates a modification of the IPSL for infinite-state ratchets. The appropriate bound, though, has already been presented in a previous work [28], which we quoted in Eq. (15). This relation shows that the work production is still bounded by the system’s entropy production; only, we must include the contribution from the ratchet’s internal state space on top of the entropy-rate difference of the input and the output HMMs.

We close by highlighting the close correspondence between information ratchets and biological enzymes. Most directly, it is possible to model the biomimetic enzymes following the design of information ratchets [59]. The correspondence goes further, though. In Sec. II, we discussed how a swarm of ratchets acting cooperatively may be more efficient than individual information ratchets, even if they are quite sophisticated. A similar phenomenon holds for enzymes where the enzymes along a metabolic pathway assemble to form a multi-enzyme complex—a “swarm”—to affect faster, efficient reaction turnover, known as *substrate channeling* [60].

ACKNOWLEDGMENTS

As an External Faculty member, JPC thanks the Santa Fe Institute for its hospitality during visits. This work was supported in part by FQXi Grant number FQXi-RFP-1609 and the U. S. Army Research Laboratory and the U. S. Army Research Office under contracts W911NF-13-1-0390 and W911NF-12-1-0234.

Appendix A: Optimally Leveraging Memoryless Inputs

It is intuitively appealing to think that memoryless inputs are best utilized by memoryless ratchets. In other words, the optimal ratchet for a memoryless input is a memoryless ratchet. We prove the validity of this intuition in the following. We start with the expression of

work production per time step:

$$\begin{aligned}
\beta\langle W \rangle &= \sum_{x,x',y,y'} \pi_{x\otimes y} M_{x\otimes y \rightarrow x'\otimes y'} \ln \frac{M_{x'\otimes y' \rightarrow x\otimes y}}{M_{x\otimes y \rightarrow x'\otimes y'}} \\
&= \sum_{x,x',y,y'} \pi_{x\otimes y} M_{x\otimes y \rightarrow x'\otimes y'} \ln \frac{\pi_{x'\otimes y'} M_{x'\otimes y' \rightarrow x\otimes y}}{\pi_{x\otimes y} M_{x\otimes y \rightarrow x'\otimes y'}} \\
&\quad - \sum_{x,x',y,y'} \pi_{x\otimes y} M_{x\otimes y \rightarrow x'\otimes y'} \ln \frac{\pi_{x'\otimes y'}}{\pi_{x\otimes y}},
\end{aligned}$$

with $\beta = 1/k_B T$. The benefit of the decomposition in the second line will be clear in the following. Let us introduce several quantities that will also be useful in the following:

$$\begin{aligned}
p(x, y, x', y') &= \pi_{x\otimes y} M_{x\otimes y \rightarrow x'\otimes y'}, \\
p_R(x, y, x', y') &= \pi_{x'\otimes y'} M_{x'\otimes y' \rightarrow x\otimes y}, \\
\pi_x^X &= \sum_y \pi_{x\otimes y}, \\
\pi_y^Y &= \sum_x \pi_{x\otimes y}, \\
p^X(x, x') &= \sum_{y,y'} p(x, y, x', y'), \text{ and} \\
p^Y(y, y') &= \sum_{x,x'} p(x, y, x', y').
\end{aligned}$$

For a memoryless input process, sequential inputs are statistically independent. This implies Y_N and X_N are independent, so the stationary distribution $\pi_{x\otimes y}$ can be written as a product of marginals:

$$\pi_{x\otimes y} = \pi_x^X \pi_y^Y. \quad (\text{A1})$$

In terms of the above quantities, we can rewrite work for a memoryless input process as:

$$\begin{aligned}
\beta\langle W \rangle &= -D_{KL}(p\|p_R) \\
&\quad - \sum_{y,y'} p^Y(y, y') \ln \frac{\pi_{y'}^Y}{\pi_y^Y} - \sum_{x,x'} p^X(x, x') \ln \frac{\pi_{x'}^X}{\pi_x^X},
\end{aligned}$$

where $D_{KL}(p\|p_R)$ is the relative entropy of the distribution p with respect to p_R [52]. Note that the last term in the expression vanishes, since the ratchet state distribution is the same before and after an interaction interval:

$$\sum_x p^X(x, x') = \sum_x p^X(x', x) = \pi_{x'}^X, \quad (\text{A2})$$

and so:

$$\begin{aligned}
&\sum_{x,x'} p^X(x, x') \ln \frac{\pi_{x'}^X}{\pi_x^X} \\
&= \sum_{x,x'} p^X(x, x') \ln \pi_{x'}^X - \sum_{x,x'} p^X(x, x') \ln \pi_x^X \\
&= \sum_{x'} \pi_{x'}^X \ln \pi_{x'}^X - \sum_x \pi_x^X \ln \pi_x^X \\
&= 0.
\end{aligned}$$

Thus, we find the average work production to be:

$$\beta\langle W \rangle = -D_{KL}(p\|p_R) - \sum_{y,y'} p^Y(y, y') \ln \frac{\pi_{y'}^Y}{\pi_y^Y}. \quad (\text{A3})$$

Let us now use the fact that the coarse graining of any two distributions, say p and q , yields a smaller relative entropy between the two [52, 61]. In the work formula, p^Y is a coarse graining of p and p_R^Y is a coarse graining of p_R , implying:

$$D_{KL}(p^Y\|p_R^Y) \leq D_{KL}(p\|p_R). \quad (\text{A4})$$

Combining the above relations, we find the inequality:

$$\beta\langle W \rangle \leq -D_{KL}(p^Y\|p_R^Y) - \sum_{y,y'} p^Y(y, y') \ln \frac{\pi_{y'}^Y}{\pi_y^Y}.$$

Now, the marginal transition probability $p^Y(y, y')$ can be broken into the product of the stationary distribution over the input variable π_y^Y and a Markov transition matrix $M_{y \rightarrow y'}^Y$ over the input alphabet:

$$p^Y(y, y') = \pi_y^Y M_{y \rightarrow y'}^Y,$$

which for any ratchet M is:

$$\begin{aligned}
M_{y \rightarrow y'}^Y &= \frac{1}{\pi_y^Y} p^Y(y, y') \\
&= \frac{1}{\pi_y^Y} \sum_{x,x'} \pi_{x\otimes y} M_{x\otimes y \rightarrow x'\otimes y'} \\
&= \frac{1}{\pi_y^Y} \sum_{x,x'} \pi_x^X \pi_y^Y M_{x\otimes y \rightarrow x'\otimes y'} \\
&= \sum_{x,x'} \pi_x^X M_{x\otimes y \rightarrow x'\otimes y'}.
\end{aligned}$$

We can treat the Markov matrix M^Y as corresponding to a ratchet in the same way as M . Note that M^Y is effectively a memoryless ratchet since we do not need to refer to the internal states of the corresponding ratchet. See Fig. 2. The resulting work production for this ratchet

$\langle W^Y \rangle$ can be expressed as:

$$\begin{aligned} \beta \langle W^Y \rangle &= \sum_{y,y'} \pi_y^Y M_{y \rightarrow y'}^Y \ln \frac{M_{y' \rightarrow y}^Y}{M_{y \rightarrow y'}^Y} \\ &= -D_{KL}(p^Y \| p_R^Y) - \sum_{y,y'} p^Y(y, y') \ln \frac{\pi_{y'}^Y}{\pi_y^Y} \\ &\geq \beta \langle W \rangle . \end{aligned}$$

Thus, for any memoryful ratchet driven by a memoryless input we can design a memoryless ratchet that extracts at least as much work as the memoryful ratchet.

There is, however, a small caveat. Strictly speaking, we must assume the case of binary input. This is due to the requirement that the matrix M be detailed balanced (see Sec. II) so that the expression of work used here is appropriate. More technically, the problem is that we do not yet have a proof that if M is detailed balanced then so is M^Y , a critical requirement above. In fact, there are examples where M^Y does not exhibit detailed balance. We do, however, know that M^Y is guaranteed to be detailed balanced if \mathcal{Y} is binary, since that means M^Y only has two states and all flows must be balanced. Thus, for memoryless binary input processes, we established that there is little point in using finite memoryful ratchets to extract work: memoryless ratchets extract work optimally from memoryless binary inputs.

Appendix B: An IPSL for Information Engines

Reference [27] proposed a generalization of the Second Law of Thermodynamics to information processing systems (IPSL, Eq. (1)) under the premise that the Second Law can be applied even when the thermodynamic entropy of the information bearing degrees of freedom is taken to be their Shannon information entropy. This led to a consistent prediction of the thermodynamics of information engines. It was also validated through numerical calculations. This appendix proves this assertion for the class of information engines considered here. The key idea is to use the irreversibility of the Markov chain dynamics followed by the engine and by the information bearing degrees of freedom to derive the IPSL inequality.

For the sake of presentation, we introduce new notation here. We refer to the engine as the demon D, following the original motivation for information engines. We refer to the information-bearing two-state systems as the bits B. According to our set up, D interacts with an infinite sequence of bits, $B_0 B_1 B_2 \dots$ as shown in Fig. 10. The figure also explains the connection of the current terminology to that in the main text. In particular, we show two snapshots of our setup, at times $t = N$ and




FIG. 10. The demon D interacts with one bit at a time for a fixed time interval; for example, with bit B_N for the time interval $t = N$ to $t = N + 1$. During this, the demon changes the state of the bit from (input) Y_N to (output) Y'_N . There is an accompanying change in D's state as well, not shown. The joint dynamics of D and B_N is governed by the Markov chain $M_{D \otimes B_N}$.

$t = N + 1$. During that interval D interacts with bit B_N and changes it from (input) symbol Y_N to (output) symbol Y'_N . The corresponding dynamics is governed by the Markov transition matrix $M_{D \otimes B_N}$ which acts *only* on the joint subspace of D and B_N .

Under Markov dynamics the relative entropy of the current distribution with respect to the asymptotic steady-state distribution is a monotonically decreasing function of time. We now use this property for the transition matrix $M_{D \otimes B_N}$ to derive the IPSL. Denote the distribution of D's states and the bits B at time t by $P_{DB_{0:\infty}}(t)$. Here, $B_{0:\infty}$ stands for all the information-bearing degrees of freedom [62]. The steady-state distribution corresponding to the operation of $M_{D \otimes B_N}$ is determined via:

$$\lim_{n \rightarrow \infty} M_{D \otimes B_N}^n P_{DB_{0:\infty}}(N) = \pi_{DB_N}^{\text{eq}} P_{B_{0:\infty/N}}(N) \quad (\text{B1})$$

$$\equiv \pi^s(N) , \quad (\text{B2})$$

where $\pi_{DB_N}^{\text{eq}}$ denotes the steady-state distribution:

$$M_{D \otimes B_N} \pi_{DB_N}^{\text{eq}} = 0$$

and $P_{B_{0:\infty/N}}(N)$ the marginal distribution of all the bits other than the N -th bit at time $t = N$. We introduce $\pi^s(N)$ in Eq. (B2) for brevity.

The rationale behind the righthand side of Eq. (B1) is that the matrix $M_{D \otimes B_N}$ acts only on D and B_N , sending to their joint distribution to the stationary distribution $\pi_{DB_N}^{\text{eq}}$ (on repeated operation), while leaving intact the marginal distribution of the rest of B. The superscript eq emphasizes the fact the distribution $\pi_{DB_N}^{\text{eq}}$ is an equilibrium distribution, as opposed to a nonequilibrium steady-state distribution, due to the assumed detailed-balance condition on $M_{D \otimes B_N}$. In other words, $\pi_{DB_N}^{\text{eq}}$ follows the

Boltzmann distribution:

$$\pi_{\text{DB}_N}^{\text{eq}}(D = x, B_N = y) = e^{\beta[F_{\text{DB}_N} - E_{\text{DB}_N}(x, y)]} \quad (\text{B3})$$

for inverse temperature β , free energy F_{DB_N} , and energy $E_{\text{DB}_N}(x, y)$. In the current notation we express the monotonicity of relative entropy as:

$$D(P_{\text{DB}_{0:\infty}}(N) \parallel \pi^s(N)) \geq D(P_{\text{DB}_{0:\infty}}(N+1) \parallel \pi^s(N)), \quad (\text{B4})$$

where $D(p \parallel q)$ denotes the relative entropy of the distribution p with respect to q :

$$D(p \parallel q) = \sum_i p(i) \ln \left[\frac{p(i)}{q(i)} \right]$$

over D's states i . The IPSL is obtained as a consequence of inequality Eq. (B4), as we now show [63].

First, we rewrite the lefthand side of Eq. (B4) as:

$$\begin{aligned} & D(P_{\text{DB}_{0:\infty}}(N) \parallel \pi^s(N)) \\ &= -H_{\text{DB}_{0:\infty}}(N) \ln 2 - \sum_{\text{DB}_{0:\infty}} P_{\text{DB}_{0:\infty}}(N) \ln \pi^s(N) \\ &= -H_{\text{DB}_{0:\infty}}(N) \ln 2 - \sum_{\text{DB}_N} P_{\text{DB}_N}(N) \ln \pi_{\text{DB}_N}^{\text{eq}} \\ &\quad - \sum_{\text{DB}_{0:\infty/N}} P_{\text{DB}_{0:\infty/N}}(N) \ln P_{\text{B}_{0:\infty/N}}(N) \\ &= -H_{\text{DB}_{0:\infty}}(N) \ln 2 - \beta F_{\text{DB}_N} + \beta \langle E_{\text{DB}_N} \rangle(N) \\ &\quad + H_{\text{B}_{0:\infty/N}}(N) \ln 2. \end{aligned} \quad (\text{B5})$$

The first line applies the definition of relative entropy. Here, H_X denotes the Shannon entropy of random variable X in *information units* of bits (base 2). The second line employs the expression of $\pi^s(N)$ given in Eq. (B2). The final line uses the Boltzmann form of $\pi_{\text{DB}_N}^{\text{eq}}$ given in Eq. (B3). Here, $\langle E_{\text{DB}_N} \rangle(N)$ denotes the average energy of D and the interacting bit B_N at time $t = N$.

Second, in a similar way, we have the following expres-

sion for the righthand side of Eq. (B4):

$$\begin{aligned} & D(P_{\text{DB}_{0:\infty}}(N+1) \parallel \pi_N^s) \\ &= -H_{\text{DB}_{0:\infty}}(N+1) \ln 2 + H_{\text{B}_{0:\infty/N}}(N) \ln 2 \\ &\quad + \beta \langle E_{\text{DB}_N} \rangle(N+1) - \beta F_{\text{DB}_N}. \end{aligned} \quad (\text{B6})$$

Note that the marginal distribution of the noninteracting bits $\text{B}_{0:\infty/N}$ does not change over the time interval $t = N$ to $t = N+1$ since the matrix $M_{\text{D} \otimes \text{B}_N}$ acts only on D and B_N , and the Shannon entropy of the noninteracting bits remains unchanged over the interval.

Third, combining Eqs. (B4), (B5), and (B6), we get the inequality:

$$\ln 2 \Delta H_{\text{DB}_{0:\infty}} - \beta \Delta \langle E_{\text{DB}_N} \rangle \geq 0, \quad (\text{B7})$$

where $\Delta H_{\text{DB}_{0:\infty}}$ is the change in the Shannon entropy of D and B and $\Delta \langle E_{\text{DB}_N} \rangle$ is the change in the average energy of D and B over the interaction interval.

Fourth, according to the ratchet's design, D and B are decoupled from the work reservoir during the interaction intervals. (The work reservoir is connected only at the end points of intervals, when one bit is replaced by another.) From the First Law of Thermodynamics, the increase in energy $\Delta \langle E_{\text{DB}_N} \rangle$ comes from the heat reservoir. In other words, we have the relation:

$$\Delta \langle E_{\text{DB}_N} \rangle = \langle \Delta Q \rangle, \quad (\text{B8})$$

where ΔQ is the heat given to the system. (In fact, Eq. (B8) is valid for each realization of the dynamics, not just on the average, since the conservation of energy holds in each realization.)

Finally, combining Eqs. (B7) and (B8), we get:

$$\ln 2 \Delta H_{\text{DB}_{0:\infty}} - \beta \langle \Delta Q \rangle \geq 0, \quad (\text{B9})$$

which is the basis of the IPSL as demonstrated in Ref. [27]; see, in particular, Eq. (A7) there.

-
- [1] C. E. Shannon and J. McCarthy, editors. *Automata Studies*. Number 34 in Annals of Mathematical Studies. Princeton University Press, Princeton, New Jersey, 1956.
 - [2] N. Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. Wiley, New York, 1949.
 - [3] N. Wiener. Nonlinear prediction and dynamics. In *Norbert Wiener, Collected Works III*. MIT Press, Cambridge, 1981.
 - [4] C. E. Shannon. A mathematical theory of communication. *Bell Sys. Tech. J.*, 27:379–423, 623–656, 1948.
 - [5] C. E. Shannon. Communication theory of secrecy systems. *Bell Sys. Tech. J.*, 28:656–715, 1949.
 - [6] C. E. Shannon. Coding theorems for a discrete source with a fidelity criterion. *IRE National Convention Record, Part 4*, 7:142–163, 623–656, 1959.
 - [7] C. E. Shannon. Two-way communication channels. In *Proc. Fourth Berkeley Symp. on Math. Statist. and Prob.*, volume 1, pages 611–644. University of California Press, 1961.
 - [8] N. Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT, Cambridge, 1948.

- [9] N. Wiener. *The Human Use Of Human Beings: Cybernetics And Society*. Da Capo Press, Cambridge, 1988.
- [10] L. von Bertalanffy. *General System Theory: Foundations, Development, Applications*. Penguin University Books, New York, revised edition, 1969.
- [11] W. R. Ashby. *Design for a Brain: The Origin of Adaptive Behavior*. Chapman and Hall, New York, second edition, 1960.
- [12] H. Quastler. The status of information theory in biology—A roundtable discussion. In H. P. Yockey, editor, *Symposium on Information Theory in Biology*, page 399. Pergamon Press, New York, 1958.
- [13] F. Conway and J. Siegelman. *Dark Hero of the Information Age: In Search of Norbert Wiener The Father of Cybernetics*. Basic Books, New York, New York, 2006.
- [14] J. P. Crutchfield. Between order and chaos. *Nature Physics*, 8(January):17–24, 2012.
- [15] R. Klages, W. Just, and C. Jarzynski, editors. *Nonequilibrium Statistical Physics of Small Systems: Fluctuation Relations and Beyond*. Wiley, New York, 2013.
- [16] W. R. Ashby. *An Introduction to Cybernetics*. John Wiley and Sons, New York, second edition, 1960.
- [17] L. Szilard. On the decrease of entropy in a thermodynamic system by the intervention of intelligent beings. *Z. Phys.*, 53:840–856, 1929.
- [18] H. Leff and A. Rex. *Maxwell's Demon 2: Entropy, Classical and Quantum Information, Computing*. Taylor and Francis, New York, 2002.
- [19] A. N. Kolmogorov. Three approaches to the concept of the amount of information. *Prob. Info. Trans.*, 1:1, 1965.
- [20] A. N. Kolmogorov. Combinatorial foundations of information theory and the calculus of probabilities. *Russ. Math. Surveys*, 38:29–40, 1983.
- [21] G. Chaitin. On the length of programs for computing finite binary sequences. *J. ACM*, 13:145, 1966.
- [22] P. M. B. Vitanyi. *Introduction to Kolmogorov complexity and its applications*. ACM Press, Reading, 1990.
- [23] E. T. Jaynes. Information theory and statistical mechanics. *Phys. Rev.*, 106:620–630, 1957.
- [24] J. M. Horowitz and S. Vaikuntanathan. Nonequilibrium detailed fluctuation theorem for repeated discrete feedback. *Phys. Rev. E*, 82:061120, 2010.
- [25] A. B. Boyd and J. P. Crutchfield. Demon dynamics: Deterministic chaos, the Szilard map, and the intelligence of thermodynamic systems. *Phys. Rev. Lett.*, 116:190601, 2016.
- [26] D. Mandal and C. Jarzynski. Work and information processing in a solvable model of Maxwell's demon. *Proc. Natl. Acad. Sci. USA*, 109(29):11641–11645, 2012.
- [27] A. B. Boyd, D. Mandal, and J. P. Crutchfield. Identifying functional thermodynamics in autonomous Maxwellian ratchets. *New J. Physics*, 18:023049, 2016.
- [28] A. B. Boyd, D. Mandal, and J. P. Crutchfield. Correlation-powered information engines and the thermodynamics of self-correction. *Phys. Rev. E*, 95(1):012152, 2017.
- [29] R. Landauer. Irreversibility and heat generation in the computing process. *IBM J. Res. Develop.*, 5(3):183–191, 1961.
- [30] C. H. Bennett. Thermodynamics of computation - a review. *Intl. J. Theo. Phys.*, 21:905, 1982.
- [31] S. Deffner and C. Jarzynski. Information processing and the second law of thermodynamics: An inclusive, Hamiltonian approach. *Phys. Rev. X*, 3:041003, 2013.
- [32] A. C. Barato and U. Seifert. Stochastic thermodynamics with information reservoirs. *Phys. Rev. E*, 90:042150, 2014.
- [33] L. Brillouin. Maxwell's demon cannot operate: Information and entropy I. *J. Appl. Phys.*, 22:334–337, 1951.
- [34] C. H. Bennett. Demons, engines and the Second Law. *Sci. Am.*, 257(5):108–116, 1987.
- [35] Z. Lu, D. Mandal, and C. Jarzynski. Engineering Maxwell's demon. *Physics Today*, 67(8):60–61, January 2014.
- [36] N. Barnett and J. P. Crutchfield. Computational mechanics of input-output processes: Structured transformations and the ϵ -transducer. *J. Stat. Phys.*, 161(2):404–451, 2015.
- [37] T. Sagawa and M. Ueda. Generalized Jarzynski equality under nonequilibrium feedback control. *Phys. Rev. Lett.*, 104:090602, 2010.
- [38] N. Merhav. Sequence complexity and work extraction. *J. Stat. Mech.*, page P06037, 2015.
- [39] D. Mandal, H. T. Quan, and C. Jarzynski. Maxwell's refrigerator: an exactly solvable model. *Phys. Rev. Lett.*, 111:030602, 2013.
- [40] A. C. Barato and U. Seifert. An autonomous and reversible Maxwell's demon. *Europhys. Lett.*, 101:60001, 2013.
- [41] A. C. Barato and U. Seifert. Unifying three perspectives on information processing in stochastic thermodynamics. *Phys. Rev. Lett.*, 112:090601, 2014.
- [42] A. N. Kolmogorov. Entropy per unit time as a metric invariant of automorphisms. *Dokl. Akad. Nauk. SSSR*, 124:754, 1959. (Russian) *Math. Rev.* vol. 21, no. 2035b.
- [43] Ja. G. Sinai. On the notion of entropy of a dynamical system. *Dokl. Akad. Nauk. SSSR*, 124:768, 1959.
- [44] J. P. Crutchfield and D. P. Feldman. Regularities unseen, randomness observed: Levels of entropy convergence. *CHAOS*, 13(1):25–54, 2003.
- [45] J. G. Brookshear. *Theory of computation: Formal languages, automata, and complexity*. Benjamin/Cummings, Redwood City, California, 1989.
- [46] L. R. Rabiner and B. H. Juang. An introduction to hidden Markov models. *IEEE ASSP Magazine*, January, 1986.
- [47] L. R. Rabiner. A tutorial on hidden Markov models and selected applications. *IEEE Proc.*, 77:257, 1989.
- [48] R. J. Elliot, L. Aggoun, and J. B. Moore. *Hidden Markov Models: Estimation and Control*, volume 29 of *Applications of Mathematics*. Springer, New York, 1995.
- [49] Y. Ephraim and N. Merhav. Hidden markov processes. *IEEE Trans. Info. Th.*, 48(6):1518–1569, 2002.
- [50] C. R. Shalizi and J. P. Crutchfield. Computational mechanics: Pattern and prediction, structure and simplicity. *J. Stat. Phys.*, 104:817–879, 2001.

- [51] J. P. Crutchfield. The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75:11–54, 1994.
- [52] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience, New York, second edition, 2006.
- [53] A. del Campo, J. Goold, and M. Parnostro. More bang for your buck: Super-adiabatic quantum engines. *Scientific Reports*, 4:6208, 2014.
- [54] M. Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, Englewood Cliffs, New Jersey, 1967.
- [55] H. R. Lewis and C. H. Papadimitriou. *Elements of the Theory of Computation*. Prentice-Hall, Englewood Cliffs, N.J., second edition, 1998.
- [56] J. G. Kemeny and J. L. Snell. *Denumerable Markov Chains*. Springer, New York, second edition, 1976.
- [57] M. Ehrenberg and C. Blomberg. Thermodynamic constraints on kinetic proofreading in biosynthetic pathways. *Biophys. J.*, 31:333–358, 1980.
- [58] A. Chapman and A. Miyake. How can an autonomous quantum Maxwell demon harness correlated information? arXiv:1506.09207.
- [59] Y. Cao, Z. Gong, and H. T. Quan. Thermodynamics of information processing based on enzyme kinetics: An exactly solvable model of an information pump. *Phys. Rev. E*, 91:062117, 2015.
- [60] R. Phillips, J. Kondev, J. Theriot, and N. Orme. *Physical biology of the cell*. Garland Science, New York, USA, 2008.
- [61] A. Gomez-Marin, J. M. R. Parrondo, and C. Van den Broeck. Lower bounds on dissipation upon coarse graining. *Phys. Rev. E*, 78:011107, 2008.
- [62] Probability distributions over infinitely many degrees of freedom ultimately require a measure-theoretic treatment. This is too heavy a burden in the current context. The difficulties can be bypassed by assuming that the number of bits in the infinite information reservoir is a large, but positive finite integer L . And so, instead of infinities in $B_{0:\infty}$ and $B_{0:\infty/N}$ we use $B_{0:L}$ and $B_{0:L/N}$, respectively, and take the appropriate limit when needed.
- [63] For a somewhat similar approach see N. Merhav. *J. Stat. Mech.: Th. Expt.* (2017) 023207.