

Work, Entropy Production, and Thermodynamics of Information under Protocol Constraints

Artemy Kolchinsky^{*} and David H. Wolpert[†]
Santa Fe Institute, Santa Fe, New Mexico 87501, USA

 (Received 21 November 2020; revised 6 July 2021; accepted 9 August 2021; published 3 November 2021)

In many real-world situations, there are constraints on the ways in which a physical system can be manipulated. We investigate the entropy production (EP) and extractable work involved in bringing a system from some initial distribution p to some final distribution p' , given that the set of master equations available to the driving protocol obeys some constraints. We first derive general bounds on EP and extractable work, as well as a decomposition of the nonequilibrium free energy into an “accessible free energy” (which can be extracted as work, given a set of constraints) and an “inaccessible free energy” (which must be dissipated as EP). In a similar vein, we consider the thermodynamics of information in the presence of constraints and decompose the information acquired in a measurement into “accessible” and “inaccessible” components. This decomposition allows us to consider the thermodynamic efficiency of different measurements of the same system, given a set of constraints. We use our framework to analyze protocols subject to symmetry, modularity, and coarse-grained constraints and consider various examples including the Szilard box, the 2D Ising model, and a multiparticle flashing ratchet.

DOI: [10.1103/PhysRevX.11.041024](https://doi.org/10.1103/PhysRevX.11.041024)

Subject Areas: Statistical Physics

I. INTRODUCTION

A. Background

One of the foundational issues in thermodynamics is quantifying how much work is required to transform a system between two thermodynamic states. Recent results in statistical physics derive general bounds on work which hold even for transformations between nonequilibrium states [1,2]. In particular, suppose one wishes to transform a system with initial distribution p and energy function E to some final distribution p' and energy function E' . For an *isothermal* process, during which the system remains in contact with a single heat bath at inverse temperature β , the work extracted during this transformation obeys

$$W(p \rightarrow p') \leq F_E(p) - F_{E'}(p'), \quad (1)$$

where $F_E(p) := \langle E \rangle_p - S(p)/\beta$ is the (*nonequilibrium*) *free energy* of distribution p given energy function E [1–3]. This inequality comes from the second law of thermodynamics,

which states that *entropy production* (EP), the total increase of the entropy of the system and all coupled reservoirs, is non-negative. For an isothermal process that carries out the transformation $p \rightarrow p'$, EP is given by

$$\Sigma(p \rightarrow p') = \beta[F_E(p) - F_{E'}(p') - W(p \rightarrow p')] \geq 0. \quad (2)$$

Equation (1) follows from Eq. (2) by a simple rearrangement.

To extract work from a system, one must manipulate the system by applying a driving protocol. There are many different driving protocols that can be used to transform some initial distribution p to some final distribution p' , which generally incur different amounts of EP and work. Achieving the fundamental bounds set by the second law, such as Eq. (1), typically requires idealized protocols, which make use of arbitrary energy functions, infinite timescales, etc. In many real-world scenarios, however, there are strong practical constraints on how one can manipulate a system, and such idealized protocols are unavailable.

The goal of this paper is to derive stronger bounds on EP and work involved in carrying out the transformation $p \rightarrow p'$, given constraints on the set of master equations available to the driving protocol. Ultimately, such stronger bounds on EP and work can provide new insights into various real-world thermodynamic processes and work-harvesting devices, ranging from biological organisms to artificial engines. They can also cast new light on some well-studied scenarios in statistical physics.

^{*}artemyk@gmail.com

[†]Also at Complexity Science Hub, Vienna; Arizona State University, Tempe, Arizona, USA; <http://davidwolpert.weebly.com>

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

For example, consider a two-dimensional Szilard box connected to a heat bath [4], which contains a single Brownian particle and a vertical partition, and suppose that the driving protocols can manipulate the horizontal position of this partition. Imagine that the particle is initially located in the *left half* of the box. How much work can be extracted by transforming this initial distribution to a uniform final distribution, assuming the system begins and ends with a uniform energy function? A simple application of Eq. (1) shows that the extractable work is upper bounded by $(\ln 2)/\beta$. This bound can be achieved by quickly moving the vertical partition to the middle of the box and then slowly expanding it rightward. Now imagine an alternative scenario, in which the particle is initially located in the *top half* of the box. By Eq. (1), the work that can be extracted by bringing this initial distribution to a uniform final distribution is again upper bounded by $(\ln 2)/\beta$. Intuitively, however, it seems that this bound should not be achievable, given the constrained set of available protocols (i.e., one can manipulate the system only by moving the vertical partition left and right). Our results make this intuition rigorous for the two-dimensional Szilard box, as well as various other systems that can be manipulated only by a constrained set of driving protocols.

This phenomenon also occurs when the starting and ending distributions can depend on the outcome of a measurement of the system. This kind of setup, which was first used to analyze the thermodynamics of information in various kinds of Maxwellian demons, is sometimes called “feedback control” in the literature [2,14]. Imagine that the state of some system X is first measured using some observation channel (i.e., a conditional distribution) $q(m|x)$, producing measurement outcome m with probability $p(m) = \sum_x p(x)q(m|x)$. The system then undergoes a driving protocol which can depend on m . For simplicity, we assume that the system’s energy function begins as E and ends as E' for all measurement outcomes. Let $p_{X|m}$ and $p'_{X'|m}$ indicate the system’s initial and final conditional distributions, respectively, given measurement outcome m , and let $p(x) = \sum_m p(m)p_{X|m}(x|m)$ and $p'(x') = \sum_m p(m)p'_{X'|m}(x'|m)$ indicate the system’s initial and final marginal distributions, respectively [for simplicity, we often use notation like p instead of $p(x)$ below]. We can then take expectations of both sides of Eq. (1) across measurement outcomes, thereby bounding the average extractable work as [15]

$$\langle W \rangle \leq \sum_m p(m) [F_E(p_{X|m}) - F_{E'}(p'_{X'|m})]. \quad (3)$$

By adding and subtracting $[S(p) - S(p')]/\beta$ on the right-hand side, we can further rewrite Eq. (3) in terms of the drop of the free energy in the marginal distribution, plus the loss of information between the measurement and the system over the course of the protocol:

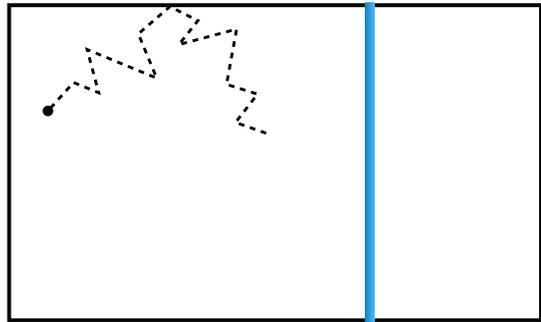


FIG. 1. A two-dimensional Szilard box with a single Brownian particle, where a vertical partition (blue) can be positioned at different horizontal locations in the box. We demonstrate that only information about the particle’s horizontal position, not its vertical position, can be used to extract work from the system.

$$\langle W \rangle \leq F_E(p) - F_{E'}(p') + [I(X;M) - I(X';M)]/\beta, \quad (4)$$

where $I(X;M)$ and $I(X';M)$ indicate the mutual information under the conditional distributions $p_{X|m}$ and $p'_{X'|m}$, respectively. Comparing Eqs. (1) and (4), the bound on average extractable work increases with the drop of mutual information. Equation (4) is a classic result from the “thermodynamics of information” [2,14], which shows that information about the state of a system can be used to increase the work extracted from this system.

Just like Eq. (1), the bound in Eq. (4) is typically saturated by idealized protocols, which have access to arbitrary energy functions, infinite timescales, etc. As mentioned above, in the real world, there are typically constraints on the available protocols, in which case the bound of Eq. (4) may not be achievable. For example, consider again the Szilard box shown in Fig. 1. Imagine measuring a bit of information about the location of the particle and then using this information to extract work from the system while driving it back to a uniform equilibrium distribution. In this case, $I(X;M) = \ln 2$ and $I(X';M) = 0$; therefore, if the system starts and ends with the uniform energy function, Eq. (4) states that $\langle W \rangle \leq (\ln 2)/\beta$. Intuitively, however, it seems that measuring the particle’s horizontal position should be useful for extracting work from the system, while measuring the particle’s vertical position should not be useful. The general bound of Eq. (4) does not distinguish between these two kinds of measurement. In fact, this bound depends only on the overall *amount* of information acquired by the measurement [as quantified by $I(X;M)$] and is completely insensitive to the *content* of that information [i.e., the particular pattern of correlations quantified by $I(X;M)$].

B. Summary of results and road map

In this paper, we derive bounds on extractable work and EP which arise when carrying out the transformation $p \rightarrow p'$ under constraints on the driving protocol. We consider a

system coupled to a single heat bath which undergoes a driving protocol over some time interval $t \in [0, 1]$ (where the units of time are arbitrary). A driving protocol is represented as a continuous-time master equation $L(t)$, where $L(t)$ refers to the (*infinitesimal*) generator at time t . For example, a driving protocol could be a trajectory of time-dependent discrete-state rate matrices or a trajectory of time-dependent Fokker-Planck operators for a continuous-state system.

We say that a driving protocol is *constrained* if there is some restricted set of generators Λ such that $L(t) \in \Lambda$ at all times $t \in [0, 1]$. As discussed below, the particular choice of Λ depends on the specific constraints being considered. For example, Λ might represent a set of generators that are invariant under some particular symmetry group (e.g., representing the dynamics of a set of indistinguishable particles or a spin system on a lattice with symmetries).

Our analysis proceeds at three different “levels” of generality, which we summarize in the following subsections.

1. Level 1: General mathematical framework

In the first level of analysis, presented in Secs. III and IV, we provide a general mathematical framework for deriving bounds on EP and work for constrained driving protocols.

To develop our framework, given some set of allowed generators Λ , we consider an associated operator ϕ over distributions which satisfies two conditions: It obeys the so-called *Pythagorean identity* from information geometry, and it commutes with the dynamics generated by elements of Λ [Eqs. (14) and (16) below]. In Sec. III, we show that, for any distribution p , the distribution $\phi(p)$ contains only that part of the free energy in p which may be turned into work by a constrained driving protocol. Formally, we decompose the nonequilibrium free energy of distribution p and energy function E as

$$F_E(p) = F_E[\phi(p)] + D[p\|\phi(p)]/\beta, \quad (5)$$

where $D(\cdot\|\cdot)$ indicates the Kullback-Leibler divergence. Then, for any constrained driving protocol that carries out the transformation $p \rightarrow p'$, the extractable work is bounded as

$$W(p \rightarrow p') \leq F_E[\phi(p)] - F_{E'}[\phi(p')]. \quad (6)$$

We also demonstrate that EP can be lower bounded by the contraction of the Kullback-Leibler (KL) divergence between p and $\phi(p)$ over the course of the protocol:

$$\Sigma(p \rightarrow p') \geq D[p\|\phi(p)] - D[p'\|\phi(p')]. \quad (7)$$

Given these bounds, it can be seen that Eq. (5) decomposes the nonequilibrium free energy $F_E(p)$ into two terms: an *accessible free energy* $F_E[\phi(p)]$, whose decrease over the course of the protocol may be extractable as work,

and an *inaccessible free energy* $D[p\|\phi(p)]/\beta$, whose decrease over the course of the protocol cannot be turned into work and must be dissipated as EP. The accessible free energy is always less than the overall free energy, $F_E[\phi(p)] \leq F_E(p)$, which follows from Eq. (5) and the non-negativity of KL divergence. We also show that the right-hand side of Eq. (7) is non-negative:

$$D[p\|\phi(p)] - D[p'\|\phi(p')] \geq 0, \quad (8)$$

which implies that our bounds on EP and work [Eqs. (6) and (7), respectively] are stronger than the general bounds provided by the second law [$\Sigma \geq 0$ and Eq. (1)]. Note that Eq. (8) also implies an irreversibility condition on the dynamics: For any two distributions p and p' , a constrained driving protocol can carry out either the transformation $p \rightarrow p'$ or the transformation $p' \rightarrow p$ but not both—unless $D[p\|\phi(p)] = D[p'\|\phi(p')]$.

In Sec. IV, we show that the general framework summarized above has important implications for thermodynamics of information. We consider the type of feedback-control setup discussed above: An observation apparatus first makes a measurement m of the system, and then the system undergoes a driving protocol (which can depend on m) that carries out the transformation $p_{X|m} \rightarrow p'_{X|m}$. Suppose that the driving protocols corresponding to all m obey bounds like Eq. (6) for the same operator ϕ . This operator then gives rise to the “mapped” initial and final conditional distributions $\phi(p_{X|m})$ and $\phi(p'_{X|m})$, respectively. We can then bound average extractable work for feedback control under constraints as

$$\langle W \rangle \leq F_E(p) - F_{E'}(p') + [I_{\text{acc}}^\phi(X; M) - I_{\text{acc}}^\phi(X'; M)]/\beta,$$

where the *accessible information* component of the initial mutual information $I(X; M)$ is defined as

$$I_{\text{acc}}^\phi(X; M) = I(X; M) - D[p_{X|M}\|\phi(p_{X|M})], \quad (9)$$

and similarly for $I_{\text{acc}}^\phi(X'; M)$. This bound is a refinement of Eq. (4) in the presence of protocol constraints, which shows that the amount of extractable work depends on the accessible information $I_{\text{acc}}^\phi(X; M)$ rather than the actual mutual information $I(X; M)$. Loosely speaking, the accessible information reflects the “alignment” between the choice of measured observable and the way the system can be manipulated, given some protocol constraints. Thus, in the presence of constraints, the thermodynamic value of information depends on not only the amount of measured information, but also the content of that information [16,17]. (See also Ref. [18] for a popular discussion of some related issues.)

It is important to note that, at this general level of analysis, we do not describe how to construct the operator

ϕ , as this construction typically depends on the structure of the set Λ . However, as described in the following subsection, we do provide explicit expressions for ϕ for three broad classes of protocol constraints, which we term symmetry, modularity, and coarse-grained constraints.

2. Level 2: Symmetry, modularity, and coarse-grained constraints

At the second level of our analysis, we apply the general framework described above to derive bounds on EP and work for three broad classes of protocol constraints.

- (i) Section V considers *symmetry constraints*, when the available generators possess some symmetry group. Examples of systems with symmetry constraints include the Szilard box in Fig. 1, spin systems on lattices, and gases of indistinguishable particles. The operator ϕ corresponding to symmetry constraints, defined in Eq. (42), maps distributions to their “symmetrized” versions (which are invariant under the action of the symmetry group).
- (ii) Section VI considers *modularity constraints*, when the available generators cause different (though possibly overlapping) subsystems of a multivariate system to evolve independently of each other. Examples of systems with modularity constraints include digital circuits [19], ideal gases, and multi-particle Maxwellian demons. The operator ϕ corresponding to modularity constraints, defined in Eq. (64), maps distributions to their “uncorrelated” versions, without statistical dependencies between independent subsystems.
- (iii) Section VII considers *coarse-grained constraints*, when the available generators exhibit closed coarse-grained dynamics which obey some constraints (e.g., coarse-grained symmetry or modularity constraints). An example is provided by the Szilard box in Fig. 1: The particle’s vertical position (the coarse-grained macrostate) evolves in a way that does not depend on the horizontal position, and the macrostate equilibrium distribution cannot be controlled by moving the partition. Given a protocol that obeys coarse-grained constraints, we show that the EP can be lower bounded in terms of a “coarse-grained EP” [Eqs. (87)–(89)] and that this coarse-grained EP can itself be lower bounded by a coarse-grained version of Eq. (7).

In addition, we also discuss how tighter bounds on work and EP can be derived by combining different kinds of constraints (e.g., when a system obeys two different symmetry groups or when it obeys both symmetry and modularity constraints).

3. Level 3: Concrete examples

At the third (and most concrete) level, we illustrate our results for symmetry, modularity, and coarse-grained constraints on several example systems.

- (i) In Sec. VA, we use symmetry constraints to derive thermodynamic bounds for the Szilard box in Fig. 1, which possesses vertical reflection symmetry.
- (ii) In Sec. VB, we use symmetry constraints to derive thermodynamic bounds for the Ising model on a 2D lattice, which possesses translational symmetry.
- (iii) In Sec. VIA, we use modularity constraints to derive thermodynamic bounds for the Szilard box in Fig. 1, which are different from the bounds derived in Sec. VA. We also demonstrate that stronger results can be derived by combining bounds arising from symmetry and modularity constraints.
- (iv) In Secs. VIB and VIC, we use modularity constraints to derive bounds on work extraction for two multi-particle feedback-control protocols that have been proposed in the literature: a multiparticle Szilard box [20] and a collective flashing ratchet [21].
- (v) In Sec. VIIA, we use coarse-grained constraints to derive thermodynamic bounds for a version of the Szilard box in Fig. 1 in the presence of gravity. We also demonstrate that stronger results can be derived by combining bounds arising from coarse-grained and modularity constraints.

4. Literature review and discussion

After presenting the results summarized above, in Sec. VIII we discuss related prior literature. We also compare and contrast our results, such as the decomposition of nonequilibrium free energy in Eq. (5), to some relevant work in quantum thermodynamics [22,23]. We conclude with a brief discussion in Sec. IX, which also touches upon how our approach generalizes beyond the assumption of a single heat bath. Proofs and derivations are in the Appendixes.

II. PRELIMINARIES

We consider a physical system with state space X , which can be either discrete or continuous ($X = \mathbb{R}^n$). The term “probability distribution” refers to a probability mass function over X in the discrete case and to a probability density function over X in the continuous case. We interchangeably use notation like $p(x)$ and p_x (as is clear from context) to indicate the probability of state x . We use \mathcal{P} to refer to the set of all probability distributions over X .

The system evolves in a stochastic manner during a driving protocol over time $t \in [0, 1]$. We write $p(t)$ to indicate the distribution at time t corresponding to some initial distribution $p(0) = p$, and $p(1) = p'$ to indicate the distribution at the end of the protocol. For a discrete-state system, the distribution at time t evolves according to the time-dependent master equation

$$\partial_t p_x(t) = \sum_{x'} [L_{xx'}(t)p_{x'}(t) - L_{x'x}(t)p_x(t)], \quad (10)$$

where $L_{x'x}(t)$ is the transition rate from state x to state x' . We assume that the system is coupled to a heat bath at inverse temperature β , and so each $L(t)$ obeys local detailed balance (see Sec. IX for a generalization of this assumption). Formally, this assumption means that $\pi_{x'}^{L(t)} L_{xx'}(t) = \pi_x^{L(t)} L_{x'x}(t)$ for all x, x' , and t , where $\pi^{L(t)}$ is the stationary distribution of rate matrix $L(t)$, which we assume is unique (though this latter assumption can be relaxed [24]).

The rate of entropy production (EP rate) incurred at time t can be written as [Eq. (33) in Ref. [25]]

$$\dot{\Sigma}[p(t), L(t)] = -\sum_x \partial_t p_x(t) \ln \frac{p_x(t)}{\pi_x^{L(t)}} \geq 0, \quad (11)$$

where $\partial_t p_x(t)$ is defined in Eq. (10). Note that the right side of Eq. (11) is sometimes called the “nonadiabatic EP rate” in stochastic thermodynamics, and it is equal to the overall EP rate for a system coupled to a single bath and obeying detailed balance [25]. The total EP incurred by a time-extended protocol over $t \in [0, 1]$ that carries out the transformation $p \rightarrow p'$ is given by the integral of the EP rate:

$$\Sigma(p \rightarrow p') = \int_0^1 \dot{\Sigma}[p(t), L(t)] dt. \quad (12)$$

The work extracted during a protocol can be calculated by using Eqs. (2) and (12), once the initial and final non-equilibrium free energies $F_E(p)$ and $F_{E'}(p')$ are specified. To define these free energies, we assume that there is some fixed pair of energy functions E and E' , which specify the Boltzmann equilibrium distributions of $L(0)$ and $L(1)$, respectively.

For a continuous-state system evolving under a continuous master equation [26,27], the sums in Eqs. (10) and (11) should be replaced by integrals [see Eq. (31) in Ref. [28]]. A prototypical example of a continuous master equation, which we use below, is a Fokker-Planck equation [26,29]:

$$\partial_t p(x, t) = -\nabla \cdot [\mathbf{A}(x, t)p(x, t) - \mathbf{D}(x, t)\nabla p(x, t)], \quad (13)$$

where \mathbf{A} and \mathbf{D} are drift and diffusion terms.

We often write dynamical equations like Eqs. (10) and (13) using the notation $\partial_t p(t) = L(t)p(t)$, where $L(t)$ is a bounded linear operator that is called the (*infinitesimal*) *generator* of the dynamics at time t . Note that, for a continuous-state system in phase space, it may be that the system is isolated from the bath for some $t \in [0, 1]$, in which case $\partial_t p(t) = L(t)p(t)$ should be understood in terms of the Liouville equation (for example, if a system is first isolated and evolves in a Hamiltonian manner and is then brought in contact with a bath at inverse temperature β and allowed to equilibrate).

III. GENERAL FRAMEWORK

We begin by presenting our general mathematical framework. The application of this framework to concrete situations is described in later sections.

A driving protocol $\{L(t) : t \in [0, 1]\}$ is said to be *constrained* if there is some restricted set of generators Λ such that $L(t) \in \Lambda$ at all t . For a given set of allowed generators Λ , we consider an associated operator $\phi : \mathcal{P} \rightarrow \mathcal{P}$ which satisfies two conditions. The first condition states that

$$D(p||q) = D[p||\phi(p)] + D[\phi(p)||q] \quad (14)$$

for all $p \in \mathcal{P}$ and $q \in \text{img } \phi$ with $D(p||q) < \infty$ [where $\text{img } \phi = \{\phi(p) : p \in \mathcal{P}\}$ is the image of the operator ϕ]. Equation (14) is sometimes called the *Pythagorean identity of KL divergence* in information geometry [30]. Any ϕ that obeys Eq. (14) can be written in terms of the following projection [31]:

$$\phi(p) = \arg \min_{q \in \text{img } \phi} D(p||q), \quad (15)$$

which shows that $D[p||\phi(p)]$ is the minimal information-theoretic distance from p to the set of distributions $\text{img } \phi$.

The second condition is that ϕ obeys the following *commutativity relation* for all $L \in \Lambda$:

$$e^{\tau L} \phi(p) = \phi(e^{\tau L} p) \quad \forall \tau \geq 0, \quad p \in \mathcal{P}. \quad (16)$$

In other words, given any initial distribution p , the same final distribution is reached regardless of whether p first relaxes under L for time τ and then undergoes ϕ or instead first undergoes ϕ and then relaxes under L for time τ .

Note that the Pythagorean identity in Eq. (14) concerns only the operator ϕ , while the commutativity relation in Eq. (16) concerns the relationship between ϕ and the generators in Λ [and, therefore, all of the generators $L(t)$ in the driving protocol, since $L(t) \in \Lambda$ at all t by assumption]. Beyond these two conditions, the operator ϕ can be arbitrary and may be linear or nonlinear. In the following sections of this paper, we show how to choose ϕ for various types of constrained protocols.

Importantly, any ϕ that satisfies the two conditions above maps any distribution p to a corresponding “accessible” distribution $\phi(p)$, which controls the amount of work that can be extracted from p by a constrained driving protocol. To prove this result, we first show that for any $L \in \Lambda$ that obeys Eq. (16), the equilibrium distribution π^L satisfies (Lemma 1 in Appendix A)

$$\pi^L \in \text{img } \phi. \quad (17)$$

We also derive the following mathematical result, central to much of what follows: If ϕ obeys Eqs. (14) and (16) for some generator L , then the EP rate incurred by any

distribution p under L can be written as the sum of two non-negative terms—the EP rate incurred by $\phi(p)$ under L and the instantaneous contraction of the KL divergence between p and $\phi(p)$.

Theorem 1.—If ϕ obeys Eqs. (14) and (16) for some generator L , then, for all $p \in \mathcal{P}$,

$$\dot{\Sigma}(p, L) = \dot{\Sigma}[\phi(p), L] - \frac{d}{dt} D\{p(t) \parallel \phi[p(t)]\},$$

and $-(d/dt)D\{p(t) \parallel \phi[p(t)]\} \geq 0$, where $\partial_t p(t) = Lp$.

We sketch the proof of this theorem in terms of a discrete-time relaxation over interval τ , as shown in Fig. 2 (see Appendix A for details). Consider some distribution p that relaxes for time τ under the generator L , thereby reaching the distribution $e^{\tau L}p$ (solid gray line). The EP incurred by this relaxation is given by the contraction of KL divergence to the equilibrium distribution π : $\Sigma(p \rightarrow e^{\tau L}p) = D(p \parallel \pi) - D(e^{\tau L}p \parallel \pi)$ (contraction of purple lines) [25,28]. Given Eq. (17), we can apply the Pythagorean identity [Eq. (14)] to both $D(p \parallel \pi)$ and $D(e^{\tau L}p \parallel \pi)$, which lets us rewrite $\Sigma(p \rightarrow e^{\tau L}p)$ as the sum of two terms: $D[p \parallel \phi(p)] - D[e^{\tau L}p \parallel \phi(e^{\tau L}p)]$ (green lines) and $D[\phi(p) \parallel \pi] - D[\phi(e^{\tau L}p) \parallel \pi]$ (red lines). Applying the commutativity relation [Eq. (16)] shows that the first term is non-negative by the data-processing inequality and that the second term is equal to $\Sigma[\phi(p) \rightarrow e^{\tau L}\phi(p)]$, the EP incurred by letting $\phi(p)$ relax freely under L . The continuous-time statement found in Theorem 1 follows by taking the appropriate $\tau \rightarrow 0$ limit, while noting that the EP rate [Eq. (11)] can be rewritten in terms of the limit $\lim_{\tau \rightarrow 0} (1/\tau)[D(p \parallel \pi) - D(e^{\tau L}p \parallel \pi)]$.

Now suppose that Eq. (16) holds, so that the assumptions of Theorem 1 are satisfied during the entire protocol. In that

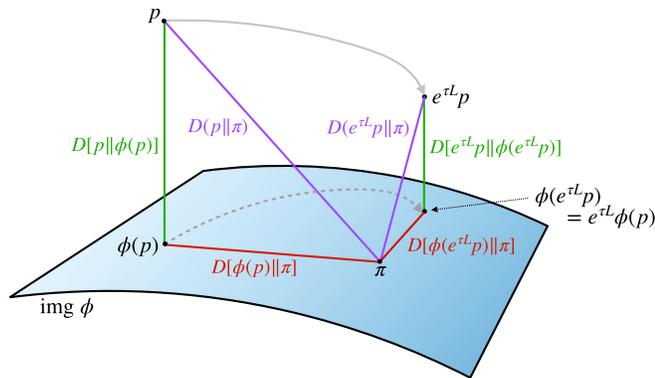


FIG. 2. Visual explanation of Theorem 1: Distribution p freely relaxes under L for time τ (solid gray line). The EP incurred during this relaxation (contraction of purple lines) can be decomposed into the contraction of the KL divergence between p and $\phi(p)$ (contraction of green lines) plus the EP incurred during the free relaxation of $\phi(p)$ (contraction of the red lines). The free relaxation of $\phi(p)$ under L is represented by the dotted gray line.

case, as we show in Lemma 3 in Appendix A, any constrained protocol that carries out the transformation $p \rightarrow p'$ must also transform the initial distribution $\phi(p)$ to the final distribution $\phi(p')$. We can then, in essence, integrate Theorem 1 over time and derive the following result about total EP.

Theorem 2.—If ϕ obeys Eqs. (14) and (16) for all $L \in \Lambda$, then, for any constrained protocol that transforms $p \rightarrow p'$,

$$\Sigma(p \rightarrow p') = \Sigma[\phi(p) \rightarrow \phi(p')] + \{D[p \parallel \phi(p)] - D[p' \parallel \phi(p')]\}$$

and $D[p \parallel \phi(p)] - D[p' \parallel \phi(p')] \geq 0$.

We use Theorem 2 to derive several useful bounds on EP and work. First, since $\Sigma[\phi(p) \rightarrow \phi(p')] \geq 0$ by the non-negativity of EP, the contraction of KL divergence between p and $\phi(p)$ bounds the EP incurred by a constrained driving protocol that carries out the transformation $p \rightarrow p'$:

$$\Sigma(p \rightarrow p') \geq D[p \parallel \phi(p)] - D[p' \parallel \phi(p')] \geq 0, \quad (18)$$

which appeared as Eq. (7) in the introduction. Furthermore, $D[p \parallel \phi(p)] - D[p' \parallel \phi(p')] \geq 0$ immediately implies that

$$\Sigma(p \rightarrow p') \geq \Sigma[\phi(p) \rightarrow \phi(p')]. \quad (19)$$

We can also derive the decomposition of free energy and the bound on extractable work, which appeared as Eqs. (5) and (6) in the introduction. Consider some transformation $p \rightarrow p'$, and write the initial nonequilibrium free energy as

$$F_E(p) = F_E(\pi) + D(p \parallel \pi)/\beta, \quad (20)$$

where $\pi \propto e^{-\beta E}$ is the Boltzmann distribution for the initial energy function E and $F_E(\pi)$ is the equilibrium free energy [3]. Using Eq. (17) and the Pythagorean identity [Eq. (14)], we decompose the nonequilibrium free energy into a sum of the accessible free energy and the inaccessible free energy:

$$\begin{aligned} F_E(p) &= F_E(\pi) + \{D[p \parallel \phi(p)] + D[\phi(p) \parallel \pi]\}/\beta \\ &= F_E[\phi(p)] + D[p \parallel \phi(p)]/\beta. \end{aligned} \quad (21)$$

Using a similar derivation, we can write the nonequilibrium free energy at the end of the protocol as

$$F_E(p') = F_E[\phi(p')] + D[p' \parallel \phi(p')]/\beta. \quad (22)$$

Subtracting Eq. (22) from Eq. (21) shows that the drop in the nonequilibrium free energy during $p \rightarrow p'$ is given by

$$\begin{aligned} F_E(p) - F_E(p') &= F_E[\phi(p)] - F_E[\phi(p')] \\ &\quad + \{D[p \parallel \phi(p)] - D[p' \parallel \phi(p')]\}/\beta. \end{aligned} \quad (23)$$

Combining this result with Theorem 2 and Eq. (2), and then rearranging, shows that the work involved in carrying out $p \rightarrow p'$ is equal to the work involved in carrying out the “accessible” transformation $\phi(p) \rightarrow \phi(p')$:

$$W(p \rightarrow p') = W[\phi(p) \rightarrow \phi(p')]. \quad (24)$$

Finally, by combining with Eq. (1), we arrive at an upper bound on work that can be extracted by a constrained protocol:

$$W(p \rightarrow p') \leq F_E[\phi(p)] - F_E'[\phi(p')], \quad (25)$$

which is tighter than the bound given by the second law [Eq. (1)].

The bounds in Eqs. (18) and (25), as well as the decomposition of free energy in Eq. (21), are the main theoretical results arising from our general framework. Figure 3 provides a schematic way of understanding these results. Theorem 2 states that, for a constrained protocol that carries out the map $p \rightarrow p'$, the EP incurred during the system’s actual trajectory (solid gray line) is given by the EP that would be incurred by a “projected trajectory” that carries out the transformation $\phi(p) \rightarrow \phi(p')$ (dashed gray line), plus the drop in the KL divergence from the system’s distribution to the set $\text{img } \phi$ over the course of the protocol (contraction of green lines). Since the EP of the projected trajectory must be non-negative, the drop in the distance from the system’s distribution to $\text{img } \phi$ serves as a lower bound on EP, as in Eq. (18). In addition, Theorem 2 states that this decrease in the KL divergence must be positive, meaning that the system’s distribution must get closer to $\text{img } \phi$ over the course of the protocol.

Following Fig. 3, it can be helpful to think of the trajectory $p \rightarrow p'$ as composed of three segments: (i) from p down to $\phi(p)$, (ii) from $\phi(p)$ to $\phi(p')$ while staying

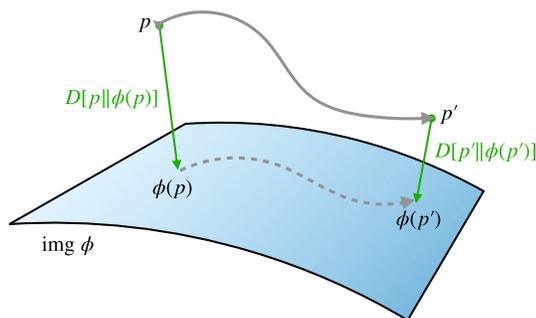


FIG. 3. Illustration of Theorem 2. Given an appropriate operator ϕ , $\Sigma(p \rightarrow p')$ (the EP incurred during some desired transformation $p \rightarrow p'$; solid gray line) is equal to $\Sigma[\phi(p) \rightarrow \phi(p')]$ [the EP incurred by that protocol when transforming $\phi(p) \rightarrow \phi(p')$; dashed gray line] plus the contraction of the KL divergence $D[p\|\phi(p)] - D[p'\|\phi(p')]$ (contraction of green lines). This contraction of KL divergence is a non-negative lower bound on $\Sigma(p \rightarrow p')$, as in Eq. (18).

within $\text{img } \phi$, and (iii) from $\phi(p')$ up to p' (note that this decomposition is useful for accounting purposes but does not generally reflect the actual trajectory the system takes in going from p to p'). The first and third segments contribute (positively and negatively, respectively) only to EP, while the projected second segment $\phi(p) \rightarrow \phi(p')$ contributes both to EP and to work. Thus, the work involved in $p \rightarrow p'$ is determined entirely by the work involved in the second segment, as stated in Eq. (24).

Note also the formal similarity between our decomposition of the drop in free energy [Eq. (23)] and the decompositions of EP in Theorem 2. Indeed, like Theorem 2, the result Eq. (23) can be illustrated with Fig. 3: During the transformation $p \rightarrow p'$ (solid gray line), the drop in free energy is given by the drop in free energy incurred by the transformation $\phi(p) \rightarrow \phi(p')$ (dotted gray line), plus the contraction of the KL divergence from the system’s distribution to the set $\text{img } \phi$ (green lines).

In general, our bounds on EP and work are not always achievable. Suppose, however, that the final distribution p' is in equilibrium, so $p' = \phi(p')$ by Eq. (17). Equation (18) then gives

$$\Sigma(p \rightarrow p') \geq D[p\|\phi(p)]. \quad (26)$$

This bound is achievable if the generators in Λ have a continuous curve of equilibrium distributions from $\phi(p)$ to $p' = \phi(p')$. Imagine a protocol in which the initial distribution p first relaxes to the equilibrium distribution $\phi(p)$ and then undergoes quasistatic driving from $\phi(p)$ to $\phi(p')$ while remaining in equilibrium throughout [in terms of Fig. 3, the system first relaxes along the green arrow connecting p to $\phi(p)$ and then follows the dashed line to $\phi(p')$ quasistatically]. The relaxation step incurs $D[p\|\phi(p)]$ of EP, while the quasistatic step incurs a vanishing amount of EP, so the bound in Eq. (26) is achieved.

A. Choice of the ϕ operator

In general, the operator ϕ associated with a given set of generators Λ is not unique. For instance, for any driving protocol, the identity map $\phi(p) = p$ always satisfies Eqs. (14) and (16). Choosing ϕ to be the identity map, however, reduces the results in Theorem 2 to trivial identities and the lower bound on EP in Eq. (18) to 0.

At a high level, those ϕ which have smaller $\text{img } \phi$ generally give tighter bounds on EP (since, given Eq. (15), a smaller image leads to larger values of $D[p\|\phi(p)]$). To illustrate this phenomenon, consider the extreme case where all $L \in \Lambda$ have the same equilibrium distribution π , so that any constrained driving protocol must be a free relaxation toward π . Then, the operator $\phi(p) = \pi$ for all p (so $\text{img } \phi$ is a singleton) satisfies Eqs. (14) and (16) and, when plugged into Eq. (18), gives the following bound on EP:

$$\Sigma(p \rightarrow p') \geq D(p\|\pi) - D(p'\|\pi). \quad (27)$$

In fact, the right-hand side is an exact expression for the EP incurred by the free relaxation, meaning that it is the tightest possible bound. If, however, the generators $L \in \Lambda$ have different equilibrium distributions, then the operator $\phi(p) = \pi$ (for whatever π) generally violates the commutativity relation in Eq. (16), and bounds like Eq. (27) no longer hold.

In the following sections, we show how to use our results to derive thermodynamic bounds for Λ that obey some kind of symmetry group, modular decomposition, or coarse-graining structure. In more general (and possibly unstructured) cases, it is an open question of whether a nontrivial operator ϕ exists and, if so, how to identify it. We explore related issues in a companion paper [32], where we use numerical optimization techniques to derive bounds on EP similar to Eq. (18).

Importantly, when there are multiple different operators that all satisfy the Pythagorean identity and the commutativity relation for the available generators Λ , one can derive tighter bounds on EP and work by applying our decompositions in an “iterative” manner. For instance, imagine that there are two different operators ϕ_1 and ϕ_2 that satisfy Esq. (14) and (16) (for example, these might represent operators arising from symmetry constraints and modularity constraints, respectively, as described below). Applying Theorem 2 iteratively leads to “stacked” bounds on EP analogous to Eq. (18):

$$\begin{aligned} \Sigma(p \rightarrow p') \geq & (D[p\|\phi_1(p)] + D\{\phi_1(p)\|\phi_2[\phi_1(p)]\}) \\ & - (D[p'\|\phi_1(p')] + D\{\phi_1(p')\|\phi_2[\phi_1(p')]\}) \geq 0. \end{aligned} \quad (28)$$

Similarly, applying Eq. (24) iteratively leads to stacked bounds on extractable work analogous to Eq. (25):

$$W(p \rightarrow p') \leq F_E\{\phi_2[\phi_1(p)]\} - F_{E'}\{\phi_2[\phi_1(p')]\} \quad (29)$$

Such stacked bounds are generally tighter than the bounds provided by either ϕ_1 or ϕ_2 alone. [Note that one can also reverse the order of operations and consider the composition $\phi_1[\phi_2(p)]$ rather than $\phi_2[\phi_1(p)]$ in Eqs. (28) and (29), which, in general, leads to different bounds.]

B. Fluctuating entropy production

As we show in detail in Appendix A 2, our results also have implications for stochastic fluctuations of trajectory-level EP, as considered in stochastic thermodynamics [33].

Consider any constrained driving protocol over $t \in [0, 1]$ with an associated operator ϕ . Let \mathbf{x} indicate some stochastically sampled trajectory of the system visited during the driving protocol, and let $\sigma_p(\mathbf{x})$ indicate the fluctuating EP incurred by trajectory \mathbf{x} when initial

states are sampled from the initial distribution p . In Appendix A 2, we consider the difference between this fluctuating EP and the fluctuating EP incurred by the same trajectory when initial states are sampled from the accessible initial distribution $\phi(p)$:

$$m_p(\mathbf{x}) := \sigma_p(\mathbf{x}) - \sigma_{\phi(p)}(\mathbf{x}). \quad (30)$$

By combining Theorem 2 with recent results in stochastic thermodynamics [34,35], we show that the expectation of $m_p(\mathbf{x})$ is equal to the difference of expected EPs, $\langle m_p(\mathbf{x}) \rangle = \Sigma(p \rightarrow p') - \Sigma[\phi(p) \rightarrow \phi(p')]$, where $\langle \cdot \rangle$ indicates expectation over trajectories sampled from initial distribution p . We also show that $m_p(\mathbf{x})$ obeys a detailed fluctuation theorem, which implies a trajectory-level version of Eq. (19): The probability that the fluctuating EP under initial distribution p is ξ less than the fluctuation EP under the accessible initial distribution $\phi(p)$ is exponentially small (i.e., it is less than $e^{-\xi}$). We leave further exploration of the connection between our framework and stochastic thermodynamics for future work.

IV. THERMODYNAMICS OF INFORMATION UNDER PROTOCOL CONSTRAINTS

The framework introduced in the previous section has implications for the thermodynamics of information under constraints. Consider the type of feedback-control setup described in the introduction: First, an observation apparatus M measures some system observable, and then the system undergoes a driving protocol that depends on the measurement outcome m . Let $L^{(m)}(t)$ indicate the driving protocol conditioned on m , and let $p_{X|m}$ and $p'_{X'|m}$ indicate the distributions over system states at the beginning and end, respectively, of the corresponding driving protocol. As standard in the literature [2], for simplicity we assume that all protocols start and end with the same energy functions E and E' , that the measurement apparatus M and the system X are energetically decoupled during the protocols, and that M does not change state.

Given the above assumptions, it is straightforward to show that the EP incurred by the joint “supersystem” $X \times M$ obeys

$$\Sigma_{XM} = \sum_m p(m) \Sigma_m, \quad (31)$$

where Σ_m is the EP incurred by protocol $L^{(m)}(t)$ in carrying out the transformation $p_{X|m} \rightarrow p'_{X'|m}$. Similarly, by taking expectations of Eq. (2) and rearranging [see the derivation of Eq. (4)], the average extracted work under feedback control can be written as

$$\langle W \rangle = \Delta F + [I(X; M) - I(X'; M)] - \sum_m p(m) \Sigma_m, \quad (32)$$

where for notational convenience we use $\Delta F = F_E(p) - F_{E'}(p')$ to indicate the drop of marginal free energy. Thus, any lower bounds on Σ_m [the EP values incurred by the individual protocols $L^{(m)}(t)$] can be translated into bounds on the overall EP and average extractable work for a feedback-control setup.

For example, suppose that there is some single set of constraints that applies to all of the driving protocols, in that there is some set of generators Λ such that $L^{(m)}(t) \in \Lambda$ for all t and m , as well as an operator ϕ that obeys the Pythagorean identity [Eq. (14)] and the commutativity relation [Eq. (16)] for all $L \in \Lambda$. In that case, the framework described in Sec. III leads to bounds on each Σ_m term. In particular, using Eqs. (18) and (31) gives the bound

$$\Sigma_{XM} \geq D[p_{X|M} \|\phi(p_{X|M})] - D[p'_{X'|M} \|\phi(p'_{X'|M})] \geq 0, \quad (33)$$

where we define the conditional KL divergence $D[p_{X|M} \|\phi(p_{X|M})] = \sum_m p(m) D[p_{X|m} \|\phi(p_{X|m})]$, and similarly for $D[p'_{X'|M} \|\phi(p'_{X'|M})]$. Plugging into Eq. (32) gives the following bound on average extractable work:

$$\langle W \rangle \leq \Delta F + [I_{\text{acc}}^\phi(X; M) - I_{\text{acc}}^\phi(X'; M)]/\beta, \quad (34)$$

where $I_{\text{acc}}^\phi(X; M)$ is given by

$$I_{\text{acc}}^\phi(X; M) = I(X; M) - D[p_{X|M} \|\phi(p_{X|M})], \quad (35)$$

and similarly for $I_{\text{acc}}^\phi(X'; M)$.

We refer to $I_{\text{acc}}^\phi(X; M)$ as the accessible information in measurement M , since any decrease in accessible information can contribute to work extraction [Eq. (34)]. We refer to the conditional KL divergence $D[p_{X|M} \|\phi(p_{X|M})]$ as the inaccessible information, since any decrease in inaccessible information must be dissipated as EP and not extracted as work [Eq. (33)]. The inaccessible information is non-negative by properties of KL divergence, so $I_{\text{acc}}^\phi(X; M) \leq I(X; M)$. In addition, whenever $p \in \text{img } \phi$ [e.g., when p is an equilibrium distribution, by Eq. (17)], the accessible information can be rewritten in a simpler form as

$$I_{\text{acc}}^\phi(X; M) = D[\phi(p_{X|M}) \| p], \quad (36)$$

as follows from Eq. (35) by writing $I(X; M) = D(p_{X|M} \| p)$ and applying the Pythagorean theorem [Eq. (14)].

In general, measurements of different observables on the same system give rise to different amounts of accessible and inaccessible information. At a high level, one should choose measurements that maximize the accessible information $I_{\text{acc}}^\phi(X; M)$, or alternatively the ‘‘efficiency’’ quantified as bits of accessible information per bit of measured information, $I_{\text{acc}}^\phi(X; M)/I(X; M) \leq 1$. Optimal measurements

satisfy $I_{\text{acc}}^\phi(X; M) = I(X; M)$, which happens when the conditional distributions over system states $p_{X|m}$ are invariant under the action of ϕ [i.e., when $\phi(p_{X|m}) = p_{X|m}$ for each m].

Note that similar results can also be derived using other kinds of bounds on Σ_m [e.g., when the individual protocols obey a combination of constraints, so that Eq. (28) holds].

V. SYMMETRY CONSTRAINTS

We now use the general framework introduced above to derive bounds on EP under symmetry constraints.

Consider a compact group \mathcal{G} that has a measurable action over X , such that each $g \in \mathcal{G}$ is a bijection $X \rightarrow X$ [36]. For continuous X , we assume that each $g \in \mathcal{G}$ is a rigid transformation. For notational convenience, for each $g \in \mathcal{G}$ we define the composition operator Φ_g , so that, for any function $f: X \rightarrow \mathbb{R}$,

$$\Phi_g(f)(x) = f[g(x)]. \quad (37)$$

We say that a set of generators Λ obeys *symmetry constraints* (with respect to the action of group \mathcal{G}) if the following commutativity relation holds for all $L \in \Lambda$:

$$\Phi_g L = L \Phi_g \quad \forall g \in \mathcal{G}. \quad (38)$$

In other words, Λ obey symmetry constraints when, for each $L \in \Lambda$ and $g \in \mathcal{G}$, it does not matter whether one first applies the generator L and then the bijection g over the state space or first applies the bijection g over the state space and then the generator L . In more concrete terms, for a (continuous or discrete) master equation L , Eq. (38) holds if the transition rates are invariant under the action of \mathcal{G} :

$$L_{xx'} = L_{g(x)g(x')} \quad \forall x, x' \in X, \quad g \in \mathcal{G}. \quad (39)$$

We can also derive simple sufficient conditions for potential-driven Fokker-Planck equations of the type

$$Lp = \nabla \cdot (\nabla E_L)p + \beta^{-1} \Delta p, \quad (40)$$

where E_L is the energy function of generator L . Then, Eq. (38) holds if all available energy functions are invariant under the action of \mathcal{G} :

$$E_L(x) = E_L[g(x)] \quad \forall x \in X, \quad g \in \mathcal{G}, \quad L \in \Lambda. \quad (41)$$

[Equation (38) is derived from Eqs. (39) and (41) in Appendix B.]

We now define a linear operator $\phi_{\mathcal{G}}$ which satisfies the Pythagorean identity and the commutativity relation [Eqs. (14) and (16)] for symmetry constraints. Let $\phi_{\mathcal{G}}$ map each $p \in \mathcal{P}$ to its average under the action of \mathcal{G} :

$$\phi_{\mathcal{G}}(p)(x) := \int_{\mathcal{G}} p[g(x)] d\mu(g), \quad (42)$$

where μ is the uniform (normalized Haar) measure over \mathcal{G} [37]. For a finite group, the integral in Eq. (42) should be replaced by a summation. Following the terminology in quantum physics, we sometimes refer to $\phi_{\mathcal{G}}$ as a *twirling operator* [23,39]. Intuitively, $\phi_{\mathcal{G}}(p)$ symmetrizes p , removing all information in p concerning the state of the system along the “coordinates” specified by the symmetry constraints.

In Appendix B, we show that $\phi_{\mathcal{G}}$ obeys the Pythagorean identity and, as long as Eq. (38) holds, the commutativity relation of Eq. (16). Thus, any protocol that carries out the transformation $p \rightarrow p'$ while obeying symmetry constraints with respect to \mathcal{G} permits the decomposition of EP found in Theorem 2, with $\phi = \phi_{\mathcal{G}}$, and satisfies all the bounds on work and EP that follow from that result.

In particular, using Eq. (21), we can decompose the free energy $F_E(p)$ of any distribution p into the accessible free energy $F_E[\phi_{\mathcal{G}}(p)]$, which is the free energy in the twirled (and, therefore, symmetric) version of p , and the inaccessible free energy $D[p\|\phi_{\mathcal{G}}(p)]/\beta$. Note that $D[p\|\phi_{\mathcal{G}}(p)]$ is a non-negative measure of the asymmetry in distribution p with respect to the symmetry group \mathcal{G} , which vanishes when p is invariant under $\phi_{\mathcal{G}}$. Thus, for any protocol that obeys symmetry constraints, the first inequality in Eq. (18) states that any “drop in asymmetry” must be dissipated as EP and not turned into work. The second inequality in Eq. (18) states that the asymmetry in the system’s distribution can only decrease during the protocol. (Some of the above results for symmetry constraints have been previously uncovered in quantum thermodynamics [22,23]; see Sec. VIII.)

We finish by discussing thermodynamics of information under symmetry constraints. In general, the results derived in Sec. IV apply to the twirling operator $\phi_{\mathcal{G}}$ as a special case. We can also exploit special properties of $\phi_{\mathcal{G}}$ to further simplify the expression of the inaccessible information term in Eqs. (33) and (35). Suppose that distribution p is invariant under $\phi_{\mathcal{G}}$, so $p = \phi_{\mathcal{G}}(p)$ (e.g., if p is an equilibrium distribution). As shown in Appendix B 4, we can then rewrite the inaccessible information term as

$$D[p_{X|M}\|\phi_{\mathcal{G}}(p_{X|M})] = \left\langle \ln \frac{q(m|x)}{\int_{\mathcal{G}} q[m|g(x)] d\mu(g)} \right\rangle, \quad (43)$$

where $q(m|x)$ is the measurement channel and $\langle \cdot \rangle$ indicates expectation under the joint distribution $p(x, m) = p(x)q(x|m)$. Equation (43) conveniently expresses the inaccessible information in terms of the asymmetry of the measurement channel relative to the action of \mathcal{G} [the right side of Eq. (43) vanishes when $q(m|x)$ is invariant under that action], which we exploit in some of our examples below.

A. Example: Szilard box with symmetry constraints

We demonstrate our results on symmetry constraints using the Szilard box shown in Fig. 1. We assume that the box is coupled to a single heat bath at inverse temperature $\beta = 1$ and that the particle inside the box has overdamped Fokker-Planck dynamics, so that all generators have the form of Eq. (40). The system’s state is represented by one horizontal and a vertical coordinate, $x = (x_1, x_2) \in \mathbb{R}^2$.

Suppose that all energy functions have the form

$$E_{\lambda}(x_1, x_2) = V_p(x_1 - \lambda) + V_w(|x_1|) + V_w(|x_2|), \quad (44)$$

where $\lambda \in \mathbb{R}$ is a controllable parameter that determines the location of the vertical partition, V_p is the partition’s repulsion potential, and V_w is the repulsion potential of the box walls:

$$V_w(a) = \begin{cases} 0 & \text{if } a \leq 1, \\ \infty & \text{otherwise,} \end{cases} \quad (45)$$

meaning that the box extends over $(x_1, x_2) \in [-1, 1]^2$ [40]. Assume that $V_p(x_1 - \lambda) = 0$ for some value of λ and all $x_1 \in [-1, 1]$ (i.e., the partition can be removed by setting λ outside the box). For such λ , let E_{\emptyset} indicate the corresponding energy function, and note that it obeys $E_{\emptyset}(x_1, x_2) = 0$ within the box (and infinity elsewhere), corresponding to a uniform equilibrium distribution $u(x_1, x_2) = \mathbf{1}_{[-1,1]^2}(x_1, x_2)/4$ (where $\mathbf{1}$ is the indicator function). This Szilard box is shown schematically in Fig. 4.

The energy functions in Eq. (44) obey the vertical reflection symmetry $E(x_1, x_2) = E(x_1, -x_2)$, corresponding to the two-element symmetric group S_2 whose action is generated by $g(x_1, x_2) = (x_1, -x_2)$. The corresponding twirling of p is the uniform mixture of p and its reflection:

$$\phi_{\mathcal{G}}(p)(x_1, x_2) = [p(x_1, x_2) + p(x_1, -x_2)]/2. \quad (46)$$

We can use our results to derive bounds on the work that can be extracted from this Szilard box. Intuitively, the set of allowed generators L —that is, Fokker-Planck operators with energy functions as in Eq. (44), corresponding to different horizontal locations of the vertical partition—all

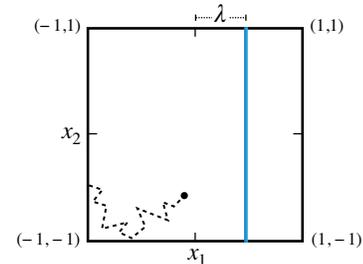


FIG. 4. A Szilard box with energy functions as in Eq. (44).

obey vertical reflection symmetry. Thus, the dynamics generated by those Fokker-Planck operators commute with ϕ_G , the twirling operator defined in Eq. (46). Using Eq. (25), we can bound the work extracted during any transformation $p \rightarrow p'$ in terms of the decrease of the accessible free energy $F_E[\phi_G(p)] - F_E[\phi_G(p')]$.

In more detail, consider some driving protocol which starts and ends with the partition removed. At intermediate times, the driving protocol manipulates the location of the partition so as to bring the system from some initial distribution p to a final equilibrium distribution $p' = u$ while extracting work. The second law gives bounds on EP, $\Sigma(p \rightarrow p') \geq 0$, and work:

$$W(p \rightarrow u) \leq F_{E^\circ}(p) - F_{E^\circ}(u) = D(p||u), \quad (47)$$

which follows from Eqs. (1) and (20). However, this bound can be too optimistic due to the protocol constraints. Given Eq. (18), as well as the fact that the final distribution obeys $\phi_G(u) = u$, we know that $\Sigma(p \rightarrow p') \geq D[p||\phi_G(p)]$. Similarly, Eq. (25) gives a tighter bound on extractable work:

$$W(p \rightarrow u) \leq F_{E^\circ}[\phi_G(p)] - F_{E^\circ}(u) = D[\phi_G(p)||u], \quad (48)$$

where the second equality follows from Eq. (20).

It is easy to use these results to resolve the question raised in the introduction: can one show that work can be extracted only from a measurement of whether the particle is in the left or right half of the box rather than a measurement of whether the particle is in the top or bottom half of the box? Suppose that the particle's initial distribution p is uniform across the left or right half of the box. Such a distribution p is invariant under vertical reflection, so $p = \phi_G(p)$ and Eq. (48) gives $W(p \rightarrow u) \leq D(p||u) = \ln 2$, the same as the bound set by the second law [Eq. (47)]. This bound can be achieved by quickly moving the partition to the middle of the box and then slowly moving it rightward. Conversely, suppose that, under the initial distribution p , the particle is uniformly distributed across the top or bottom half of the box. The twirling of such a distribution is a uniform distribution over the box, $\phi_G(p) = u$. In this case, Eq. (48) gives $W(p \rightarrow u) \leq 0$, meaning that no work can be extracted.

We now demonstrate the power of our approach by analyzing extractable work given a more complex family of initial distributions (while using the same energy functions as above). Suppose that the initial distribution is concentrated within half the box, as determined by a separating line that is rotated by an arbitrary angle $\theta \in [-\pi, \pi]$ [see Fig. 5(a)]. This initial distribution can be written formally as

$$p_\theta(x_1, x_2) = \frac{\mathbf{1}_{[-1,1]}(x_1, x_2)}{2} \Theta(x_2 \sin \theta - x_1 \cos \theta), \quad (49)$$

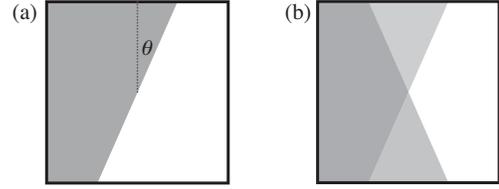


FIG. 5. (a) A nonequilibrium distribution p_θ that is “rotated” by an arbitrary angle θ [Eq. (49)]. (b) The distribution in (a) under the action of the vertical reflection twirling operator $\phi_G(p_\theta)$.

where Θ is the Heaviside function. For instance, p_θ for $\theta = 0$ corresponds to the particle being in the left half of the box, while p_θ for $\theta = \pi/2$ corresponds to the particle being in the top half of the box.

Because we are considering the same set of generators as above, we can bound the extractable work in a given p_θ using the same twirling operator as defined above in Eq. (46). [For a sample p_θ , the twirling $\phi_G(p_\theta)$ is illustrated in Fig. 5(b).] Using Eq. (48), the extractable work obeys $W(p_\theta \rightarrow u) \leq D[\phi_G(p_\theta)||u]$. Moreover, as we show in Appendix B 5, this KL divergence can be written in closed form as

$$D[\phi_G(p_\theta)||u] = \ln 2 \cdot \begin{cases} \frac{1}{2} \left| \tan\left(\theta - \frac{\pi}{2}\right) \right| & |\theta| \in \left(\frac{\pi}{4}, \frac{3\pi}{4}\right), \\ 1 - \frac{1}{2} |\tan \theta| & \text{otherwise.} \end{cases} \quad (50)$$

This result is plotted as a function of θ in Fig. 6.

We can also analyze the thermodynamics of information for different measurements of the Szilard box. Imagine that, starting from a uniform equilibrium distribution, one measures which side of the box contains the particle, as determined by a separating line at some arbitrary angle $\theta \in [-\pi, \pi]$. For this measurement, the conditional distribution over system states $p_{X|m}$ is equal to p_θ half the time [as in Fig. 5(a)] and equal to $p_{\theta+\pi}$ the other half the time. Then, for both measurement outcomes, one manipulates the vertical partition so as to drive the particle back to the equilibrium distribution $p' = u$

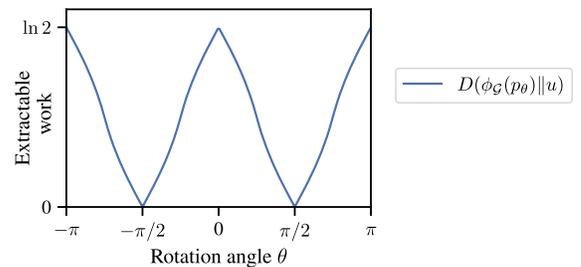


FIG. 6. Szilard box with symmetry constraints: the bound on extractable work as a function of θ [Eq. (50)].

while extracting work. For simplicity, we assume that the initial and final energy functions are the same.

The general bound on average extractable work for feedback control [Eq. (4)] gives

$$\langle W \rangle \leq I(X; M) = \ln 2, \quad (51)$$

where we use that $p = p'$ and $I(X'; M) = 0$. Our results provide a tighter bound, showing that the average extractable work is bounded by the accessible information in the measurement:

$$\langle W \rangle \leq I_{\text{acc}}^{\phi}(X; M) = \frac{D[\phi_{\mathcal{G}}(p_{\theta})||u] + D[\phi_{\mathcal{G}}(p_{\theta+\pi})||u]}{2}, \quad (52)$$

where we use Eqs. (34) and (36). It can be verified from Eq. (50) that $D[\phi_{\mathcal{G}}(p_{\theta})||u] = D[\phi_{\mathcal{G}}(p_{\theta+\pi})||u]$. Thus, the accessible information for a given θ is simply equal to $D[\phi_{\mathcal{G}}(p_{\theta})||u]$, the right side of Eq. (50), and shown in Fig. 6. As expected, the accessible information achieves a maximum of $\ln 2$ at $\theta = 0$ (or $\theta = \pm\pi$), which corresponds to a measurement of whether the particle is on the left or right side of the box. The accessible information falls nonlinearly (but continuously) to a minimum of 0 at $\theta = \pm\pi/2$, which corresponds to a measurement of whether the particle is on the top or bottom of the box.

In the example above, the accessible information quantifies in a very literal way the alignment between the choice of measurement and the way the system can be manipulated. More generally, this example illustrates how our bounds on EP and work depend on the interplay between the operator ϕ , the initial and final distributions p and p' , and (for feedback-control protocols) the choice of measurement M . This interplay can give rise to highly nontrivial thermodynamic bounds, such as in Eq. (50) and Fig. 6, even for very simple operators ϕ , such as in Eq. (46).

Finally, we note that our analysis above assumes only that the energy functions are vertically symmetric, which includes many energy functions that do not have the form of the vertical partition defined in Eq. (44). Furthermore, while the bounds on work and EP which we derive here are achievable by *some* vertically symmetric energy functions, they are not necessarily achievable by manipulating the location of a vertical partition. For instance, achieving the extractable work bound for a given θ [Eq. (50)] generally requires that the corresponding twirled distribution $\phi_{\mathcal{G}}(p)$, such as the one shown in Fig. 5(b), is an equilibrium distribution for some available energy function.

We analyze the same system using a different set of constraints in Secs. VIA and VIIA below. (Also see Ref. [42] for a different recent analysis of the thermodynamics of the Szilard box with rotated measurements, though from the point of view of partial observability rather than protocol constraints.)

B. Example: Feedback control on the Ising model

Our bounds on symmetry constraints can be useful for various multiparticle systems with symmetries, such as gases of indistinguishable particles and spin systems with symmetries. As a demonstration, we now analyze the thermodynamics of feedback control on an Ising model. The reader may also be interested in B 6, where we analyze a simpler (and more pedagogical) example of a discrete-state system with symmetry constraints.

Consider a 2D Ising model on a square lattice on a torus, containing a total of $N^2 = N \times N$ spins. The state of the lattice is indicated as $x \equiv (x_1, \dots, x_{N^2})$, where $x_i \in \{-1, 1\}$ is the state of the spin at location i . We assume that the energy functions have the following form:

$$E(x) = -J \sum_{(i,j) \in \mathcal{N}} x_i x_j - H \sum_i x_i, \quad (53)$$

where \mathcal{N} is the set of all nearest neighbors on the lattice, J is the coupling strength, and H is the external magnetic field.

Energy functions like these are invariant under the symmetry group \mathcal{G} corresponding to horizontal and vertical translations of the lattice (for simplicity, we ignore other symmetries of the lattice, such as reflections and rotations). The action of this group is given by a set of N^2 bijections $g_{a,b}: X \rightarrow X$ for $a, b \in \{0, \dots, N-1\}$, where $g_{a,b}(x)$ translates the lattice state x to the right by a spins and upward by b spins (with periodic boundary conditions). We assume that the system evolves according to Glauber dynamics [43], or some other dynamics that respects the translational symmetry of the 2D lattice, such that Eq. (39) is satisfied.

Given these assumption, we can derive thermodynamic bounds for the 2D Ising model in terms of the following twirling operator:

$$\phi_{\mathcal{G}}(p)(x) = N^{-2} \sum_{a=0}^{N-1} \sum_{b=0}^{N-1} g_{a,b}(x). \quad (54)$$

We use this twirling operator to analyze the thermodynamics of the following feedback-control setup on the Ising model, also shown in Fig. 7. The lattice is initially in equilibrium p at some temperature β and $J = 1$, $H = 0$ (no external field). The state of the spin at location 1 is then measured under the measurement channel $q(m|x) = \delta_m(x_1)$, where δ is the Kronecker delta. Since there is no initial external field, the two outcomes $m \in \{-1, 1\}$ have equal probability and $I(X; M) = \ln 2$. The measured outcome is then used to select a driving protocol, which extracts work from the system by manipulating the control parameters J and H . At the end of the protocol corresponding to each outcome, the system is brought back to the original equilibrium (so $p'_{X'|m} = p$ for all m). For simplicity, we assume that the initial and final energy functions are the same.

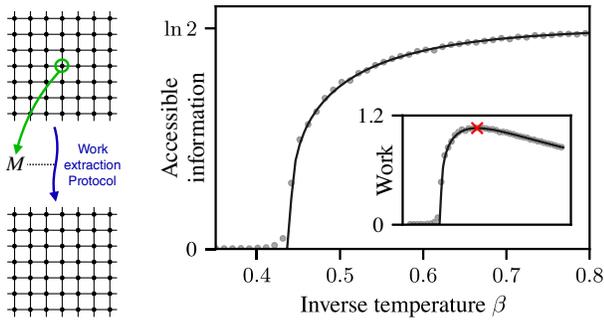


FIG. 7. Thermodynamics of information on a 2D Ising model. Left: A measurement M is made of the state of a single spin (green) and then used to drive the system while extracting work (blue). Right: The accessible information $I_{\text{acc}}^{\phi_G}(X; M)$ increases with inverse temperature after the critical value $\beta_c \approx 0.44$ [gray circles from Monte Carlo simulations, black line from closed-form expression in Eq. (56)]. The inset shows the bound on extractable work, $I_{\text{acc}}^{\phi_G}(X; M)/\beta$, which peaks at $\beta \approx 0.547$ (red cross).

Under this setup, one can verify that $I_{\text{acc}}^{\phi_G}(X'; M) = 0$ and $F_E(p) = F_E(p')$, so Eq. (34) bounds average extractable work as $\langle W \rangle \leq I_{\text{acc}}^{\phi_G}(X; M)/\beta$, where $I_{\text{acc}}^{\phi_G}(X; M)$ is the accessible information from Eq. (35). Using Eqs. (35) and (43), we can write this accessible information as

$$I_{\text{acc}}^{\phi_G}(X; M) = \ln 2 - \left\langle \ln \frac{q(m|x)}{N^{-2} \sum_{a,b} q[m|g_{a,b}(x)]} \right\rangle, \quad (55)$$

where $\langle \cdot \rangle$ indicates expectation over the joint distribution $p(x)q(m|x)$, where $p(x)$ is the initial equilibrium distribution at inverse temperature β and $J = 1$, $H = 0$. We emphasize that the accessible information depends on β (though we leave this dependence implicit in the notation).

In general, one can estimate the accessible information in Eq. (55) using various numerical techniques (e.g., by sampling from the initial equilibrium distribution using Monte Carlo methods). It is also possible to use Onsager's well-known solution of the 2D Ising model to calculate the accessible information in closed form. In particular, in Appendix B 7, we show that in the thermodynamic limit $N \rightarrow \infty$,

$$I_{\text{acc}}^{\phi_G}(X; M) = \begin{cases} 0 & \text{for } \beta \leq \beta_c, \\ \ln 2 - h_2 \left(\frac{1 + \sqrt{1 - (\sinh 2\beta)^{-4}}}{2} \right) & \text{for } \beta > \beta_c, \end{cases} \quad (56)$$

where $h_2(x) = -x \ln x - (1-x) \ln(1-x)$ is the binary entropy function and $\beta_c = \ln(1 + \sqrt{2})/2 \approx 0.44$ is the critical inverse temperature of the 2D Ising model. This result is verified in Fig. 7, where we compare Eq. (56) with a Monte Carlo estimate of Eq. (55) on a 100×100 lattice. It can be seen that, in the high-temperature (low- β) regime,

the accessible information vanishes. In the low-temperature (high- β) regime, the amount of accessible information increases, approaching $\ln 2$ as $\beta \rightarrow \infty$.

We also plot the bound on average extractable work, $\langle W \rangle \leq I_{\text{acc}}^{\phi_G}(X; M)/\beta$, in the inset in Fig. 7. This bound is the ratio of two terms: the accessible information $I_{\text{acc}}^{\phi_G}(X; M)$ and the inverse temperature β , both of which are increasing in β . In fact, it can be seen from Fig. 7 that the bound on extractable work peaks at a finite value of β , the optimal inverse temperature for work extraction. Using Eq. (56) and numerical techniques, we find this optimal value to be $\beta \approx 0.547$, leading to the bound $\langle W \rangle \leq 1.06$ joules.

This result shows that the amount of accessible information provided by a given measurement can depend on the structure of correlations in the system and, therefore, vary dramatically as the system undergoes a phase transition. At a high level, any driving protocol that is restricted to energy functions like Eq. (53) can extract work only from “global” (i.e., translationally invariant) information. If the measurement acquires such information (e.g., if it directly measures the spatially averaged magnetization), then, in principle, all of the acquired information may be extractable as work. Measurement of the state of a single spin, however, generally provides only local information. The temperature dependence observed in Eq. (56) and Fig. 7 arises from the presence of long-range order in the magnetic regime ($\beta > \beta_c$). In this regime, the state of each spin is highly correlated with the magnetization of the entire lattice, so local and global information are equivalent. In the high-temperature regime ($\beta < \beta_c$), the state of a single spin is not correlated with any kind of global information, and so most of the measured information is inaccessible.

For a different kind of analysis of the thermodynamics of a 1D Ising model under constraints, see Ref. [44].

VI. MODULARITY CONSTRAINTS

Many systems of interest exhibit modular organization, meaning that their degrees of freedom can be grouped into decoupled subsystems. Examples of modular systems include computational devices such as digital circuits [19,45,46], regulatory networks in biology [47], and brain networks [48].

We use our framework to derive bounds on work and EP for modular systems. We begin by introducing some terminology and notation. Consider a system whose degrees of freedom are indexed by the set V , such that the overall state space can be written as $X = \times_{v \in V} X_v$, where X_v is the state space of degree of freedom v . We use the term *subsystem* to refer to any subset of the degrees of freedom, $A \subseteq V$. We use X_A to indicate the random variable representing the state of subsystem A and x_A to indicate an actual state of A . Given some distribution p over the entire system, we use p_A to indicate a marginal distribution over

subsystem A and $[Lp]_A$ to indicate the derivative of the marginal distribution of subsystem A under the generator L .

We use the term *modular decomposition* to refer to a set of subsystems \mathcal{C} , such that each $v \in V$ belongs to at least one subsystem $A \in \mathcal{C}$. Note that some of the degrees of freedom $v \in V$ can belong to more than one subsystem in \mathcal{C} . We use

$$O(\mathcal{C}) = \bigcup_{A, B \in \mathcal{C}: A \neq B} (A \cap B) \quad (57)$$

to indicate those degrees of freedom that belong to more than one subsystem in \mathcal{C} , which we refer to as the *overlap*. We often write O instead of $O(\mathcal{C})$ for notational simplicity.

We say that the available driving protocols obey *modularity constraints* (with respect to the modular decomposition \mathcal{C}) if each generator $L \in \Lambda$ can be written as a sum of generators of the different subsystems in \mathcal{C} :

$$L = \sum_{A \in \mathcal{C}} L^{(A)}, \quad (58)$$

and each $L^{(A)}$ obeys two properties: The dynamics over the marginal distribution p_A are closed under $L^{(A)}$ (i.e., depend only on the marginal distribution over A):

$$p_A = q_A \Rightarrow [L^{(A)}p]_A = [L^{(A)}q]_A \quad \forall p, q \in \mathcal{P}, \quad (59)$$

and the distribution over other subsystems besides A does not change under $L^{(A)}$:

$$[L^{(A)}p]_B = 0 \quad \forall p \in \mathcal{P}, \quad B \in \mathcal{C} \setminus \{A\}. \quad (60)$$

In other words, we require that each subsystem evolves independently and does not affect the other subsystems.

The role of the degrees of freedom in the overlap is somewhat subtle. It can be verified that Eq. (60) implies that the degrees of freedom in the overlap cannot change state when evolving under L . Importantly, however, the overlap may influence the dynamics of those degrees of freedom that *can* change state. For example, consider an inclusive model of a feedback-control setup: There are two nested subsystems, $\mathcal{C} = \{A, B\}$ with $B \subseteq A$, and the degrees of freedom in $O = B$ (the controller) cannot change state but can influence the evolution of $A \setminus B$. More elaborate feedback-control setups, in which the same controller can control multiple subsystems, can be modeled using decompositions with multiple non-nested subsystems. Other examples of modular decompositions with overlap include circuits [19], spin systems where some spins are pinned by local magnetic fields, and many-particle systems where some particles have no mobility.

We can also provide more concrete conditions when Eqs. (59) and (60) hold for discrete-state master equations and Fokker-Planck equations. For discrete-state master equations, it can be verified by inspection that Eqs. (59) and (60) hold when all $L \in \Lambda$ can be written in the form

$$L_{x'x} = \sum_{A \in \mathcal{C}} R_{x'_A, x_A}^{(A)} \delta_{x'_{V \setminus A}}(x'_{V \setminus A}), \quad (61)$$

where δ is the Kronecker delta and $R^{(A)}$ is some rate matrix over subsystem A that does not allow the degrees of freedom in the overlap to change state ($R_{x'_A, x_A}^{(A)} = 0$ if $x_{A \cap O} \neq x'_{A \cap O}$).

For Fokker-Planck equations, for simplicity, consider overdamped dynamics of the form

$$Lp = \sum_{v \in V} \gamma_v^L \partial_{x_v} [(\partial_{x_v} E_L)p + \beta^{-1} \partial_{x_v} p], \quad (62)$$

where γ_v^L is the mobility coefficient along dimension v and E_L is the potential energy function associated with generator L . Such equations can represent potential-driven Brownian particles coupled to a heat bath, where the different mobility coefficients represent different particle masses or sizes [49]. Now imagine that, for all $L \in \Lambda$, the energy functions are additive over the subsystems and that the degrees of freedom in the overlap have no mobility:

$$E_L(x) = \sum_{A \in \mathcal{C}} E_L^{(A)}(x_A), \quad \gamma_v^L = 0 \quad \forall v \in O. \quad (63)$$

In that case, Eq. (62) can be rewritten in the form of Eq. (58), with $L^{(A)}p = \sum_{v \in A \setminus O} \gamma_v^L \partial_{x_v} [(\partial_{x_v} E_L^{(A)})p_A + \beta^{-1} \partial_{x_v} p_A]$, and satisfies Eqs. (59) and (60).

We now define the following nonlinear operator $\phi_{\mathcal{C}}$:

$$\phi_{\mathcal{C}}(p) = p_O \prod_{A \in \mathcal{C}} p_{A \setminus O | A \cap O}. \quad (64)$$

This operator preserves the statistical correlations within each subsystem $A \in \mathcal{C}$, as well as within the overlap O , while destroying all other statistical correlations. As a simple example, if all the subsystems in \mathcal{C} are nonoverlapping, then $\phi_{\mathcal{C}}(p)$ has the product form $\phi_{\mathcal{C}}(p) = \prod_{A \in \mathcal{C}} p_A$. In Appendix C, we show that $\phi_{\mathcal{C}}$ obeys the Pythagorean identity [Eq. (14)]. We also show that, if some generator $L(t)$ obeys Eqs. (59) and (60), then $e^{\tau L(t)}$ commutes with $\phi_{\mathcal{C}}$, so Eq. (16) holds.

As a result, for any protocol that carries out the transformation $p \rightarrow p'$ while obeying modularity constraints, the decompositions and bounds for EP and work derived in Sec. III are satisfied for $\phi = \phi_{\mathcal{C}}$. In particular, using Eq. (21), we can decompose the free energy $F_E(p)$ of any distribution p into the accessible free energy $F_E[\phi_{\mathcal{C}}(p)]$ and the inaccessible free energy $D[p || \phi_{\mathcal{C}}(p)]/\beta$. Note that $D[p || \phi_{\mathcal{C}}(p)]$ is a non-negative measure of the amount of statistical correlations between the subsystems of \mathcal{C} under distribution p , which vanishes when each subsystem is conditionally independent given the overlap O . Thus, for a protocol that obeys modularity constraints, Eq. (18) states

that the drop in those statistical correlations is a lower bound on EP, and that the amount of statistical correlation between the subsystems of \mathcal{C} cannot increase over the course of the protocol. (There is a fair amount of closely related prior work; see Sec. VIII.)

A particularly simple application of our bounds occurs when \mathcal{C} contains two (possibly overlapping) subsystems, $\mathcal{C} = \{A, B\}$. In that case, the bounds in Eq. (18) can be rewritten in terms of the drop of a conditional mutual information between the two subsystems:

$$\Sigma(p \rightarrow p') \geq I(X_A; X_B | X_{A \cap B}) - I(X'_A; X'_B | X'_{A \cap B}) \geq 0. \quad (65)$$

If the subsystems do not overlap, this expression can be further rewritten as the drop of the regular mutual information:

$$\Sigma(p \rightarrow p') \geq I(X_A; X_B) - I(X'_A; X'_B) \geq 0. \quad (66)$$

More generally, if \mathcal{C} contains an arbitrary number of nonoverlapping subsystems, the EP can be bound as

$$\Sigma(p \rightarrow p') \geq \mathcal{I}(p) - \mathcal{I}(p') \geq 0, \quad (67)$$

where $\mathcal{I}(p) = [\sum_{A \in \mathcal{C}} S(p_A)] - S(p)$ is the multi-information in distribution p with respect to partition \mathcal{C} [50].

We finish by discussing thermodynamics of information under modularity constraints. In general, the results derived in Sec. IV apply to modularity constraints as a special case. However, we can also exploit special properties of the operator $\phi_{\mathcal{C}}$ to further simplify the expression of accessible information. Suppose that the distribution p is invariant under $\phi_{\mathcal{C}}$, so $p = \phi_{\mathcal{C}}(p)$ [e.g., if p is an equilibrium distribution; see Eq. (17)]. Using Eq. (64), we can then rewrite Eq. (36) as

$$I_{\text{acc}}^{\phi_{\mathcal{C}}}(X; M) = I(X_O; M) + \sum_{A \in \mathcal{C}} I(X_A; M | X_{A \cap O}). \quad (68)$$

Thus, the accessible information in measurement M is the information that M provides about the overlap, plus the conditional mutual information between each subsystem and M given the relevant part of the overlap. This result means that only information about individual subsystems—not about intersubsystem correlations—can be turned into work. If there is no overlap, Eq. (68) can be further simplified as

$$I_{\text{acc}}^{\phi_{\mathcal{C}}}(X; M) = \sum_{A \in \mathcal{C}} I(X_A; M). \quad (69)$$

We use these expressions in some of our examples below.

A. Example: Szilard box with modularity constraints

We illustrate our results for modularity constraints on a Szilard box. In doing so, we demonstrate two important

concepts: first, how the same set of generators Λ can be analyzed under different constraints, resulting in different bounds on work and EP (compare this section to Sec. VA); second, how bounds arising from multiple constraints can be stacked on top of each in an iterative manner, as in Eq. (28) (we combine bounds from modularity and symmetry constraints).

We consider the same setup as in Sec. VA: There is a single overdamped particle in a box coupled to a bath at inverse temperature $\beta = 1$, which evolves under potential energy functions as in Eq. (44). This system is driven from some initial distribution p to a final uniform equilibrium distribution $p' = u$ while extracting work.

Note that the energy functions in Eq. (44) have no interaction terms between x_1 (the horizontal position of the particle) and x_2 (the vertical position of the particle). That means that the allowed driving protocols obey modularity constraints for a decomposition of the system into two subsystems, $\mathcal{C} = \{\{X_1\}, \{X_2\}\}$, since Eq. (63) is satisfied for the decomposition. This result allows us to analyze EP and work using an operator $\phi_{\mathcal{C}}$ which maps each joint distribution over $X_1 \times X_2$ into a product distribution:

$$\phi_{\mathcal{C}}(p)(x_1, x_2) = p(x_1)p(x_2). \quad (70)$$

In particular, using the same derivation as in Eq. (48), we can bound the extractable work in terms of the accessible free energy in p :

$$W(p \rightarrow u) \leq D[\phi_{\mathcal{C}}(p) \| u]. \quad (71)$$

As discussed in Sec. VA, this system also obeys symmetry constraints, corresponding to the vertical reflection twirling operator $\phi_{\mathcal{G}}$ defined in Eq. (46). We can use Eq. (29) to bound the extractable work using a combination of $\phi_{\mathcal{C}}$ and $\phi_{\mathcal{G}}$:

$$W(p \rightarrow u) \leq D\{\phi_{\mathcal{C}}[\phi_{\mathcal{G}}(p)] \| u\}, \quad (72)$$

$$W(p \rightarrow u) \leq D\{\phi_{\mathcal{G}}[\phi_{\mathcal{C}}(p)] \| u\}. \quad (73)$$

For concreteness, imagine that the initial distribution p is concentrated within half the box, as determined by a separating line rotated by some arbitrary angle $\theta \in [-\pi, \pi]$, so $p = p_{\theta}$ from Eq. (49) [see Fig. 5(a) for an illustration].

We consider the extractable work bound in Eq. (71) for the initial distribution p_{θ} . For a given p_{θ} , the corresponding decorrelated initial distribution $\phi_{\mathcal{C}}(p_{\theta})$ is illustrated in Fig. 8(a). Then, the accessible free energy in Eq. (71) can be expressed in closed form as (see Appendix C 3)

$$D[\phi_{\mathcal{C}}(p_{\theta}) \| u] = \ln 4 - \frac{1}{2} (\min\{|\tan \theta|, |\tan(\pi/2 - \theta)|\} + f[\max\{|\tan \theta|, |\tan(\pi/2 - \theta)|\}]), \quad (74)$$

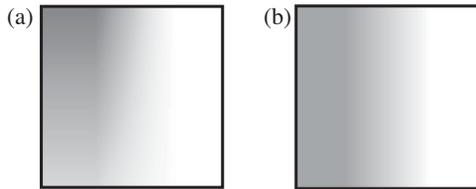


FIG. 8. (a) Given a “rotated” distribution p_θ , as shown above in Fig. 5(a), the decorrelated distribution $\phi_C(p_\theta)$, as in Eq. (70). (b) The decorrelated and twirled distribution $\phi_G[\phi_C(p_\theta)]$.

where for notational convenience we define

$$f(x) = 1 - \frac{1+x^2}{2x} \ln \frac{x+1}{x-1} - \ln \frac{x^2-1}{4x^2}. \quad (75)$$

Equation (74) is plotted in Fig. 9 in green. Note that this function peaks at both $\theta \in \{-\pi, 0, \pi\}$ (i.e., when the particle is in the left or right half of the box) as well as $\theta \in \{-\pi/2, \pi/2\}$ (i.e., when the particle is in the top or bottom half of the box)—precisely those θ for which p_θ has no correlations between the horizontal and vertical position of the particle.

Next, we consider the extractable work bound in Eq. (72) for the initial distribution p_θ . It can be verified that $\phi_G[\phi_C(p_\theta)](x_1, x_2) = p_\theta(x_1)u(x_2)$, which is illustrated in Fig. 8(a). The right-hand side of Eq. (72) can again be expressed in closed form as (see Appendix C 3)

$$D\{\phi_G[\phi_C(p_\theta)]||u\} = \ln 2 - \frac{1}{2} \begin{cases} f(|\tan \theta|) & \text{if } |\theta| \in \left(\frac{\pi}{4}, \frac{3\pi}{4}\right), \\ |\tan \theta| & \text{otherwise} \end{cases} \quad (76)$$

with f defined as in Eq. (75). This result is shown in Fig. 9 in orange. Note also that $\phi_G[\phi_C(p_\theta)] = \phi_C[\phi_G(p_\theta)]$ for all p_θ , so the bounds in Eqs. (72) and (73) are equivalent.

For comparison, we also plot the extractable work bound derived using symmetry constraints [Eq. (50)] (Fig. 9 in blue). It is clear that the bound derived by exploiting a combination of modularity and symmetry constraints (in orange) is strictly tighter than the bounds derived by using either only modularity (green) or only symmetry constraints (blue) individually.

One can also use the bounds derived in this section to analyze the accessible information in a measurement of the Szilard box. Imagine that, starting from a uniform equilibrium distribution, one measures which side of the box contains the particle, as determined by a separating line at some arbitrary angle $\theta \in [-\pi, \pi]$. For this measurement, the conditional distribution over system states $p_{X|m}$ is equal to p_θ half the time and equal to $p_{\theta+\pi}$ the other half the time. One can then derive bounds on accessible information such

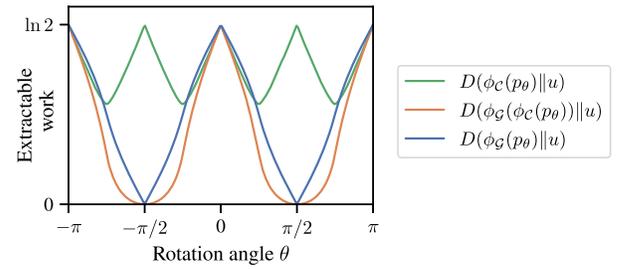


FIG. 9. Bounds on extractable work as a function of θ , as derived from only modularity constraints [in green, Eq. (74)], a combination of modularity + symmetry constraints [in orange, Eq. (76)], and only symmetry constraints [in blue, Eq. (50)].

as Eq. (52), while using the bounds derived in this section [Eqs. (71)–(73)].

B. Example: Generalized Szilard box

Our results on modularity constraints can be useful for analyzing the thermodynamics of multiparticle systems. As an example, consider the “generalized Szilard box” feedback-control scenario analyzed in Ref. [20]. Here, a box containing an ideal gas of N particles, which are indexed by $v \in V$, begins in uniform equilibrium with a heat bath at inverse temperature β . Several partitions are inserted into the box, separating the box into separate volumes, and a measurement M is made of the number of particles in each volume (see the illustration in Fig. 10). The box is then separated from the bath, and, depending on the outcome of the measurement, the partitions are moved so as to equalize the pressure within each volume while extracting work. To make the process repeatable, suppose that at the end of the protocol, the partitions are removed and the box is again equilibrated with the bath (note that this last step does not contribute to extracted work).

The ideal gas assumption means that the particles do not interact, so by Eqs. (59) and (60) the protocol obeys modularity constraints with respect to a decomposition in which each particle is a separate subsystem. The corresponding operator ϕ_C is given by

$$\phi_C(p)(x) = \prod_{v=1}^N p(x_v). \quad (77)$$

Given Eq. (34), the average extractable work for the above feedback-control scenario is bounded by $\langle W \rangle \leq I_{\text{acc}}^{\phi_C}(X; M)/\beta$, which can also be written in terms of the information provided by the measurement M about each individual particle:

$$\langle W \rangle \leq \sum_{v=1}^N I(X_v; M)/\beta, \quad (78)$$

as follows from Eq. (69). In fact, by symmetry of the initial distribution, the measurement provides the same information

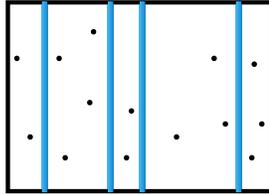


FIG. 10. A generalized Szilard box with multiple particles [20].

about each particle, $I(X_v; M) = I(X_1; M)$ for all v , so we can further rewrite Eq. (78) as $\langle W \rangle \leq N \cdot I(X_1; M)/\beta$.

This result shows that Eq. (78), which is reported as one of the main results of Ref. [20] [Eq. (5)], follows immediately from our framework. Moreover, our derivation holds under a broader set of conditions than those considered in Ref. [20], since it does not rely on any of the details of setup (such as the type of partitions, the particular work extraction protocol, or even the assumption that the particles are identical).

C. Example: Collective flashing ratchet

As a final example of modularity constraints, we consider the ‘‘collective flashing ratchet,’’ a classic model in the literature on the thermodynamics of information [21,52]. This system involves N overdamped particles evolving under an additive potential

$$E(x) = \lambda \sum_{v=1}^N V(x_v), \quad (79)$$

where V is a single-particle potential and $\lambda \in \{0, 1\}$ is a control parameter that can be used to turn the potential on and off. The single-particle potential V is chosen as an asymmetrical sawtooth ‘‘ratchet’’ pattern, shown in Fig. 11, where $\alpha \in [0, 1/2]$ parameterizes the degree of asymmetry.

By manipulating λ over time, possibly in a way that depends on measurements of the system, the particles can be driven so as to have a net directional flux or to do work against the externally applied force [53]. For instance, in a feedback-control setup, λ is determined by the outcome of some measurement M . The most common strategy involves turning the ratchet potential on when the net force on the particles is positive and turning it off otherwise, according to the following measurement channel [21]:

$$q(m|x) = \delta_m \{ \Theta[\sum_v V'(x_v)] \}, \quad (80)$$

where Θ is the Heaviside function. Note that this system has been experimentally realized [54].

Suppose that starting from some initial distribution p , the measurement in Eq. (80) is performed. As is common in the literature [21], we assume that under p the particles are identically and independently distributed and that each

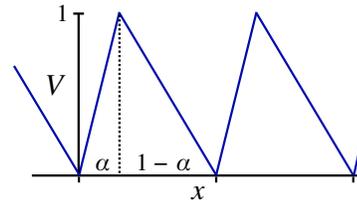


FIG. 11. The sawtooth potential of the flashing ratchet, from Ref. [21].

particle is in the increasing part of the potential [$V'(x_v) \geq 0$] with probability α (see Fig. 11). The measurement outcome is then used to drive the system back to distribution p while extracting work by manipulating the system’s energy function, all while coupled to a heat bath at inverse temperature β . We assume that the driving protocols start and end on the same energy function and that only additive potentials (without interaction terms) are applied to the system during the driving [this assumption allows for potentials such as Eq. (79), as well as many others].

The driving protocols obey Eq. (63) for a decomposition where each particle is its own subsystem, corresponding to the same type of ϕ_C as in Eq. (77), $\phi_C(p)(x) = \prod_{v \in V} p(x_v)$. As in Sec. VI A, we can use Eq. (34) to bound average extractable work as $\langle W \rangle \leq I_{\text{acc}}^{\phi_C}(X; M)/\beta$. Using Eq. (69),

$$I_{\text{acc}}^{\phi_C}(X; M) = \sum_{v=1}^N I(X_v; M) = N \cdot I(X_1; M), \quad (81)$$

where we use that the measurement provides the same information about each particle, $I(X_v; M) = I(X_1; M)$ for all v (as follows from a symmetry argument).

In Appendix C 4, we show that $I_{\text{acc}}^{\phi_C}(X; M)$ can be computed in closed form. Values of $I_{\text{acc}}^{\phi_C}(X; M)$ for different values of N (the number of particles) and α (the asymmetry parameter) are plotted in Fig. 12 (left). Note that the accessible information shows a nonmonotonic behavior in the number of particles for $\alpha \neq 0.5$. This behavior occurs because, for a highly asymmetric potential, the total amount of acquired information grows with N : $I(X; M)$ grows from a minimum value of $h_2(\alpha)$ for $N = 1$ to a maximum value of $\ln 2$ as $N \rightarrow \infty$. Given this observation, we also calculate the ‘‘efficiency’’ of the measurements in terms of the ratio $I_{\text{acc}}^{\phi_C}(X; M)/I(X; M)$. This calculation is shown in Fig. 12 (right) for various values of N and α . Interestingly, lower values of α (higher values of asymmetry) have higher efficiency values.

In the $N \rightarrow \infty$ limit, accessible information and efficiency converge to a single value, irrespective of α . In Appendix C 4, we show that the accessible information $I_{\text{acc}}^{\phi_C}(X; M)$ converges to $1/\pi \approx 0.32$ nats, while the efficiency $I_{\text{acc}}^{\phi_C}(X; M)/I(X; M)$ converges to $1/(\pi \ln 2) \approx 0.46$ (dotted lines in Fig. 12).

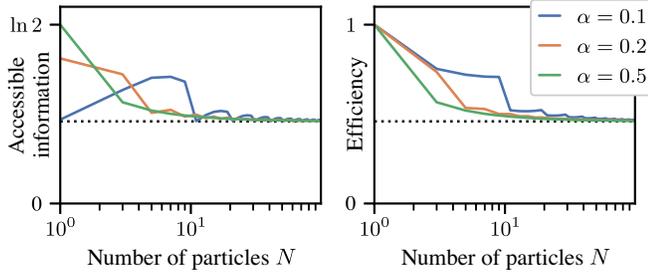


FIG. 12. Left: accessible information $I_{\text{acc}}^{\phi_c}(X; M)$ for the collective flashing ratchet, as a function of N (number of particles) and α (asymmetry). Right: the efficiency of the measurements, $I_{\text{acc}}^{\phi_c}(X; M)/I(X; M)$.

For a different (and complementary) theoretical analysis of extracted work in a feedback-controlled flashing ratchet, see Ref. [53].

VII. COARSE-GRAINED CONSTRAINTS

In our final results section, we consider bounds on EP and work that arise from coarse-grained constraints.

We begin by introducing some notation and preliminaries. Let $\xi: X \rightarrow Z$ be some coarse graining of the microscopic state space X , where Z is a set of macrostates. For any distribution p over X , we use $p_Z(z) = \int \delta_{\xi(x)}(z) p(x) dx$ to indicate the corresponding distribution over the macrostates Z , $p_{X|Z}(x|z) = p(x)/p_Z(z)$ to indicate the conditional probability distribution of microstates within macrostates, and $\mathcal{P}_Z := \{p_Z: p \in \mathcal{P}\}$ to indicate the set of all coarse-grained distributions. Finally, for any generator L and distribution p , we use $[Lp]_Z$ to indicate the resulting instantaneous dynamics of the coarse-grained distribution p_Z .

To derive our bounds, we suppose that the dynamics over the coarse-grained distributions are closed; i.e., for all $L \in \Lambda$,

$$p_Z = q_Z \Rightarrow [Lp]_Z = [Lq]_Z \quad \forall p, q \in \mathcal{P}. \quad (82)$$

Given this assumption, the evolution of the coarse-grained distribution p_Z can be represented by a coarse-grained generator, which we write as $\partial_t p_Z = \hat{L} p_Z$ (discussed in detail below).

We can specify more concrete conditions that guarantee that (82) holds for a given generator L (see Appendix D for details). For a discrete-state rate matrix L , it is satisfied when

$$\sum_{x: \xi(x)=z} L_{xx'} = \hat{L}_{z, \xi(x')} \quad \forall x', \quad z \neq \xi(x'), \quad (83)$$

where $\hat{L}_{z, z'}$ is some coarse-grained transition rate from macrostate z' to macrostate z . Equation (83) states that, for

each microstate x' , the total rate of transitions from x' to microstates located in another macrostate $z \neq \xi(x')$ depends only on the macrostate $\xi(x')$, not on x' directly. This condition has been sometimes called “lumpability” in the literature [55].

For a continuous-state master equation, Eq. (82) is satisfied when a continuous-state version of Eq. (83) (with sums replaced by integrals) holds. Moreover, for certain Fokker-Planck equation and linear coarse-graining functions, Eq. (83) can be replaced by a simple coarse-graining condition on the energy functions. Suppose each $L \in \Lambda$ is a Fokker-Planck operator like

$$Lp = \nabla \cdot (\nabla E_L) p + \beta^{-1} \Delta p \quad (84)$$

and that ξ is a linear function: $\xi(x) = Wx$ (where W is some full-rank $m \times n$ matrix, $m \leq n$). Without loss of generality, we assume that W is scaled so that $WW^T = I$ [56]. In addition, suppose that each energy function satisfies

$$W \nabla E_L(x) = -\hat{F}[\xi(x)] \quad \forall x \quad (85)$$

for some arbitrary macrostate drift function $\hat{F}: Z \rightarrow \mathbb{R}$. Then, the coarse-grained generator \hat{L} itself will have a Fokker-Planck form (see Ref. [57] and Appendix D):

$$\hat{L} p_Z = -\nabla \cdot \hat{F} p_Z + \beta^{-1} \Delta p_Z. \quad (86)$$

The right side of Eq. (86) depends only on p_Z and not the full microstate distribution p , so Eq. (82) is satisfied.

Importantly, if Eq. (82) holds, the EP rate at time t can be bounded as (see Appendix D)

$$\dot{\Sigma}[p(t), L(t)] \geq -\sum_z \partial_t p_Z(z, t) \ln \frac{p_Z(z, t)}{\pi_Z^{L(t)}(z)} \geq 0, \quad (87)$$

where $\partial_t p_Z(t) = \hat{L} p_Z(t)$ and $\pi_Z^{L(t)}$ is the coarse-grained version of $\pi^{L(t)}$, the stationary distribution of $L(t)$. The right-hand side of Eq. (87) is the coarse-grained version of Eq. (11), which arises from the macrostate distribution p_Z being out of equilibrium. We then define the total “coarse-grained EP” over the course of the protocol as the time integral of the middle term in Eq. (87):

$$\hat{\Sigma}(p_Z \rightarrow p'_Z) = \int_0^1 -\sum_z \partial_t p_Z(z, t) \ln \frac{p_Z(z, t)}{\pi_Z^{L(t)}(z)} dt. \quad (88)$$

Given (87), the coarse-grained EP serves as a non-negative lower bound on the total EP:

$$\Sigma(p \rightarrow p') \geq \hat{\Sigma}(p_Z \rightarrow p'_Z) \geq 0. \quad (89)$$

Note that Ref. [58] previously derived a coarse-grained EP rate for discrete-state master equations, which differs from

the one that appears on the right-hand side in Eq. (87); however, Eq. (87) can be seen as the “nonadiabatic component” of the coarse-grained EP rate from Ref. [58] and is a lower bound on it [25].

We say that the available driving protocols obey *coarse-grained constraints* if the generators $L \in \Lambda$ exhibit closed dynamics over Z [Eq. (82)] and there is some operator $\hat{\phi}: \mathcal{P}_Z \rightarrow \mathcal{P}_Z$ that obeys the Pythagorean identity [Eq. (14)] and the commutativity relation [Eq. (16)] with respect to all \hat{L} . For example, this coarse-grained operator $\hat{\phi}$ might reflect the presence of symmetry or modularity constraints on the coarse-grained dynamics.

We can then use Eq. (89) and the framework developed in Sec. III to derive bounds on work and EP. In particular, Eq. (18) implies the following bound on coarse-grained EP: $\hat{\Sigma}(p_Z \rightarrow p'_Z) \geq D[p_Z \|\hat{\phi}(p_Z)] - D[p'_Z \|\hat{\phi}(p'_Z)] \geq 0$. Using Eq. (89), we can then bound overall EP as

$$\Sigma(p \rightarrow p') \geq D[p_Z \|\hat{\phi}(p_Z)] - D[p'_Z \|\hat{\phi}(p'_Z)] \geq 0. \quad (90)$$

Using Eq. (2), we can also bound extractable work as

$$W(p \rightarrow p') \leq F_E(p) - F_E(p') - \{D[p_Z \|\hat{\phi}(p_Z)] - D[p'_Z \|\hat{\phi}(p'_Z)]\} / \beta. \quad (91)$$

Equations (90) and (91) can also be used to derive bounds on average work extraction in feedback-control protocols, using the strategy described in Sec. IV.

If $\hat{\phi}$ represents coarse-grained symmetry or modularity constraints, then Eq. (90) implies that any asymmetry or intersubsystem correlation in the macrostate distribution can only be dissipated away, not turned into work. Another simple application occurs when all $L \in \Lambda$ have the same coarse-grained equilibrium distribution; i.e., there is some π_Z such that $\hat{L}\pi_Z = 0$ for all L . In this case, $\hat{\phi}(p) = \pi_Z$ satisfies Eqs. (14) and (16) at the coarse-grained level [compare to the derivation of Eq. (27) above]. Applying Eq. (90) then gives

$$\Sigma(p \rightarrow p') \geq D(p_Z \|\pi_Z) - D(p'_Z \|\pi_Z) \geq 0, \quad (92)$$

as well as a corresponding extractable work bound, as in Eq. (91). This result shows that, if the coarse-grained equilibrium distribution π_Z cannot change, then any deviation between the actual coarse-grained distribution p_Z and π_Z must be dissipated as EP, not turned into work.

A. Example: Szilard box

We demonstrate our results on coarse-grained constraints using the Szilard box. We consider a similar setup as in Secs. VA and VIA, where there is a single overdamped particle in a box coupled to a bath at inverse temperature $\beta = 1$. However, we now assume that there is a vertical gravitational force, as illustrated in Fig. 13. Formally, this

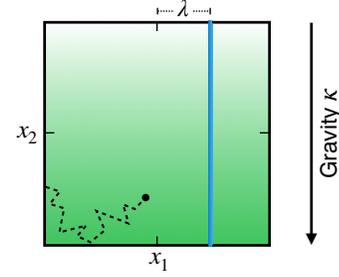


FIG. 13. A two-dimensional Szilard box with a Brownian particle, in the presence of gravity.

assumption means that the available potential energy functions have the form

$$E_\lambda(x_1, x_2) = V_p(x_1 - \lambda) + V_w(|x_1|) + V_w(|x_2|) + \kappa x_2, \quad (93)$$

where κ is a fixed constant that determines the strength of gravity. Unlike Eq. (44), this energy function in Eq. (93) no longer obeys the reflection symmetry $(x_1, x_2) \mapsto (x_1, -x_2)$.

The microstate of the particle is represented by the horizontal and vertical position $x = (x_1, x_2)$. We consider a coarse graining in which the macrostate is the vertical coordinate of the particle $Z = X_2$, corresponding to the coarse-graining function $\xi(x_1, x_2) = Wx = x_2$ with $W = [0 \ 1]$. It is easy to check that the potential energy functions in Eq. (93) satisfy

$$W\nabla E_\lambda(x) = \partial_{x_2}[V_w(|x_2|) + \kappa x_2], \quad (94)$$

which obeys Eq. (85) and, therefore, guarantees that the coarse-grained dynamics are closed. In fact, the coarse-grained generators have the Fokker-Planck form of Eq. (86) with the coarse-grained drift function $\hat{F}(x_2) = -\partial_{x_2}[V_w(|x_2|) + \kappa x_2]$, which leads to the following Boltzmann stationary distribution:

$$\begin{aligned} \pi_{X_2}(x_2) &\propto e^{-\beta[V_w(|x_2|) + \kappa x_2]} \\ &= \mathbf{1}_{[-1, 1]}(x_2) e^{-\beta \kappa x_2}, \end{aligned} \quad (95)$$

where in the second line we used the form of $V_w(\cdot)$ from Eq. (45). Since the coarse-grained equilibrium distribution is the same for all energy functions having the form Eq. (93), we can use the EP bound in Eq. (92).

Suppose that the system starts from some initial distribution p and is then driven to a final equilibrium distribution p' while extracting work. We assume that the partition is removed at the beginning and end of the protocol, corresponding to the energy function $E^\emptyset(x_1, x_2) = V_w(|x_1|) + V_w(|x_2|) + \kappa x_2$, with the Boltzmann distribution

$$\pi^\emptyset(x_1, x_2) \propto \mathbf{1}_{[-1, 1]^2}(x_1, x_2) e^{-\beta \kappa x_2}. \quad (96)$$

We also assume that the final distribution is in equilibrium, so $p' = \pi^\varnothing$. Then, the extractable work involved in this transformation can be expressed as

$$\begin{aligned} W(p \rightarrow \pi^\varnothing) &= F_{E^\varnothing}(p) - F_{E^\varnothing}(\pi^\varnothing) - \Sigma(p \rightarrow \pi^\varnothing) \\ &= D(p \parallel \pi^\varnothing) - \Sigma(p \rightarrow \pi^\varnothing), \end{aligned} \quad (97)$$

where we use Eqs. (2) and (5). We can then upper bound extractable work by combining Eq. (97) with various lower bounds on $\Sigma(p \rightarrow \pi^\varnothing)$.

For instance, the second law states that $\Sigma(p \rightarrow \pi^\varnothing) \geq 0$, so

$$W(p \rightarrow \pi^\varnothing) \leq D(p \parallel \pi^\varnothing). \quad (98)$$

We can also derive a stronger bound by exploiting coarse-graining constraints. For the coarse graining discussed above, Eq. (92) implies that $\Sigma(p \rightarrow \pi^\varnothing) \geq D(p_{X_2} \parallel \pi_{X_2})$, which gives the bound

$$\begin{aligned} W(p \rightarrow \pi^\varnothing) &\leq D(p \parallel \pi^\varnothing) - D(p_{X_2} \parallel \pi_{X_2}) \\ &= D(p_{X_1|X_2} \parallel \pi_{X_1|X_2}^\varnothing), \end{aligned} \quad (99)$$

where we've used the chain rule of KL divergence.

We can also bound EP and work using other kinds of constraints. For instance, the energy functions in Eq. (93) have no interaction terms between x_1 and x_2 and, therefore, obey modularity constraints for the decomposition $\mathcal{C} = \{\{X_1\}, \{X_2\}\}$ (see the analysis in Sec. VIA). We can then bound EP and work using the operator $\phi_{\mathcal{C}}$, as defined above in Eq. (70). In particular, using Theorem 2, we have that

$$\begin{aligned} \Sigma[p \rightarrow \pi^\varnothing] &= D[p \parallel \phi_{\mathcal{C}}(p)] + \Sigma[\phi_{\mathcal{C}}(p) \rightarrow \pi^\varnothing] \\ &\geq D[p \parallel \phi_{\mathcal{C}}(p)], \end{aligned} \quad (100)$$

which implies the extractable work bound

$$\begin{aligned} W(p \rightarrow \pi^\varnothing) &\leq D(p \parallel \pi^\varnothing) - D[p \parallel \phi_{\mathcal{C}}(p)] \\ &= D[\phi_{\mathcal{C}}(p) \parallel \pi^\varnothing]. \end{aligned} \quad (101)$$

Finally, we can also combine modularity and coarse-grained constraints. The coarse-grained constraints imply that $\Sigma[\phi_{\mathcal{C}}(p) \rightarrow \pi^\varnothing] \geq D[\phi_{\mathcal{C}}(p)_{X_2} \parallel \pi_{X_2}]$ by Eq. (92). Plugging into Eq. (100) gives the bound

$$\Sigma[p \rightarrow \pi^\varnothing] \geq D[p \parallel \phi_{\mathcal{C}}(p)] + D[\phi_{\mathcal{C}}(p)_{X_2} \parallel \pi_{X_2}], \quad (102)$$

resulting in the extractable work bound

$$W[p \rightarrow \pi^\varnothing] \leq D[\phi_{\mathcal{C}}(p)_{X_1|X_2} \parallel \pi_{X_1|X_2}^\varnothing], \quad (103)$$

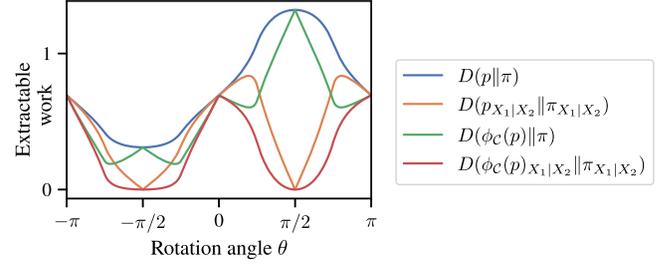


FIG. 14. Szilard box with gravity: bounds on extractable work as a function of θ , as derived from the second law [in blue, Eq. (98)], coarse-grained constraints [in orange, Eq. (99)], modularity constraints [in green, Eq. (101)], and a combination of modularity+coarse-grained constraints [in red, Eq. (103)].

which follows from combining Eqs. (97) and (102), and then using the Pythagorean theorem and the chain rule for KL divergence.

We now illustrate these bounds using a concrete set of initial distributions. Imagine that the initial distribution p is the equilibrium distribution π^\varnothing restricted to half the box, as determined by a rotated separating line at some angle $\theta \in [-\pi, \pi]$:

$$p_\theta(x_1, x_2) = \frac{1}{2} \pi^\varnothing(x_1, x_2) \Theta(x_2 \sin \theta - x_1 \cos \theta). \quad (104)$$

[Compare to Eq. (49), for the Szilard box without gravity.] For these initial distributions and gravity parameter $\kappa = 1$, we plot the four extractable work bounds derived above [Eqs. (98), (99), (101), and (103)], as a function of θ in Fig. 14 (values are calculated numerically). Note that, unlike the results presented in Figs. 6 and 9, the plots are no longer symmetric under the transformation $\theta \mapsto -\theta$. This result arises because gravity breaks the vertical reflection symmetry, so the nonequilibrium free energy of a distribution concentrated on the top half of the box ($\theta = \pi/2$) is greater than the nonequilibrium free energy of a distribution concentrated on the bottom half of the box ($\theta = -\pi/2$). It can also be seen that work bounds derived from coarse-grained constraints [Eq. (99) (orange)] can be either weaker or stronger than the work bounds derived from modularity constraints [Eq. (101) (green)], depending on the value of θ . For all θ , however, the work bound derived by combining both constraints [Eq. (103) (red)] is stronger than the work bound derived from either constraint individually.

VIII. RELEVANT LITERATURE

In previous work on the general topic of thermodynamic bounds under constraints, Wilming, Gallego, and Eisert [59] consider how extractable work depends on constraints on the Hamiltonian, given a quantum system coupled to a finite-sized heat bath. That paper derives an upper bound on the work that could be extracted by

carrying out a physical process which consists of sequences of (i) unitary transformations of the system and bath and (ii) total relaxations of the system to some equilibrium Gibbs state (see also a similar setup for closed systems in Ref. [60]). Building on Ref. [59], Ref. [44] analyzes the efficiency of a heat engine coupled to two baths and subject to “local control” constraints (i.e., a many-particle system where local Hamiltonians can be changed but the interaction Hamiltonians cannot). In contrast to these works, we consider a classical system coupled to idealized reservoir(s). We then derive bounds on EP and work for a much broader set of protocols.

At a high level, our approach complements previous research on the relationship between EP, extractable work, and different aspects of the driving protocol, such as temporal duration [61–67], stochasticity of control parameters [68], nonidealized work reservoirs [69], cyclic protocols [67,70], the presence of additional conservation laws [71], and the design of “optimal protocols” [72–74].

There is also previous work related to our analysis of thermodynamics of information under constraints in Sec. IV. Reference [75] recently analyzed the thermodynamics of feedback control under a somewhat different formulation of constraints [76]. In this work, we analyze the thermodynamics of information for a broader set of constraints. It is not immediately clear how the framework in Ref. [75] compares to ours, or whether it could be applied to the examples considered in this paper, although such a comparison is an interesting direction for future work.

Some of our results concerning work extraction under modularity constraints in Sec. VI have appeared in prior literature. Equation (66) is derived in Ref. [46] for the special case of an isothermal processes with two nonoverlapping subsystems, where one of the subsystems is held fixed. For the more general case of an arbitrary discrete-state system coupled to one or more reservoirs which have rate matrices as in Eq. (61), Eq. (66) was also previously derived in Refs. [19,77], while Eq. (67) was previously derived in Refs. [19,77,78]. Decompositions with overlap were previously considered in Refs. [79,80]. In addition, Example 1 in Ref. [81] can be used to derive the first inequality Eq. (65) for discrete-state systems [82].

Those papers also derive some results that are more general than the ones derived here, in that they apply even if the overlap changes state. Our paper goes beyond this previous work, though, to include continuous-state systems and to derive inequalities such as $D[p||\phi_C(p)] - D[p'||\phi_C(p')] \geq 0$, albeit for the more restricted scenario where the overlap does not change state.

Some of our results concerning work extraction under symmetry constraints, presented in Sec. V, appear in previous work on quantum thermodynamics. For a finite-state quantum system coupled to a work reservoir and heat bath, Vaccaro *et al.* [23] investigate how much work can be

extracted by bringing some initial quantum state ρ to a maximally mixed state, with a uniform initial and final Hamiltonian, using discrete-time operations that commute with the action of some symmetry group \mathcal{G} . It is shown that the work that can be extracted from ρ under such transformations is equal to the work that can be extracted from the (quantum) twirling $\phi_{\mathcal{G}}(\rho)$, analogous to Eq. (24) for symmetry constraints. This research also derives an operational measure of asymmetry that is the quantum equivalent of $D[p||\phi_{\mathcal{G}}(p)]$ and shows that asymmetry can only decrease under operations that commute with \mathcal{G} . Janzing [22] extends Ref. [23] to consider arbitrary Hamiltonians, in the process deriving analogs of our decomposition of free energy [Eq. (21)] for the special case of the twirling operator $\phi_{\mathcal{G}}$. A similar decomposition of free energy into coherent and incoherent components has recently appeared in Refs. [83,84] (which is a special case of the result in Ref. [22], since a decohering map is a twirling operator [85]). Finally, the idea of probability distributions that are invariant under symmetry groups, as well as a version of the twirling operator $\phi_{\mathcal{G}}$, is a topic of research in probability and statistics; for details, see Chap. 3 in Ref. [38].

While our approach is restricted to classical systems, in some respects our results for symmetry constraints are more general than this earlier work, since they hold for arbitrary (discrete and/or uncountably infinite) state spaces and for systems coupled to more than one reservoir (see Sec. IX). Moreover, for Fokker-Planck dynamics, we derive simple conditions for symmetry constraints stated in terms of the energy functions, which makes these results applicable to a large set of problems in stochastic thermodynamics and biophysics.

More fundamentally, one of the ways in which we go beyond previous literature on symmetry and modularity constraints is by providing a unified mathematical framework that applies to a broad set of constraints, including symmetry, modularity, and coarse-grained constraints (as well as their combinations) as special cases. A key idea in our framework is that the information-geometric Pythagorean identity [Eq. (14)] is the essential property that allows an operator ϕ to uncover the thermodynamically accessible part of any distribution p (assuming also that ϕ commutes with the dynamics). The Pythagorean identity is satisfied by many ϕ , including both linear operators such as twirling operators $\phi_{\mathcal{G}}$ and nonlinear operators such as modular decomposition operators ϕ_C . We believe this idea can be extended to the quantum domain, though we leave this extension for future work.

Finally, our approach is also related to “resource theories,” which are an active area of research in various areas of quantum physics [86], including quantum thermodynamics [59,87–91]. A resource theory quantifies a physical resource in an operational way, in terms of what transformations are possible when the resource is available. Most resource theories are based on a common set of formal elements,

such as a *resource quantifier* (a real-valued function that measures the amount of a resource), a set of *free states* (statistical states that lack the resource), and *free operations* (transformations between statistical states that do not increase the amount of resource). In fact, some previous work on symmetry constraints in quantum thermodynamics [22,23] can be seen as part of a broader literature on the resource theory of asymmetry [92–94].

Our approach has similar operational motivations as resource theories; for example, we define accessible free energy in an operational way, as a quantity that governs extractable work under protocol constraints. Moreover, many elements of our framework are analogous to elements of the resource theory framework: The set of allowed generators (which we call Λ) plays the role of the free operations, the image of the operator ϕ plays the role of the set of free states, and the KL divergence $D[p||\phi(p)]$ serves as the resource quantifier. In addition, the commutativity relation Eq. (16) (see Sec. III) has recently appeared in work on so-called resource-destroying maps [95]. However, unlike most resource theories, our focus is on the thermodynamics of classical systems modeled as driven continuous-time open systems. Further exploration of the connection between our approach and resource theories is left for future work.

IX. DISCUSSION

In this paper, we analyzed the EP and work incurred by a driving protocol that carries out some transformation $p \rightarrow p'$ while subject to constraints on the set of available generators. We constructed a general framework that allowed us derive several decompositions and bounds on EP and extractable work and demonstrated that this framework has implications for the thermodynamics of feedback control under constraints. Finally, we used our framework to analyze three broad classes of protocol constraints: reflecting symmetry, modularity, and coarse graining.

Note that our bounds on EP and extractable work, such as Eqs. (18) and (25), are expressed in terms of state functions; i.e., they depend only on the initial and final distributions p and p' and not on the path that the system takes in going from p to p' . In general, it may be possible to derive other bounds on work and EP that are not written in this form, which may be tighter. Nonetheless, bounds written in terms of state functions have some important advantages. In particular, they allow one to quantify the inherent “thermodynamic value” (in terms of EP and work) of a distribution p relative to a set of available generators, irrespective of what protocol brought the system there or what future protocols that system may undergo (as long as those protocols obey the relevant constraints).

For simplicity, our results were derived for isothermal protocols, where the system is coupled to a single heat

bath at a constant inverse temperature β and obeys local detailed balance (LDB). Nonetheless, many of our results continue to hold for more general protocols, in which the system is coupled to any number of thermodynamic reservoirs and/or violates LDB. For a general protocol, our EP rate in Eq. (11) refers to the so-called *non-adiabatic EP rate* [25,28,96], which is a non-negative quantity that reflects the contribution to EP that is due to the system being out of the stationary distribution. In the general case, our decompositions in Secs. III and III, as well as EP lower bounds in Eqs. (18) and (33), apply to nonadiabatic EP rather than overall EP. Importantly, the nonadiabatic EP rate is a lower bound on the overall EP rate whenever the stationary distribution of L is symmetric under conjugation of odd-parity variables [96], which holds in most cases of interest such as discrete-state master equations (which typically have no odd variables), overdamped dynamics (which have no odd variables), and many types of underdamped dynamics. In such cases, Eqs. (18) and (33) provide lower bounds not only on the nonadiabatic EP, but also on the overall EP, regardless of the number of coupled reservoirs or LDB. However, the relationship between work and EP in Eq. (2), as well as our bounds on work which make use of this relationship such as Eqs. (24) and (25), hold only for isothermal protocols. Note that our EP bound for closed coarse-grained dynamics [Eq. (87)] concerns the overall EP rate, *not* the nonadiabatic EP rate, even for nonisothermal protocols (see Appendix D 2 for details).

There are several possible directions for future research.

First, it remains an open question of whether our framework can also be used to analyze other classes of constraints, beyond the three classes (symmetry, modularity, and coarse graining) considered in this paper.

Second, our results point to a novel connection between entropy production, which plays a central role in non-equilibrium thermodynamics, and the Pythagorean identity in Eq. (14), which plays a central role in information geometry. This connection contributes to the growing number of existing results that demonstrate formal relationships between information geometry and nonequilibrium thermodynamics [97–102]. One direction for future work would be to extend the framework developed in this work for classical to quantum systems. In this extension, one would derive bounds on quantum work and EP by considering a quantum operator ϕ over density matrices which obeys quantum analogs of the Pythagorean identity in Eq. (14) (see p. 44 in Ref. [103]) and the commutativity relation in Eq. (16).

Finally, our results may also lead to some new treatments of foundational questions in thermodynamics. In stochastic thermodynamics, probability distributions over system states are usually interpreted in a “subjective” sense, in that the distribution p assigned to a system typically reflects *what one knows* about the system (for this reason,

this distribution changes once a measurement is made of the system's state [2]). At the same time, our results show that, for constrained driving protocols, one can often assign a different distribution to the system, $\phi(p)$, which reflects *what one can control* about the system. This assignment also leads to the difference between the overall nonequilibrium free energy, defined in terms of the distribution p , and the accessible free energy, defined in terms of the distribution $\phi(p)$. Note that thermodynamic entropy is often understood in an operational way, e.g., in terms of constrained macroscopic control, as has been previously discussed by Jaynes [104] and others. An interesting direction for future work would explore whether the distinction between the distributions p and $\phi(p)$ maps onto the distinction between (microscopic) statistical mechanical entropy and (macroscopic) thermodynamic entropy. In particular, one might ask whether this mapping can resolve some classic paradoxes concerning the relationship between statistical mechanical and thermodynamic entropy, such as the Gibbs paradox [104] (mixing of indistinguishable particles increases statistical mechanical entropy but not thermodynamic entropy) and Loschmidt's paradox (for an isolated Hamiltonian system, statistical mechanical entropy remains constant while the thermodynamic entropy can increase). This direction could also be related to a recent axiomatic treatment of thermodynamic entropy which has been developed within the framework of quantum resource theory [105].

ACKNOWLEDGMENTS

We thank Massimiliano Esposito and Henrik Wilming for helpful discussions. This research was supported by Grant No. FQXi-RFP-IPW-1912 from the Foundational Questions Institute and Fetzer Franklin Fund, a donor-advised fund of Silicon Valley Community Foundation. The authors thank the Santa Fe Institute for helping to support this research.

APPENDIX A: DERIVATIONS FOR SECS. III AND IV

1. Proofs of Secs. III and III

We first prove a few helpful lemmas.

Lemma 1.—If L obeys $e^L \phi(p) = \phi(e^L p)$ for all $p \in \mathcal{P}$, then L has a stationary distribution $\pi \in \text{img } \phi$.

Proof.—Let q be some stationary distribution of L . Then,

$$e^L \phi(q) = \phi(e^L q) = \phi(q). \quad (\text{A1})$$

Thus, $\phi(q) \in \text{img } \phi$ is stationary under L . ■

Lemma 2.—If $e^{\tau L} \phi(p) = \phi(e^{\tau L} p)$ for all $p \in \mathcal{P}$ and $\tau \geq 0$, then, for any $r, s \in \mathcal{P}$,

$$-\frac{d}{dt} D\{r(t) \parallel \phi[s(t)]\} \geq 0,$$

where $\partial_t r = Lr$ and $\partial_t s = Ls$.

Proof.—Expand the derivative as

$$\begin{aligned} & -\frac{d}{dt} D\{r(t) \parallel \phi[s(t)]\} \\ &= \lim_{\tau \rightarrow 0} \frac{1}{\tau} \{D[r \parallel \phi(s)] - D[e^{\tau L} r \parallel \phi(e^{\tau L} s)]\} \\ &= \lim_{\tau \rightarrow 0} \frac{1}{\tau} \{D[r \parallel \phi(s)] - D[e^{\tau L} r \parallel e^{\tau L} \phi(s)]\} \geq 0, \end{aligned}$$

where we first use the commutativity relation and then use the data-processing inequality for KL divergence [106]. ■

Lemma 3.—Consider a protocol $\{L(t) : t \in [0, 1]\}$ and an operator ϕ that obeys Eqs. (14) and (16). Then

$$\phi[p(t)] = \phi(p)(t),$$

where $p(t)$ is the distribution at time t given initial distribution p and $\phi(p)(t)$ is the distribution at time t given initial distribution $\phi(p)$.

Proof.—Using Lemma 2 with $r = \phi(p)(t)$ and $s = p(t)$,

$$\frac{d}{dt} D\{\phi(p)(t) \parallel \phi[p(t)]\} \leq 0. \quad (\text{A2})$$

Note that

$$D\{[\phi(p)](0) \parallel \phi[p(0)]\} = D[\phi(p) \parallel \phi(p)] = 0$$

and that $D\{\phi(p)(t) \parallel \phi[p(t)]\} \geq 0$ for all t by non-negativity of KL divergence. Combined with Eq. (A2), this result implies $D\{\phi(p)(t) \parallel \phi[p(t)]\} = 0$ for all t , and, therefore, $\phi(p)(t) = \phi[p(t)]$ (see Theorem 8.6.1 in Ref. [107]). ■

We are now ready to prove Theorems 1 and 2. Note that, in the proof of Theorem 1, we make the assumption that there is some stationary distribution π^L of L such that $D(p \parallel \pi^L) < \infty$, and similarly, in Theorem 2, we make the assumption that $D[p(t) \parallel \pi^{L(t)}] < \infty$ at all $t \in [0, 1]$. These are weak and physically realistic assumptions, which essentially mean that we restrict our attention to distributions with finite nonequilibrium free energy [see Eq. (20)].

In addition, in these proofs, we use that the EP rate incurred by distribution p under the generator L with stationary distribution π can be written as

$$\dot{\Sigma}(p, L) = \lim_{\tau \rightarrow 0} \frac{1}{\tau} [D(p \parallel \pi) - D(e^{\tau L} p \parallel \pi)]. \quad (\text{A3})$$

This expression can be derived from Eq. (11), by noting that the KL divergence can be written as

$$D(p||\pi) = -S(p) - \mathbb{E}_p[\ln \pi], \quad (\text{A4})$$

where \mathbb{E}_p indicates expectation under the distribution p , and then using that

$$-\sum_x \partial_t p_x(t) \ln p_x = \lim_{\tau \rightarrow 0} \frac{1}{\tau} [S(e^{\tau L} p) - S(p)], \quad (\text{A5})$$

$$\sum_x \partial_t p_x(t) \ln \pi_x = \lim_{\tau \rightarrow 0} \frac{1}{\tau} \{ \mathbb{E}_{e^{\tau L} p}[\ln \pi] - \mathbb{E}_p[\ln \pi] \}, \quad (\text{A6})$$

where $\partial_t p_x(t)$ is defined as in Eq. (10). (As usual, summations should be replaced by integrals for continuous-state systems.)

Proof of Theorem 1.—Consider a generator L with a stationary distribution π and some distribution $p \in \mathcal{P}$ such that $D(p||\pi) < \infty$. By Lemma 1, $\phi(\pi) \in \text{img } \phi$ is also a stationary distribution of L . If L has a unique stationary distribution, then $\pi = \phi(\pi)$ and so $\pi \in \text{img } \phi$; otherwise, as long as $D[p||\phi(\pi)] < \infty$ (see Ref. [24]), we can assume that $\phi(\pi) = \pi$ in Eq. (A3). Then, assuming that $\pi \in \text{img } \phi$, we rewrite the term in the brackets in Eq. (A3) as

$$\begin{aligned} & D[p||\phi(p)] + D[\phi(p)||\pi] \\ & - D[e^{\tau L} p||\phi(e^{\tau L} p)] - D[\phi(e^{\tau L} p)||\pi] \\ & = D[p||\phi(p)] - D[e^{\tau L} p||\phi(e^{\tau L} p)] \\ & + D[\phi(p)||\pi] - D[\phi(e^{\tau L} p)||\pi] \\ & = D[p||\phi(p)] - D[e^{\tau L} p||\phi(e^{\tau L} p)] \\ & + D[\phi(p)||\pi] - D[e^{\tau L} \phi(p)||\pi], \end{aligned}$$

where we use the Pythagorean identity of Eq. (14), rearranged, and then use the commutativity relation of Eq. (16). Plugging into Eq. (A3) gives

$$\begin{aligned} \dot{\Sigma}(p, L) &= \lim_{\tau \rightarrow 0} \frac{1}{\tau} \{ D[p||\phi(p)] - D[e^{\tau L} p||\phi(e^{\tau L} p)] \} \\ &+ \lim_{\tau \rightarrow 0} \frac{1}{\tau} \{ D[\phi(p)||\pi] - D[e^{\tau L} \phi(p)||\pi] \} \\ &= -\frac{d}{dt} D\{p(t)||\phi[p(t)]\} + \dot{\Sigma}[\phi(p), L]. \end{aligned}$$

The non-negativity of $-(d/dt)D\{p(t)||\phi[p(t)]\}$ follows by taking $r = s = p$ in Lemma 2. ■

Proof of Theorem 2.—Using Eq. (12) and Theorem 1, write

$$\begin{aligned} \Sigma(p \rightarrow p') &= \int_0^1 \dot{\Sigma}[p(t), L(t)] dt \\ &= -\int_0^1 \frac{d}{dt} D\{p(t)||\phi[p(t)]\} dt \\ &+ \int_0^1 \dot{\Sigma}\{\phi[p(t)], L(t)\} dt. \end{aligned}$$

Both integrals have a simple expression. First, by the fundamental theorem of calculus,

$$-\int_0^1 \frac{d}{dt} D\{p(t)||\phi[p(t)]\} dt = D[p||\phi(p)] - D[p'||\phi(p')].$$

This expression is non-negative, since $-(d/dt)D\{p(t)||\phi[p(t)]\} \geq 0$ by Lemma 2. Second, using Lemma 3,

$$\begin{aligned} \int_0^1 \dot{\Sigma}\{\phi[p(t)], L(t)\} dt &= \int_0^1 \dot{\Sigma}\{\phi(p)(t), L(t)\} dt \\ &= \Sigma[\phi(p) \rightarrow \phi(p')]. \end{aligned}$$

■

2. Trajectory-level version of Eq. (19)

Stochastic thermodynamics shows that thermodynamic properties of physical processes (such as heat, work, and EP) can be defined as stochastically fluctuating quantities at the level of individual trajectories. We first briefly review the basic concepts of stochastic thermodynamics (for more details, the reader should consult Refs. [33,108–110]).

Let $\mathbf{x} = (x, \dots, x')$ indicate a continuous-time trajectory of system states \mathbf{x} over time interval $t \in [0, 1]$, where x and x' indicate the initial and final system states, respectively, and let $P(\mathbf{x}|x)$ indicate the conditional probability of observing trajectory \mathbf{x} given initial state x . For a given initial distribution $p(x)$, the probability of observing trajectory \mathbf{x} is then given by $p(\mathbf{x}) = p(x)P(\mathbf{x}|x)$, and the corresponding final distribution is given by $p'(x') = \int P(x'|x)p(x)dx$. In addition, let $\tilde{P}(\tilde{\mathbf{x}}|x')$ indicate the conditional probability of observing the time-reversed and trajectory $\tilde{\mathbf{x}} = (x', \dots, x)$ given the final state x' under a “time-reversed” driving protocol [33].

Trajectory-level EP is then defined in terms of the asymmetry between forward and reversed trajectory probabilities:

$$\sigma_p(\mathbf{x}) = \ln p(\mathbf{x}) - \ln p'(\mathbf{x}') + \ln \frac{P(\mathbf{x}|x)}{\tilde{P}(\tilde{\mathbf{x}}|x')}, \quad (\text{A7})$$

which is sometimes referred to as a “detailed fluctuation theorem.” (The above expression should be slightly modified in the presence of odd-parity variables such as momentum, though in a way which does not change our derivations; see Ref. [111].) The expectation of trajectory-level EP across all trajectories is equal to the standard expression for integrated EP as used in the main text:

$$\langle \sigma_p(\mathbf{x}) \rangle = \Sigma(p \rightarrow p'), \quad (\text{A8})$$

where $\langle \cdot \rangle$ refers to expectations under the trajectory distribution $p(\mathbf{x})$. Furthermore, by a simple manipulation, the detailed fluctuation theorem in Eq. (A7) leads to the following integral fluctuation theorem for EP:

$$\begin{aligned} \langle e^{-\sigma_p} \rangle &= \int_{p(x)>0} p(x) P(\mathbf{x}|x) \frac{p'(x') \tilde{P}(\tilde{\mathbf{x}}|x')}{p(x) P(\mathbf{x}|x)} D\mathbf{x} \\ &= \int_{p(x)>0} p'(x') \tilde{P}(\tilde{\mathbf{x}}|x') D\mathbf{x} = \gamma, \end{aligned} \quad (\text{A9})$$

where $\int \cdot D\mathbf{x}$ is the path integral. In this result, $\gamma \in (0, 1]$ reflects the absolute irreversibility of the process under initial distribution p [112]. When p has full support, $\gamma = 1$, giving the standard integral fluctuation theorem $\langle e^{-\sigma_p} \rangle = 1$.

Now consider the extra trajectory-level EP incurred by some trajectory \mathbf{x} on initial distribution p , additional to the trajectory-level EP incurred by the same trajectory on initial distribution $\phi(p)$:

$$m(\mathbf{x}) := \sigma_p(\mathbf{x}) - \sigma_{\phi(p)}(\mathbf{x}) \quad (\text{A10})$$

$$= \ln \frac{p(x)}{\phi(p)(x)} - \ln \frac{p'(x')}{\phi(p')(x')} \quad (\text{A11})$$

$$= \ln \frac{p(x)}{\phi(p)(x)} - \ln \frac{p'(x')}{\phi(p')(x')} \quad (\text{A12})$$

where in Eq. (A11) we use that the last term in Eq. (A7) cancels (as it does not depend on the initial or final distributions) and in Eq. (A10) we use that $\phi(p)' = \phi(p')$ by Lemma 3. Equation (A10) appears in the main text as Eq. (30). It is easy to verify that $m(\mathbf{x})$ agrees in expectation with the contraction of KL divergence between p and $\phi(p)$:

$$\langle m \rangle = D[p \parallel \phi(p)] - D[p' \parallel \phi(p')], \quad (\text{A13})$$

where, as before, $\langle \cdot \rangle$ refers to expectations under the trajectory distribution $p(\mathbf{x})$. Given Theorem 2, Eq. (A13) implies that the expectation $m(\mathbf{x})$ is also equal to the extra total EP incurred by initial distribution p rather than the accessible distribution $\phi(p)$:

$$\langle m \rangle = \Sigma(p \rightarrow p') - \Sigma[\phi(p) \rightarrow \phi(p')]. \quad (\text{A14})$$

In Ref. [34], it is shown that $m(\mathbf{x})$ obeys a fluctuation theorem (see also Ref. [35]). We rederive the relevant results here. First, a simple rearrangement of Eq. (A11) gives the following detailed fluctuation theorem:

$$m(\mathbf{x}) := \ln \frac{p(x)}{p'(x')} + \ln \frac{P(\mathbf{x}|x)}{Q(\tilde{\mathbf{x}}|x')}, \quad (\text{A15})$$

where the conditional distribution $Q(\tilde{\mathbf{x}}|x')$ is given by

$$Q(\tilde{\mathbf{x}}|x') := \frac{P(\mathbf{x}|x)\phi(p)(x)}{\phi(p)'(x')}.$$

In words, $Q(\tilde{\mathbf{x}}|x')$ is the Bayesian posterior probability of the trajectory \mathbf{x} given final state x' , when the process begins on initial distribution $\phi(p)$. A similar derivation as in Eq. (A9) shows that m obeys an integral fluctuation theorem:

$$\langle e^{-m} \rangle = \int_{p(x)>0} p'(x') Q(\tilde{\mathbf{x}}|x') D\mathbf{x} = \chi. \quad (\text{A16})$$

Here, $\chi \in (0, 1]$ indicates the absolute irreversibility of the process on initial distribution p relative to initial distribution $\phi(p)$. χ is equal to 1 when p and $\phi(p)$ have the same support, which then leads to a standard integral fluctuation theorem $\langle e^{-m} \rangle = 1$.

Importantly, Eq. (A16) implies that the probability that the trajectory-level EP on initial distribution p is ξ less than the trajectory-level EP on initial distribution $\phi(p)$ is exponentially suppressed:

$$P[\sigma_p < \sigma_{\phi(p)} - \xi] \stackrel{(a)}{=} P[m < -\xi] \stackrel{(b)}{\leq} \chi e^{-\xi} \stackrel{(c)}{\leq} e^{-\xi}. \quad (\text{A17})$$

Here, (a) uses the definition of $m(\mathbf{x})$, and (b) uses a standard derivation in stochastic thermodynamics (see Ref. [113] or the appendix in Ref. [34]), while (c) uses that $\chi \in (0, 1]$.

APPENDIX B: SYMMETRY CONSTRAINTS

1. $\phi_{\mathcal{G}}$ obeys the Pythagorean identity [Eq. (14)]

In the following derivations, all integrals should be understood in the Lebesgue sense. For discrete state systems, integrals over X can be replaced by summations.

The state space X is assumed to be Borel measurable. Similarly, we assume that the action of the group \mathcal{G} [i.e., the function $\mathcal{G} \times X \rightarrow X: (g, x) \mapsto g(x)$] is Borel measurable. Note that these assumptions imply that, for any probability distribution $p \in \mathcal{P}$, the function $(g, x) \mapsto p[g(x)]$ is measurable, since it is the composition of two Borel measurable functions: $(g, x) \mapsto g(x)$ and $x \mapsto p(x)$.

We begin with a few intermediate results.

Lemma 4.—For any $p \in \mathcal{P}$, $g \in \mathcal{G}$, and $x \in X$,

$$\phi_{\mathcal{G}}(p)(x) = \phi_{\mathcal{G}}(p)[g(x)].$$

Proof.—Using the definition of $\phi_{\mathcal{G}}$ in Eq. (42), write

$$\begin{aligned} \phi_{\mathcal{G}}(p)[g(x)] &= \int_{\mathcal{G}} p\{g'[g(x)]\} d\mu(g') \\ &= \int_{\mathcal{G}} p[g'(x)] d\mu(g') = \phi_{\mathcal{G}}(p)(x), \end{aligned}$$

where we perform a change of variables $x \mapsto g^{-1}(x)$ and use the invariance properties \mathcal{G} and the Haar measure μ . ■

Lemma 5.—For any $p \in \mathcal{P}$, measurable set $\Omega \subseteq X$, and function $f: X \rightarrow \mathbb{R}$,

$$\int_{\Omega} p(x)f(x) = \int_{\Omega} \phi_{\mathcal{G}}(p)(x)f(x)dx \quad (\text{B1})$$

if the following three conditions hold: (i) $g(\Omega) = \Omega$ for all $g \in \mathcal{G}$, (ii) $f(x) = f[g(x)]$ for all $x \in X$ and $g \in \mathcal{G}$, and (iii) either $|\int_{\Omega} p(x)f(x)dx| < \infty$ or f is measurable and non-negative.

Proof.—To begin, write the left-hand side of Eq. (B1) as

$$\begin{aligned} \int_{\Omega} p(x)f(x)dx &= \int_{\mathcal{G}} \left[\int_{\Omega} p(x)f(x)dx \right] d\mu(g) \\ &= \int_{\mathcal{G}} \left[\int_{g^{-1}(\Omega)} p[g(x)]f[g(x)]dx \right] d\mu(g) \\ &= \int_{\mathcal{G}} \left[\int_{\Omega} p[g(x)]f(x)dx \right] d\mu(g). \end{aligned} \quad (\text{B2})$$

In the second line, we substitute $x \mapsto g(x)$ within each inner integral while using that each g is a rigid transformation (so the absolute value of its Jacobian is 1). In the last line, we use conditions (i) and (ii).

We now show that we can exchange the order of integrals in Eq. (B2) using condition (iii) and Tonelli's theorem. First, if f is measurable and non-negative, then the function $x \mapsto p[g(x)]f(x)$ is non-negative and measurable (since it is a product of two non-negative measurable functions), so the integrals can be exchanged by Theorem 3.7.7 in Ref. [114]. Alternatively, assume that $|\int_{\Omega} p(x)f(x)dx| < \infty$, which means that the function $x \mapsto p(x)f(x)$ is integrable. Then,

$$\begin{aligned} \infty &> \int_{\Omega} p(x)|f(x)|dx \\ &= \int_{\mathcal{G}} \left[\int_{\Omega} p(x)|f(x)|dx \right] d\mu(g) \quad (\text{B3}) \\ &= \int_{\mathcal{G}} \left[\int_{g^{-1}(\Omega)} p[g(x)]|f[g(x)]|dx \right] d\mu(g) \\ &= \int_{\mathcal{G}} \left[\int_{\Omega} p[g(x)]|f(x)|dx \right] d\mu(g), \end{aligned} \quad (\text{B4})$$

where the first line follows from definition of Lebesgue integrability, while the rest follows from the same steps as Eq. (B2). Given Eq. (B4), the function $(g, x) \mapsto p[g(x)]f(x)$ must be integrable, which again allows us to exchange the order of the integrals in Eq. (B2) (see Theorem 3.7.8 in Ref. [114]).

We then derive our result by rewriting Eq. (B2) as

$$\begin{aligned} \int_{\Omega} p(x)f(x)dx &= \int_{\Omega} \left[\int_{\mathcal{G}} p[g(x)]f(x)d\mu(g) \right] dx \\ &= \int_{\Omega} \phi_{\mathcal{G}}(p)(x)f(x)dx, \end{aligned}$$

where we use the definition of $\phi_{\mathcal{G}}$. \blacksquare

Finally, we prove that $\phi_{\mathcal{G}}$ obeys the Pythagorean identity.

Proposition 1.—For any $p, q \in \mathcal{P}$ such that $D[p\|\phi_{\mathcal{G}}(q)] < \infty$,

$$D[p\|\phi_{\mathcal{G}}(q)] = D[p\|\phi_{\mathcal{G}}(p)] + D[\phi_{\mathcal{G}}(p)\|\phi_{\mathcal{G}}(q)]. \quad (\text{B5})$$

Proof.—For any $p \in \mathcal{P}$, we indicate the support set as $\text{supp } p = \{x \in X: p(x) > 0\}$. We first prove that

$$\text{supp } p \subseteq \text{supp } \phi_{\mathcal{G}}(p) \subseteq \text{supp } \phi_{\mathcal{G}}(q). \quad (\text{B6})$$

By the definition of $\phi_{\mathcal{G}}$ in Eq. (42), if $\phi_{\mathcal{G}}(p)(x) > 0$ for some $x \in X$, then $p[g(x)] > 0$ for that x and some $g \in \mathcal{G}$. In addition, the assumption that $D[p\|\phi_{\mathcal{G}}(q)] < \infty$ implies that $\text{supp } p \subseteq \text{supp } \phi_{\mathcal{G}}(q)$ [107] (except for a set of measure 0, which we can safely ignore). Combining these facts implies that, if $\phi_{\mathcal{G}}(p)(x) > 0$ for some x , then $\phi_{\mathcal{G}}(q)[g(x)] > 0$ for that x —and, therefore, also $\phi_{\mathcal{G}}(q)(x) > 0$, since $\phi_{\mathcal{G}}(q)$ is invariant under \mathcal{G} (Lemma 4). This result proves that $\text{supp } \phi_{\mathcal{G}}(p) \subseteq \text{supp } \phi_{\mathcal{G}}(q)$. Finally, by Lemmas 4 and 5,

$$\int_{\text{supp } \phi_{\mathcal{G}}(p)} p(x)dx = \int_{\text{supp } \phi_{\mathcal{G}}(p)} \phi_{\mathcal{G}}(p)(x)dx = 1,$$

which implies that $\text{supp } p \subseteq \text{supp } \phi_{\mathcal{G}}(p)$ (up to a set of measure 0).

Next, write the KL divergence on the left-hand side of Eq. (B5) as [see Eq. (8.58) in Ref. [107]]

$$\begin{aligned} D[p\|\phi_{\mathcal{G}}(q)] &= \int_{\text{supp } p} p(x) \ln \frac{p(x)}{\phi_{\mathcal{G}}(q)(x)} dx \\ &= D[p\|\phi_{\mathcal{G}}(p)] + \int_{\text{supp } p} p(x) \ln \frac{\phi_{\mathcal{G}}(p)(x)}{\phi_{\mathcal{G}}(q)(x)} dx \\ &= D[p\|\phi_{\mathcal{G}}(p)] + \int_{\text{supp } \phi_{\mathcal{G}}(p)} p(x) \ln \frac{\phi_{\mathcal{G}}(p)(x)}{\phi_{\mathcal{G}}(q)(x)} dx, \end{aligned} \quad (\text{B7})$$

where the last line uses Eq. (B6) [in particular, that $\text{supp } p \subseteq \text{supp } \phi_{\mathcal{G}}(p)$ and $p(x) \ln[\phi_{\mathcal{G}}(p)(x)/\phi_{\mathcal{G}}(q)(x)] = 0$ for $x \in \text{supp } \phi_{\mathcal{G}}(p) \setminus \text{supp } p$].

The integral in Eq. (B7) is bounded from above by $D[p\|\phi_{\mathcal{G}}(q)] < \infty$, since $D[p\|\phi_{\mathcal{G}}(p)] \geq 0$. We also show

that this integral is bounded from below. Note that $\phi_{\mathcal{G}}(p)(x)$ and $\phi_{\mathcal{G}}(q)(x)$ are both non-negative measurable functions, which follows from the fact that $x \mapsto p[g(x)]$ and $x \mapsto p[g(x)]$ are non-negative measurable functions, the definition of $\phi_{\mathcal{G}}$, and Tonelli's theorem (Theorem 3.7.7 in Ref. [114]). Thus, the function $x \mapsto [\phi_{\mathcal{G}}(q)(x)/\phi_{\mathcal{G}}(p)(x)]$ is also non-negative and measurable, letting us bound the integral in the following way:

$$\begin{aligned} & \int_{\text{supp } \phi_{\mathcal{G}}(p)} p(x) \ln \frac{\phi_{\mathcal{G}}(p)(x)}{\phi_{\mathcal{G}}(q)(x)} dx \\ & \geq -\ln \left[\int_{\text{supp } \phi_{\mathcal{G}}(p)} p(x) \frac{\phi_{\mathcal{G}}(q)(x)}{\phi_{\mathcal{G}}(p)(x)} dx \right] \\ & = -\ln \left[\int_{\text{supp } \phi_{\mathcal{G}}(p)} \phi_{\mathcal{G}}(p)(x) \frac{\phi_{\mathcal{G}}(q)(x)}{\phi_{\mathcal{G}}(p)(x)} dx \right] \\ & = -\ln \left[\int_{\text{supp } \phi_{\mathcal{G}}(p)} \phi_{\mathcal{G}}(q)(x) dx \right] \geq -\ln 1 = 0. \end{aligned}$$

where in the second line we use Jensen's inequality, while in the third line we apply Lemma 5. Finally, we use Lemma 5 to rewrite the integral in Eq. (B7) as

$$\begin{aligned} & \int_{\text{supp } \phi_{\mathcal{G}}(p)} p(x) \ln \frac{\phi_{\mathcal{G}}(p)(x)}{\phi_{\mathcal{G}}(q)(x)} dx \\ & = \int_{\text{supp } \phi_{\mathcal{G}}(p)} \phi_{\mathcal{G}}(p)(x) \ln \frac{\phi_{\mathcal{G}}(p)(x)}{\phi_{\mathcal{G}}(q)(x)} dx = D[\phi_{\mathcal{G}}(p) \parallel \phi_{\mathcal{G}}(q)]. \end{aligned}$$

■

2. $\phi_{\mathcal{G}}$ obeys the commutativity relation [Eq. (16)]

It is easy to verify that Φ_g is a linear operator. It then follows that, if Φ_g commutes with the linear operator L , as in Eq. (38), then it also commutes with the exponential $e^{\tau L} = \sum_k (1/k!) \tau^k L^k$. We then have

$$\begin{aligned} e^{\tau L} \phi_{\mathcal{G}}(p) & = e^{\tau L} \int \Phi_g p d\mu(g) \\ & = \int e^{\tau L} \Phi_g p d\mu(g) \\ & = \int \Phi_g e^{\tau L} p d\mu(g) \\ & = \phi_{\mathcal{G}}(e^{\tau L} p), \end{aligned}$$

where in the second line we exchange the bounded operator $e^{\tau L}$ and the (Bochner) integral and in the third line we use that Φ_g and $e^{\tau L}$ commute.

3. Derivation of Eq. (38) from Eqs. (39) and (41)

Consider some $f: X \rightarrow \mathbb{R}$ and a continuous-state master equation L such that

$$[Lf](x) = \int [L_{xx'} f(x') - L_{x'x} f(x)] dx'. \quad (\text{B8})$$

[The derivation for discrete-state master equations, as in Eq. (10), is the same, but with integrals replaced with summations.] Then,

$$\begin{aligned} [\Phi_g Lf](x) & = [Lf][g(x)] \\ & = \int \{L_{g(x)x'} f(x') - L_{x'g(x)} f[g(x)]\} dx' \end{aligned} \quad (\text{B9})$$

$$= \int \{L_{g(x)g(x')} f[g(x')] - L_{g(x')g(x)} f[g(x)]\} dx' \quad (\text{B10})$$

$$= \int \{L_{xx'} f[g(x')] - L_{x'x} f[g(x)]\} dx' \quad (\text{B11})$$

$$= \int \{L_{xx'} [\Phi_g f](y) - L_{x'x} [\Phi_g f](x)\} dx' \quad (\text{B12})$$

$$= [L\Phi_g f](x), \quad (\text{B13})$$

which implies $\Phi_g L = L\Phi_g$ [Eq. (38)]. Here, we use the definition of Φ_g in the first line and Eq. (B8) in (B9). In Eq. (B10), we use the variable substitution $x' \mapsto g(x')$, along with the fact that g is volume preserving. In Eq. (B11), we use Eq. (39).

Next, we show that Eq. (41) is sufficient for Eq. (38) to hold, assuming that all $g \in \mathcal{G}$ are rigid transformations and the $L \in \Lambda$ refer to Fokker-Planck equations of the form Eq. (40). First, given some (sufficiently smooth) function $f: X \rightarrow \mathbb{R}$, write Eq. (40) as

$$\partial_t f = Lf = \nabla \cdot [(\nabla E)f] + \beta^{-1} \Delta f. \quad (\text{B14})$$

For any $g \in \mathcal{G}$, write the diffusion term in Eq. (B14) as

$$\Delta f = \Delta(\Phi_g f \circ g^{-1}) = \Delta(\Phi_g f) \circ g^{-1}, \quad (\text{B15})$$

where we use the identity $f = \Phi_{g^{-1}} \Phi_g f = \Phi_g f \circ g^{-1}$ and that the Laplace operator commutes with rigid transformations. Now consider the drift term in Eq. (B14). Using the product rule,

$$\nabla \cdot [(\nabla E)f] = (\nabla f)^T (\nabla E) + f \Delta E. \quad (\text{B16})$$

We can rewrite the second term above as

$$\begin{aligned} f \Delta E & = (\Phi_g f \circ g^{-1}) \Delta E \\ & = (\Phi_g f \circ g^{-1}) \Delta(E \circ g^{-1}) \\ & = (\Phi_g f \circ g^{-1}) [(\Delta E) \circ g^{-1}] \\ & = [(\Phi_g f)(\Delta E)] \circ g^{-1}, \end{aligned} \quad (\text{B17})$$

where we use $f = \Phi_g f \circ g^{-1}$, the invariance of E under \mathcal{G} [Eq. (41)], and in the third line that the Laplace operator commutes with rigid transformations. Now consider the first term on the right-hand side of Eq. (B16):

$$\begin{aligned} (\nabla f)^T (\nabla E) &= [\nabla(\Phi_g f \circ g^{-1})^T \nabla(E \circ g^{-1})] \\ &= \{J^T [\nabla(\Phi_g f) \circ g^{-1}]\}^T \{J^T [(\nabla E) \circ g^{-1}]\} \\ &= [\nabla(\Phi_g f) \circ g^{-1}]^T J J^T [(\nabla E) \circ g^{-1}] \\ &= [\nabla(\Phi_g f) \circ g^{-1}]^T [(\nabla E) \circ g^{-1}] \\ &= [\nabla(\Phi_g f)^T (\nabla E)] \circ g^{-1}, \end{aligned} \quad (\text{B18})$$

where J indicates the Jacobian of g^{-1} . In the first line, we again use the identity $f = \Phi_g f \circ g^{-1}$ and the invariance of E under \mathcal{G} , in the second line we use the chain rule, and in the fourth line we use that $J J^T = I$ for rigid transformations. Plugging Eqs. (B17) and (B18) back into Eq. (B16) and rearranging gives

$$\nabla \cdot [(\nabla E) f] = \nabla \cdot [(\nabla E) (\Phi_g f)] \circ g^{-1}. \quad (\text{B19})$$

Combined with Eqs. (B15) and (B14), this expression implies that $L f = (L \Phi_g f) \circ g^{-1}$ or, in other words, that

$$\Phi_g L f = L \Phi_g f.$$

4. Derivation of Eq. (43)

First, write the inaccessible information term in Eq. (35) as

$$\begin{aligned} D[p_{X|M} \| \phi_{\mathcal{G}}(p_{X|M})] &= \sum_m p(m) D[p_{X|M} \| \phi_{\mathcal{G}}(p_{X|M})] \\ &= \sum_m p(m, x) \ln \frac{p(x|m)}{\int p[g(x)|m] \mu g} \\ &= \sum_m p(m, x) \ln \frac{p(x) q(m|x) / p(m)}{\int p[g(x)] q[m|g(x)] / p_g(m) \mu(g)}, \end{aligned} \quad (\text{B20})$$

where we define $p(m) = \sum_x p(x) q(m|x)$ and $p_g(m) = \sum_x p[g(x)] q(m|x)$ and use the definition of $\phi_{\mathcal{G}}$ in Eq. (42). [Here, we assume for simplicity that both X and M are discrete valued; otherwise, the summations in Eq. (B20) should be replaced with integrals.]

Recall that we assume that p is invariant under \mathcal{G} , so $\phi_{\mathcal{G}}(p) = p$. By Lemma 4, $p(x) = p[g(x)]$ for all x and $g \in \mathcal{G}$, which, in turn, implies that $p(m) = p_g(m)$. Plugging into Eq. (B20) then gives

$$D[p_{X|M} \| \phi_{\mathcal{G}}(p_{X|M})] = \sum_m p(m, x) \ln \frac{q(m|x)}{\int q[m|g(x)] \mu(g)},$$

which appears in the main text as Eq. (43).

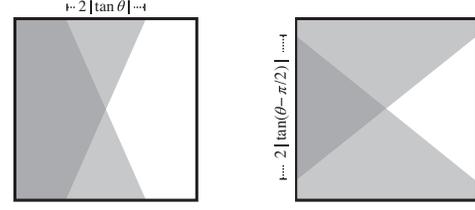


FIG. 15. The twirling $\phi_{\mathcal{C}}(p_{\theta})$ for two cases. Left: $|\theta| \in [(\pi/4), (3\pi/4)]$. Right: $\phi_{\mathcal{C}}(p_{\theta})$ for $|\theta| \in [-\pi, \pi] \setminus [(\pi/4), (3\pi/4)]$.

5. Example: Szilard box, derivation of Eq. (50)

We derive Eq. (50) using a simple geometric argument.

Consider the twirling of p_{θ} , as shown in Fig. 5(b). From the definition of $\phi_{\mathcal{G}}$ and Eq. (49), it is easy to see the following.

- (1) The dark gray areas in Fig. 5(b) [where both $p_{\theta}(x_1, x_2) = 1/2$ and $p_{\theta}(x_1, -x_2) = 1/2$] have probability density $\phi_{\mathcal{G}}(p_{\theta})(x_1, x_2) = 1/2$.
- (2) The light gray areas in Fig. 5(b) [where either $p_{\theta}(x_1, x_2) = 1/2$ or $p_{\theta}(x_1, -x_2) = 1/2$, but not both] have probability density $\phi_{\mathcal{G}}(p_{\theta})(x_1, x_2) = 1/4 = u(x_1, x_4)$.
- (3) The white areas in Fig. 5(b) [where $p_{\theta}(x_1, x_2) = 0$ and $p_{\theta}(x_1, -x_2) = 0$] have probability density $\phi_{\mathcal{G}}(p_{\theta})(x_1, x_2) = 0$.

Given these facts,

$$D[\phi_{\mathcal{G}}(p_{\theta}) \| u] = \ln 2 \cdot P_{\theta}, \quad (\text{B21})$$

where P_{θ} is the probability assigned by p to the dark gray areas [i.e., those (x_1, x_2) where $p_{\theta}(x_1, x_2) = 1/2 = p_{\theta}(x_1, -x_2) = 1/2$].

To calculate the value of P_{θ} , it suffices to consider two separate cases:

- (1) $|\theta| \in [-\pi, \pi] \setminus [(\pi/4), (3\pi/4)]$;
- (2) $|\theta| \in [(\pi/4), (3\pi/4)]$,

which are shown visually in Fig. 15. Using this figure and a bit of trigonometry, it can be shown that $P_{\theta} = 1 - \frac{1}{2} |\tan \theta|$ in the first case and $P_{\theta} = \frac{1}{2} |\tan(\theta - \pi/2)|$ in the second case. Combining these results with Eq. (B21) gives Eq. (50).

6. Example: Symmetry constraints on a discrete-state master equation

Here, we demonstrate our results on symmetry constraints using a simple finite-state system. The system contains n states, $x = \{0, \dots, n-1\}$. We consider a group generated by circular shifts, representing m -fold circular symmetry:

$$g(x) = x + n/m \pmod{n}. \quad (\text{B22})$$

Assume that the driving protocol obeys the following symmetry group at all $t \in [0, 1]$:

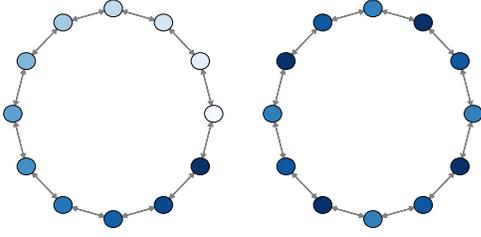


FIG. 16. A unicyclic master equation over 12 states with fourfold symmetry, as in Eq. (B23). Left: an initial distribution $p(x) \propto x$, which does not respect the fourfold symmetry. Right: the twirling $\phi_G(p)$, which is invariant to the symmetry. (Colors indicate relative probability assigned to each of the 12 states.) The extractable work depends on the accessible free energy in p , which is given by $F_E[\phi_G(p)]$.

$$L_{x'x}(t) = L_{g(x')g(x)}(t). \quad (\text{B23})$$

An example of such a master equation would be a unicyclic network, where the n states are arranged in a ring and transitions between nearest-neighbor states obey Eq. (B23). Such unicyclic networks are often used to model biochemical oscillators and similar biological systems [115]. This kind of system is illustrated in Fig. 16, with $n = 12$ and $m = 4$.

Imagine that this system starts from the initial distribution $p(x) \propto x$, so the probability grows linearly from 0 (for $x = 0$) to maximal (for $x = n$). For the 12-state system with fourfold symmetry, this initial distribution is given by

$$p(x) = \frac{x}{\sum_{x'=0}^{11} x'} = \frac{x}{66}$$

and is shown on the left-hand side in Fig. 16. How much work can be extracted by bringing this initial distribution to some other distribution p' , while using rate matrices of the form Eq. (B23)? This distribution is bounded by the drop of the accessible free energy, via Eq. (25):

$$W(p \rightarrow p') \leq F_E[\phi_G(p)] - F_E[\phi_G(p')]. \quad (\text{B24})$$

Using the example system with 12 states and fourfold symmetry, the twirled distribution $\phi_G(p)$ is given by

$$\begin{aligned} \phi_G(p)(x) \\ = \frac{x + (x + 3 \bmod 12) + (x + 6 \bmod 12) + (x + 9 \bmod 12)}{4 \times 66}. \end{aligned}$$

For example, for the distribution $p(x) = x/66$,

$$\begin{aligned} \phi_G(p)(0) &= (0 + 3 + 6 + 9)/(4 \times 66) \approx 0.068, \\ \phi_G(p)(1) &= (1 + 4 + 7 + 10)/(4 \times 66) \approx 0.083, \\ \phi_G(p)(2) &= (2 + 5 + 8 + 11)/(4 \times 66) \approx 0.098, \\ \phi_G(p)(3) &= (3 + 6 + 9 + 0)/(4 \times 66) \approx 0.068, \\ &\dots \end{aligned}$$

This twirled distribution is shown in the right panel in Fig. 16.

7. Example: 2D Ising model, derivation of Eq. (56)

We begin by recalling the expression for accessible information in our feedback-control protocol over the 2D Ising model, which appears as Eq. (55) in the main text:

$$I_{\text{acc}}^{\phi_G}(X; M) = \ln 2 - \left\langle \ln \frac{q(m|x)}{N^{-2} \sum_{a,b} q[m|g_{a,b}(x)]} \right\rangle. \quad (\text{B25})$$

Using $q(m|x) = \delta_m(x_1)$, the expectation term in Eq. (B25) can be rewritten as

$$-\sum_x p(x) \sum_{m \in \{-1,1\}} \delta_m(x_1) \ln \left[N^{-2} \sum_{a,b} \delta_m[g_{a,b}(x)_1] \right]. \quad (\text{B26})$$

Let $z(x) = (1 + \sum_i x_i/N^2)/2$ indicate the magnetization of lattice state x , normalized to lie between 0 and 1. Note that, for any lattice state x , the frequency that spin 1 is in state 1 averaged across all translations is equal to the magnetization of x :

$$N^{-2} \sum_{a,b} \delta_1[g_{a,b}(x)_1] = z(x).$$

In addition, by symmetry, the probability that spin 1 is in state 1 averaged across all states that have magnetization z is equal to z :

$$\sum_x p(x|z) \delta_1(x_1) = z.$$

Using these results and $\delta_{-1}(x) = 1 - \delta_1(x)$, we can rewrite the expression in Eq. (B26) as

$$\begin{aligned} &-\sum_x p(x) \{ \delta_1(x_1) \ln z(x) + [1 - \delta_1(x_1)] \ln [1 - z(x)] \} \\ &= \sum_z p(z) [-z \ln z - (1 - z) \ln (1 - z)] \equiv \langle h_2(z) \rangle, \quad (\text{B27}) \end{aligned}$$

where $p(z') = \sum_x p(x) \delta_{z'}[z(x)]$ is the probability that the system has magnetization z' and h_2 is the binary entropy function.

We now consider the $N \rightarrow \infty$ limit and use Onsager's expression for the spontaneous magnetization for the 2D Ising model [116]. When β is below the critical

inverse temperature, $\beta_c = \ln(1 + \sqrt{2})/2 \approx 0.44$, the magnetization distribution $p(z)$ concentrates at $z = 1/2$, so Eq. (B27) approaches $h_2(1/2) = \ln 2$. When $\beta > \beta_c$, the magnetization distribution concentrates on a uniform mixture of two delta functions at $z = f(\beta)$ and $z = 1 - f(\beta)$, where $f(\beta) = [1 + \sqrt[3]{1 - (\sinh 2\beta)^{-4}}]/2$. In this case, Eq. (B27) approaches $\{h_2[f(\beta)] + h_2[1 - f(\beta)]\}/2 = h_2[f(\beta)]$. Combining these results with Eq. (B25) implies that $I_{\text{acc}}^{\phi_G}(X; M) = 0$ for $\beta \leq \beta_c$ and $I_{\text{acc}}^{\phi_G}(X; M) = \ln 2 - h_2[f(\beta)]$ for $\beta > \beta_c$, which appears as Eq. (56) in the main text.

APPENDIX C: MODULARITY CONSTRAINTS

1. ϕ_C obeys the Pythagorean identity [Eq. (14)]

We show that ϕ_C obeys the Pythagorean identity:

$$D[p \parallel \phi_C(q)] = D[p \parallel \phi_C(p)] + D[\phi_C(p) \parallel \phi_C(q)] \quad (\text{C1})$$

for all $p, q \in \mathcal{P}$ such that $D[p \parallel \phi_C(q)] < \infty$. For any $p, r \in \mathcal{P}$,

$$\begin{aligned} \mathbb{E}_p[\ln \phi_C(r)] &= \mathbb{E}_p[\ln r_O] + \sum_{A \in \mathcal{C}} \mathbb{E}_p[\ln r_{A \setminus O | A \cap O}] \\ &= \mathbb{E}_{\phi_C(p)}[\ln r_O] + \sum_{A \in \mathcal{C}} \mathbb{E}_{\phi_C(p)}[\ln r_{A \setminus O | A \cap O}] \quad (\text{C2}) \end{aligned}$$

$$= \mathbb{E}_{\phi_C(p)}[\ln \phi_C(r)], \quad (\text{C3})$$

where a_O and $a_{A \setminus O | A \cap O}$ indicate marginal and conditional distributions, respectively. In Eq. (C2), we use that p and $\phi_C(p)$ have the same marginals over all subsystems all $A \in \mathcal{C}$ as well as the overlap O [this result can be verified from the definition of ϕ_C , Eq. (64)]. Then,

$$\begin{aligned} D[p \parallel \phi_C(q)] &= D[p \parallel \phi_C(p)] + \mathbb{E}_p[\ln \phi_C(p) - \ln \phi_C(q)] \\ &= D[p \parallel \phi_C(p)] + \mathbb{E}_{\phi_C(p)}[\ln \phi_C(p) - \ln \phi_C(q)] \\ &= D[p \parallel \phi_C(p)] + D[\phi_C(p) \parallel \phi_C(q)], \end{aligned}$$

where the second line follows by applying Eq. (C3) twice, first taking $r = p$ and then taking $r = q$.

2. ϕ_C commutes with $e^{\tau L}$

We show that if, for some generator L , Eqs. (59) and (60) hold for all $A \in \mathcal{C}$, then ϕ_C and $e^{\tau L}$ obey the commutativity relation of Eq. (16). We assume that all $L^{(A)}$ in Eq. (60) are bounded linear operators.

Before deriving our result, we introduce some helpful notation.

- (1) $\delta_x(x')$ indicates the delta function distribution over X centered at x (the Dirac delta for continuous X and the Kronecker delta for discrete X). For any

subsystem $S \subseteq V$, $\delta_{x_S}(x'_S)$ indicates the delta function distribution over X_S centered at x_S .

- (2) $T_\tau^{(A)}(x'|x) = [e^{\tau L^{(A)}} \delta_x](x')$ indicates the conditional distribution over X , given that the system starts on state x and then evolves under $L^{(A)}$ for time τ .
- (3) For any $A \in \mathcal{C}$,

$$A := A \setminus \left(\bigcup_{B \in \mathcal{C} \setminus \{A\}} B \right) = A \setminus O(\mathcal{C})$$

indicates the set of degrees of freedom that belong exclusively to $A \in \mathcal{C}$ (and no other subsystems), and

$$A^c := V \setminus A = \bigcup_{B \in \mathcal{C} \setminus \{A\}} B$$

indicates the complement of A , which is the set of degrees of freedom that fall into at least one of the other subsystem besides A .

To derive the commutativity relation, we proceed in three steps, which are described in detail in the subsections below. In the first step, we show that, for all $\tau \geq 0$ and $A \in \mathcal{C}$, the conditional distribution $T_\tau^{(A)}(x'|x)$ can be written in the following product form:

$$T_\tau^{(A)}(x'|x) = T_\tau^{(A)}(x'_A | x_A) \delta_{x_{A^c}}(x'_{A^c}). \quad (\text{C4})$$

In the second step, we show that Eq. (C4) implies the following commutativity relation for any $p \in \mathcal{P}$ and each $A \in \mathcal{C}$:

$$e^{\tau L^{(A)}} \phi_C(p) = \phi_C(e^{\tau L^{(A)}} p). \quad (\text{C5})$$

In the third step, we show that the generators corresponding to all subsystems commute:

$$L^{(A)} L^{(B)} = L^{(B)} L^{(A)} \quad \forall A, B \in \mathcal{C}. \quad (\text{C6})$$

We then combine these three results to show that ϕ_C and $e^{\tau L}$ commute. Write

$$e^{\tau L} \phi_C(p) = e^{\sum_{A \in \mathcal{C}} \tau L^{(A)}} \phi_C(p) = \prod_{A \in \mathcal{C}} e^{\tau L^{(A)}} \phi_C(p),$$

where we use Eqs. (58) and (C6) to expand the operator exponential. Then, using Eq. (C5), write

$$\prod_{A \in \mathcal{C}} e^{\tau L^{(A)}} \phi_C(p) = \phi_C \left(\prod_{A \in \mathcal{C}} e^{\tau L^{(A)}} p \right) = \phi_C(e^{\tau L} p).$$

Combining these two results implies that $e^{\tau L} \phi_C(p) = \phi_C(e^{\tau L} p)$ for all $p \in \mathcal{P}$ and $\tau \geq 0$, as in Eq. (16).

a. Derivation of Eq. (C4)

To derive Eq. (C4), consider the conditional distribution over A given initial state x , as induced by $L^{(A)}$:

$$\begin{aligned} T_\tau^{(A)}(x'_A|x) &= [e^{\tau L^{(A)}} \delta_x]_A(x'_A) \\ &= [\delta_x]_A(x'_A) + \sum_{k \geq 1} \frac{\tau^k}{k!} [L^{(A)k} \delta_x]_A(x'_A) \\ &= \delta_{x_A}(x'_A) + \sum_{k \geq 1} \frac{\tau^k}{k!} [L^{(A)k} \delta_x]_A(x'_A), \end{aligned} \quad (\text{C7})$$

where in the second line we expand the operator exponential as $e^{\tau L^{(A)}} = \sum_k \tau^k L^{(A)k} / k!$. Note that $A \subseteq A$, so $[L^{(A)} \delta_x]_A$ is a function of $[L^{(A)} \delta_x]_A$, which, in turn, is a function of x_A by Eq. (59). Similarly, $\delta_{x_A}(x'_A)$ depends on only x_A , not x . This result means the right-hand side of Eq. (C7) depends on only x_A , which we indicate by

$$T_\tau^{(A)}(x'_A|x) = T_\tau^{(A)}(x'_A|x_A). \quad (\text{C8})$$

Now consider the conditional distribution over any other subsystem $B \neq A$ given initial state x , as induced by $L^{(A)}$:

$$\begin{aligned} T_\tau^{(A)}(x'_B|x) &= \delta_{x_B}(x'_B) + \sum_{k \geq 1} \frac{\tau^k}{k!} [L^{(A)k} \delta_x]_B(x'_B) \\ &= \delta_{x_B}(x'_B), \end{aligned} \quad (\text{C9})$$

where we use that $[L^{(A)} \delta_x]_B = 0$ by Eq. (60).

Now, it is straightforward to show that if some distribution p over X_V has delta function marginals $p_B = \delta_{x_B}$ for all $B \neq A$, then p must have following product form:

$$p(x') = p_A(x'_A) \delta_{x_{A^c}}(x'_{A^c}), \quad (\text{C10})$$

where we use that $A^c = \bigcup_{B \in \mathcal{C} \setminus \{A\}} B$. Equation (C4) follows by taking $p(x') = T_\tau^{(A)}(x'|x)$ in Eq. (C10), while using Eq. (C8).

b. Derivation of Eq. (C5)

Consider any $\tau \geq 0$ and $A \in \mathcal{C}$. Using Eq. (59) and the identity $e^{\tau L^{(A)}} = \sum_k \tau^k L^{(A)k} / k!$, one can show that, whenever two distributions $p, q \in \mathcal{P}$ obey $p_A = q_A$, it must be that $[e^{\tau L^{(A)}} p]_A = [e^{\tau L^{(A)}} q]_A$. Since $p_A = [\phi_C(p)]_A$ [see the definition of ϕ_C in Eq. (64)],

$$[e^{\tau L^{(A)}} p]_A = [e^{\tau L^{(A)}} \phi_C(p)]_A. \quad (\text{C11})$$

In addition, given Eq. (C9), we have $[e^{\tau L^{(A)}} p]_{A^c} = p_{A^c}$. Given that $B \subseteq A^c$ for each $B \neq A$, we have

$$[e^{\tau L^{(A)}} p]_B = p_B = \phi_C(p)_B = [e^{\tau L^{(A)}} \phi_C(p)]_B. \quad (\text{C12})$$

Similarly, $O(\mathcal{C}) \subseteq A^c$ and, therefore,

$$[e^{\tau L^{(A)}} p]_{O(\mathcal{C})} = [e^{\tau L^{(A)}} \phi_C(p)]_{O(\mathcal{C})}. \quad (\text{C13})$$

Now, observe that the distribution $\phi_C(p)$ does not depend on the full distribution p but only on the marginal distributions $p_{O(\mathcal{C})}$ and $\{p_A\}_{A \in \mathcal{C}}$. By Eqs. (C11)–(C13), these marginals are the same for $e^{\tau L^{(A)}} p$ and $e^{\tau L^{(A)}} \phi_C(p)$, which means that

$$\phi_C(e^{\tau L^{(A)}} p) = \phi_C[e^{\tau L^{(A)}} \phi_C(p)]. \quad (\text{C14})$$

Next, using Eq. (C4) and some simple (but rather tedious) algebra, it can be shown that

$$e^{\tau L^{(A)}} \phi_C(p) = p'_{A \setminus O|A \cap O} p_O \prod_{B \neq A} p_{B \setminus O|B \cap O}, \quad (\text{C15})$$

where

$$\begin{aligned} p'_{A \setminus O|A \cap O}(x'_A \setminus O | x'_{A \cap O}) \\ = \int T_\tau^{(A)}(x'_A | x_A, x_{A \cap O'}) p(x_A | x'_{A \cap O}) dx_A \end{aligned} \quad (\text{C16})$$

and we use the conditional distribution $T_\tau^{(A)}(x'_A | x_A, x_{A \cap O'})$ from Eq. (C8). The right-hand side of Eq. (C15) has the form of the right-hand side of Eq. (64), so it is invariant under ϕ_C :

$$\phi_C[e^{\tau L^{(A)}} \phi_C(p)] = e^{\tau L^{(A)}} \phi_C(p). \quad (\text{C17})$$

Equation (C5) follows by combining Eqs. (C14) and (C17).

c. Derivation of Eq. (C6)

Using Eq. (C4) and some algebra, one can verify that, for all $\tau \geq 0$ and $A, B \in \mathcal{C}$,

$$\int T_\tau^{(A)}(x''|x') T_\tau^{(B)}(x'|x) dx' = \int T_\tau^{(B)}(x''|x') T_\tau^{(A)}(x'|x) dx', \quad (\text{C18})$$

which in operator notation can be written as

$$e^{\tau L^{(A)}} e^{\tau L^{(B)}} \delta_x = e^{\tau L^{(B)}} e^{\tau L^{(A)}} \delta_x. \quad (\text{C19})$$

Then, for any function $f = \int f(x) \delta_x dx$, write

$$\begin{aligned}
e^{\tau L^{(A)}} e^{\tau L^{(B)}} f &= e^{\tau L^{(A)}} e^{\tau L^{(B)}} \int f(x) \delta_x dx \\
&= \int f(x) e^{\tau L^{(A)}} e^{\tau L^{(B)}} \delta_x dx \\
&= \int f(x) e^{\tau L^{(B)}} e^{\tau L^{(A)}} \delta_x dx \\
&= e^{\tau L^{(B)}} e^{\tau L^{(A)}} \int f(x) \delta_x dx \\
&= e^{\tau L^{(B)}} e^{\tau L^{(A)}} f,
\end{aligned}$$

where we exchange the order of the bounded operators $e^{\tau L^{(A)}} e^{\tau L^{(B)}}$ and $e^{\tau L^{(B)}} e^{\tau L^{(A)}}$ with the (Bochner) integral $\int f(x) \delta_x dx$ and use Eq. (C19). This expression shows that $e^{\tau L^{(A)}}$ and $e^{\tau L^{(B)}}$ commute for all $\tau \geq 0$, so their inverses $e^{-\tau L^{(A)}}$ and $e^{-\tau L^{(B)}}$ must also commute. Given that $e^{\tau L^{(A)}}$ and $e^{\tau L^{(B)}}$ commute for all $\tau \in \mathbb{R}$, $L^{(A)}$ and $L^{(B)}$ must commute (see p. 23 in Ref. [117]).

3. Szilard box: Derivation of Eqs. (74) and (76)

We first derive Eq. (74). Using Eq. (70) and some rearrangement, write

$$D[\phi_C(p_\theta)||u] = \ln 4 - S[p_\theta(X_1)] - S[p_\theta(X_2)], \quad (\text{C20})$$

where $S[p_\theta(X_1)]$ and $S[p_\theta(X_2)]$ refer to the marginal entropies under p_θ . It is easy to see that, by symmetry,

$$S[p_\theta(X_1)] = S[p_{(\pi/2)-\theta}(X_2)]. \quad (\text{C21})$$

Therefore, we derive a closed-form expression for $D[\phi_C(p_\theta)||u]$ by finding a closed-form expression for

$$S[p_\theta(X_1)] := - \int_{-1}^1 p_\theta(x_1) \ln p_\theta(x_1) dx_1. \quad (\text{C22})$$

First, consider the case of $\theta \in [-\pi/2, \pi/2]$, and define $A_\theta := |\tan \theta|$. It can be verified from Eq. (49) that the marginal distribution $p_\theta(x_1)$ always has a piecewise linear form. In particular, if $A_\theta < 1$, then, for any $x_1 \in [-1, 1]$,

$$p_\theta(x_1) = \begin{cases} 1 & \text{if } -1 \leq x_1 \leq -A_\theta, \\ \frac{A_\theta - x_1}{2A_\theta} & \text{if } -A_\theta \leq x_1 \leq A_\theta, \\ 0 & \text{if } x_1 > A_\theta. \end{cases} \quad (\text{C23})$$

Otherwise, if $A_\theta > 1$, then, for any $x_1 \in [-1, 1]$,

$$p_\theta(x_1) = \frac{A_\theta - x_1}{2A_\theta}. \quad (\text{C24})$$

Plugged into Eq. (C22), this expression gives

$$S[p_\theta(X_1)] = \begin{cases} - \int_{-1}^1 \frac{A_\theta - x_1}{2A_\theta} \ln \frac{A_\theta - x_1}{2A_\theta} dx_1 & \text{if } A_\theta > 1, \\ - \int_{-A_\theta}^{A_\theta} \frac{A_\theta - x_1}{2A_\theta} \ln \frac{A_\theta - x_1}{2A_\theta} dx_1 & \text{otherwise.} \end{cases}$$

Integrating these two cases separately in *Mathematica*, and plugging in the definition of A_θ , gives

$$S[p_\theta(X_1)] = \frac{1}{2} \begin{cases} f(|\tan \theta|) & \text{if } |\tan \theta| > 1, \\ |\tan \theta| & \text{otherwise,} \end{cases} \quad (\text{C25})$$

where for convenience we define

$$f(x) := 1 - \frac{1+x^2}{2x} \ln \frac{x+1}{x-1} - \ln \frac{x^2-1}{4x^2}. \quad (\text{C26})$$

Recall that so far we assume that $\theta \in [-\pi/2, \pi/2]$. However, by Eq. (49), $p_\theta(x_1, x_2) = p_{\pm\pi-\theta}(-x_1, x_2)$, which implies that $p_\theta(x_1) = p_{\pi-\theta}(-x_1) = p_{-\pi-\theta}(-x_1)$ and $S[p_\theta(X_1)] = S[p_{\pi-\theta}(X_1)] = S[p_{-\pi-\theta}(X_1)]$. It can also be verified that $|\tan \theta| = |\tan(\pi - \theta)| = |\tan(-\pi - \theta)|$, so, in fact, Eq. (C25) holds for all $\theta \in [-\pi, \pi]$.

Finally, if $|\theta| \in [(\pi/4), (3\pi/4)]$, then Eqs. (C25) and (C21) imply

$$\begin{aligned}
|\tan \theta| > 1, \quad S[p_\theta(X_1)] &= \frac{1}{2} f(|\tan \theta|), \\
\left| \tan\left(\frac{\pi}{2} - \theta\right) \right| \leq 1, \quad S[p_\theta(X_2)] &= \frac{1}{2} \left| \tan\left(\frac{\pi}{2} - \theta\right) \right|.
\end{aligned}$$

Conversely, if $|\theta| \in [0, \pi] \setminus [(\pi/4), (3\pi/4)]$, then

$$\begin{aligned}
|\tan \theta| \leq 1, \quad S[p_\theta(X_1)] &= \frac{1}{2} |\tan \theta|, \\
\left| \tan\left(\frac{\pi}{2} - \theta\right) \right| > 1, \quad S[p_\theta(X_2)] &= \frac{1}{2} f\left[\left| \tan\left(\frac{\pi}{2} - \theta\right) \right| \right].
\end{aligned}$$

Equation (74) follows by combining these results and rearranging.

To derive Eq. (76), use $\phi_G[\phi_C(p_\theta)](x_1, x_2) = p_\theta(x_1) u(x_2)$ to write

$$\begin{aligned}
D\{\phi_G[\phi_C(p_\theta)]||u\} &= \ln 4 - S[p_\theta(X_1)] - S[u(X_2)] \\
&= \ln 2 - S[p_\theta(X_1)], \quad (\text{C27})
\end{aligned}$$

where we use that $S[u(X_2)] = \ln 2$. Equation (76) then follows by combining Eqs. (C27) and (C25).

4. Example: Feedback-controlled flashing ratchet

Here, we derive a closed-form expression for the accessible information in the feedback-controlled collective flashing ratchet.

For notational convenience, let $a = 1/\alpha$ indicate the slope of the increasing part of V in Fig. 10(b) and $b = -1/(1 - \alpha)$ indicate the slope of the decreasing part of V . Note that the net force $\sum_v V'(x_v)$ can be seen as the sum of N random variables, where by assumption each $V'(x_v)$ is equal to $a = 1/\alpha$ with probability α and equal to $b = -1/(1 - \alpha)$ with probability $1 - \alpha$. This result implies that the expectation of $V'(x_v)$ is 0 and the variance is $1/[\alpha(1 - \alpha)]$.

We first compute the accessible information $I_{\text{acc}}^{\phi_c}(X; M) = \sum_v I(X_v; M) = N \cdot I(X_1; M)$. The mutual information between M and the state of a single particle X_1 is given by

$$\begin{aligned} I(X_1; M) &= S(M) - S(M|X) \\ &= h_2[p(1)] - \alpha h_2[p(1|a)] \\ &\quad - (1 - \alpha) h_2[p(1|b)], \end{aligned} \quad (\text{C28})$$

where $p(1)$ is the probability that the net force is positive, $p(1|a)$ is the probability that the net force is positive given that particle X_1 experiences force a , and $p(1|b)$ is the probability that the net force is positive given that particle

X_1 experiences force b . We can compute $p(1)$ by considering the case when $k = 0, 1, 2, \dots$ particles experience force a . Assuming the particles are independent, this case is given by

$$p(1) = \sum_{k=0}^N B_{N,\alpha}(k) \Theta[ka + (N - k)b], \quad (\text{C29})$$

where $B_{N,\alpha}$ is the binomial probability of k successes, given N trials with success probability α . To compute $p(1|a)$, note that, given that X_1 experiences force a , $M = 1$ whenever the other $N - 1$ particles experience a net force larger than $-a$. The probability of this event is

$$p(1|a) = \sum_{k=0}^{N-1} B_{N-1,\alpha}(k) \Theta[ka + (N - 1 - k)b + a]. \quad (\text{C30})$$

Conversely, if X_1 experiences force b , then $M = 1$ if the other $N - 1$ particles experience a net force larger than $-b$, which has probability

$$p(1|b) = \sum_{k=0}^{N-1} B_{N-1,\alpha}(k) \Theta[ka + (N - 1 - k)b + b]. \quad (\text{C31})$$

Plugging Eqs. (C29)–(C31) into Eq. (C28) gives $I(X_1; M)$. Multiplying by N gives the accessible information:

$$\begin{aligned} I_{\text{acc}}^{\phi_c}(X; M) &= N \cdot I(X_1; M) \\ &= N \left[h_2 \left(\sum_{k=0}^N B_{N,\alpha}(k) \Theta[ka + (N - k)b] \right) - \alpha h_2 \left(\sum_{k=0}^{N-1} B_{N-1,\alpha}(k) \Theta[ka + (N - 1 - k)b + a] \right) \right. \\ &\quad \left. - (1 - \alpha) h_2 \left(\sum_{k=0}^{N-1} B_{N-1,\alpha}(k) \Theta[ka + (N - 1 - k)b + b] \right) \right]. \end{aligned} \quad (\text{C32})$$

This result is shown in Fig. 12(left) for different values of N and α .

To compute the efficiency values in Fig. 12(right), we simply divide $I_{\text{acc}}^{\phi_c}(X; M)$ by $I(X; M)$, the total mutual information between the measurement and all particles. Since the measurement in Eq. (80) is deterministic, this mutual information is given by the entropy of M :

$$I(X; M) = S(M) = h_2[p(1)], \quad (\text{C33})$$

which can be computed using Eq. (C29).

We now compute the asymptotic value of accessible information and efficiency in the $N \rightarrow \infty$ limit. The sum of a large number of independent random variables with mean 0 and variance $1/[\alpha(1 - \alpha)]$ approaches a Gaussian with mean 0 and variance $N/[\alpha(1 - \alpha)]$. Thus, in the $N \rightarrow \infty$ limit, the probability that the force is positive converges to

$p(1) = 1/2$, so $I(X; M) = S(M)$ converges to $\ln 2$. Recall that $p(1|a)$ is given by the probability that $N - 1$ particles experience a net force larger than $-a$. In the $N \rightarrow \infty$ limit, this conditional probability converges to

$$p(1|a) = 1 - \Phi_{\alpha, N-1}(-a) = \Phi_{\alpha, N-1}(a),$$

where $\Phi_{\alpha, N-1}$ is the cumulative distribution function of a Gaussian with mean 0 and variance $N/[\alpha(1 - \alpha)]$. We can similarly calculate

$$p(1|b) = 1 - \Phi_{\alpha, N-1}(-b) = \Phi_{\alpha, N-1}(b).$$

Plugging into Eq. (C28) gives

$$I(X_1; M) = \ln 2 - \alpha h_2[\Phi_{\alpha, N-1}(a)] - (1 - \alpha) h_2[\Phi_{\alpha, N-1}(b)]. \quad (\text{C34})$$

Using $a = 1/\alpha$ and $b = -1/(1 - \alpha)$ and some analysis (e.g., by taking limits in *Mathematica*) shows that

$$\lim_{N \rightarrow \infty} N \cdot I(X_1; M) = \frac{1}{\pi}, \quad (\text{C35})$$

irrespective of α . This result is the asymptotic accessible information, which appears as the dotted line in Fig. 12 (left). The asymptotic efficiency, which appears as the dotted line in Fig. 12 (right), is given by $1/(\pi \ln 2)$ [since $I(X; M) = \ln 2$ in the $N \rightarrow \infty$ limit].

APPENDIX D: COARSE-GRAINED CONSTRAINTS

1. Derivation of Eq. (82) from Eqs. (83) and (85)

In general, the microstate distribution p evolves according to some generator L , $\partial_t p(t) = Lp(t)$, and the macrostate distribution p_Z evolves according to a coarse-grained generator \hat{L}^p . In general, the coarse-grained dynamics are not closed, meaning that \hat{L}^p can depend on the microstate distribution p . In this section, we provide concrete conditions on the generators that guarantee that the coarse-grained dynamics are closed. In the following derivations, for notational simplicity, we omit the dependence of $p(x, t)$ and $p(z, t)$ on t .

For discrete-state master equations, the coarse-grained dynamics are given by [58]

$$\begin{aligned} \partial_t p_Z(z) &= \hat{L}^p p_Z(z) \\ &= \sum_z [\hat{L}_{zz'}^p] p_Z(z') - \hat{L}_{zz'}^p p_Z(z), \end{aligned} \quad (\text{D1})$$

where $\hat{L}_{zz'}^p$ is the transition rate from macrostate z' to z :

$$\hat{L}_{zz'}^p = \sum_{x'} p(x'|z') \sum_x \delta_{\xi(x)}(z) L_{xx'}. \quad (\text{D2})$$

By plugging Eq. (83) into Eq. (D2) and simplifying, one can verify that $\hat{L}_{zz'}^p$ does not depend on the microstate distribution p ; therefore, Eq. (82) holds.

A similar approach can be used for continuous-state master equations.

We now consider Fokker-Planck equations of the form Eq. (84), given a linear coarse-graining function $\xi(x) = Wx$. Using Proposition 2.8 in Ref. [57], we write the evolution of the coarse-grained distribution p_Z as

$$\partial_t p_Z(z) = \nabla \cdot [\hat{A}(z) p_Z(z)] + \beta^{-1} \text{tr} \{ H^T [\hat{D}(z) p_Z(z)] \}, \quad (\text{D3})$$

where H is the Hessian matrix of second derivative operators and we define

$$\hat{A}(z) := \int [p(x|z) W \nabla E(x) - \beta^{-1} \Delta \xi(x)] dx \quad (\text{D4})$$

$$= \int [p(x|z) W \nabla E(x)] dx \quad (\text{D5})$$

$$= -\hat{F}(z), \quad (\text{D6})$$

$$\hat{D}(z) := \int p(x|z) W W^T dx = I. \quad (\text{D7})$$

We use Eq. (2.29) from Ref. [57] in Eq. (D4), the linearity of ξ in Eq. (D5), and Eq. (85) in Eq. (D6). We use Eq. (2.30) from Ref. [57] and the assumption that $W W^T = I$ in Eq. (D7). It is easy to check that $\text{tr}[H^T(I p_Z)] = \Delta p_Z$, which can be combined with Eqs. (D6), (D7), and (D3) to give Eq. (86). Since the right-hand side of Eq. (86) does not depend on the microstate distribution, the coarse-grained dynamics are closed.

2. Derivation of Eq. (87)

Our derivation below does not assume isothermal protocols, so the inequality in Eq. (87) holds both for isothermal protocols and for protocols connected to any number of thermodynamic reservoirs.

To derive this result for a given L , we make two assumptions. First, as described in the main text, we assume that the coarse-grained dynamics are closed [Eq. (82)]. Second, we assume that the coarse-grained stationary distribution π_Z (where π is the stationary distribution of L) is invariant under conjugation of odd-parity variables:

$$\pi_Z[\xi(x)] = \pi_Z[\xi(x^\dagger)] \quad \forall x \in X, \quad (\text{D8})$$

where x^\dagger indicate the conjugation of state x in which all odd-parity variables (such as momentum) have their sign flipped. For an isothermal protocol, the stationary distributions are equilibrium distributions, and Eq. (D8) is satisfied [96]. For more general protocols, Eq. (D8) holds if there are no odd-parity variables (e.g., overdamped dynamics), so $x = x^\dagger$. It also holds if the coarse-graining function maps each x and its conjugate to the same macrostate, $\xi(x) = \xi(x^\dagger)$, as well as some other cases.

Now imagine a system that starts from some initial distribution p at time $t = 0$, and then undergoes free relaxation under L toward a (possibly nonequilibrium) stationary distribution π , reaching a final distribution p' by time $t = \tau$. Next, we use existing results in stochastic thermodynamics [96, 110] and write the EP incurred over time interval $t \in [0, \tau]$ as

$$\Sigma(\tau) = D[p(\mathbf{x}, \boldsymbol{\nu}) \| \tilde{p}(\tilde{\mathbf{x}}^\dagger, \tilde{\boldsymbol{\nu}})] \quad (\text{D9})$$

(see also Appendix A 2), where

- (1) $\mathbf{x} = (x, \dots, x')$ indicate a continuous-time trajectory of system states over time interval $t \in [0, \tau]$, where x and x' indicate the initial and final system states, respectively, and $\tilde{\mathbf{x}}^\dagger = (x'^\dagger, \dots, x^\dagger)$ is the corresponding time-reversed and conjugated trajectory;
- (2) $\boldsymbol{\nu}$ is a sequence of reservoirs which exchange conserved quantities with the system during $t \in [0, \tau]$, and $\tilde{\boldsymbol{\nu}}$ is the corresponding time-reversed sequence [25,28,110];
- (3) $p(\mathbf{x}, \boldsymbol{\nu}) = P(\mathbf{x}, \boldsymbol{\nu}|x)p(x)$ is the probability of forward trajectory $(\mathbf{x}, \boldsymbol{\nu})$ given initial distribution p , where $P(\mathbf{x}, \boldsymbol{\nu}|x)$ is the conditional distribution generated by the free relaxation;
- (4) $\tilde{p}(\tilde{\mathbf{x}}^\dagger, \tilde{\boldsymbol{\nu}}) = P(\tilde{\mathbf{x}}^\dagger, \tilde{\boldsymbol{\nu}}|x'^\dagger)p'(x')$ is the probability of reverse trajectory $(\tilde{\mathbf{x}}^\dagger, \tilde{\boldsymbol{\nu}})$ under a free relaxation that starts with the following distribution:

$$p'(x') = \int P(x'|x)p(x)dx. \quad (\text{D10})$$

Using the fact that EP decreases under state-space and temporal coarse-graining [58,118], we bound Eq. (D9) as

$$\Sigma(\tau) \geq D[p(\mathbf{x}) \| p(\tilde{\mathbf{x}}^\dagger)] \geq D[p(z, z') \| \tilde{p}(z^\dagger, z'^\dagger)], \quad (\text{D11})$$

where $z = \xi(x)$, $z' = \xi(x')$, $z^\dagger = \xi(x^\dagger)$, and $z'^\dagger = \xi(x'^\dagger)$. The final KL divergence can be decomposed as

$$\begin{aligned} & D[p(z, z') \| \tilde{p}(z^\dagger, z'^\dagger)] \\ &= [D(p_Z \| \pi_Z) - D(p'_Z \| \pi_Z)] \\ &+ \int p(z, z') \ln \left[\frac{p(z, z') \pi_Z(z) p'_Z(z')}{\tilde{p}(z^\dagger, z'^\dagger) p_Z(z) \pi_Z(z')} \right] dz dz'. \end{aligned} \quad (\text{D12})$$

Using Jensen's inequality, we lower bound the integral term as

$$\begin{aligned} & \int p(z, z') \ln \left[\frac{p(z, z') \pi_Z(z) p'_Z(z')}{\tilde{p}(z^\dagger, z'^\dagger) p_Z(z) \pi_Z(z')} \right] dz dz' \\ &= - \int p(z, z') \ln \left[\frac{\tilde{p}(z^\dagger, z'^\dagger) p_Z(z) \pi_Z(z')}{p(z, z') \pi_Z(z) p'_Z(z')} \right] dz dz' \\ &\geq - \ln \left[\int \frac{\tilde{p}(z^\dagger, z'^\dagger) p_Z(z) \pi_Z(z')}{\pi_Z(z) p'_Z(z')} dz dz' \right]. \end{aligned} \quad (\text{D13})$$

Note that $\pi_Z(z') = \pi_Z(z'^\dagger)$ by Eq. (D8) and $\tilde{p}(z^\dagger, z'^\dagger) = p'_Z(z')$ by the definition of p'_Z in Eq. (D10), allowing us to rewrite the rhs of Eq. (D13) as

$$- \ln \left\{ \int \frac{p_Z(z)}{\pi_Z(z)} \left[\int \tilde{p}(z^\dagger | z'^\dagger) \pi_Z(z'^\dagger) dz' \right] dz \right\}. \quad (\text{D14})$$

The inner integral can be further rewritten as

$$\begin{aligned} & \int \tilde{p}(z^\dagger | z'^\dagger) \pi_Z(z'^\dagger) dz' \\ &= \int P(z^\dagger | x'^\dagger) \tilde{p}(x'^\dagger | z'^\dagger) \pi_Z(z'^\dagger) dx' \\ &= \pi_Z(z^\dagger) \\ &= \pi_Z(z), \end{aligned}$$

where in the third line we use the assumption of closed dynamics [Eq. (82)] and the stationarity of π under $P(\cdot|\cdot)$ and in the fourth line we use Eq. (D8). We can then rewrite Eq. (D14) as

$$- \ln \left[\int \frac{\pi_Z(z)}{\pi_Z(z)} \pi_Z(z) dz \right] = 0.$$

Combined with Eq. (D13), this expression implies that the integral term in Eq. (D12) is non-negative. Combining with Eq. (D11) gives

$$\Sigma(\tau) \geq D(p_Z \| \pi_Z) - D(p'_Z \| \pi_Z).$$

Finally, using the definition of the EP rate and the results above,

$$\begin{aligned} \dot{\Sigma}(p, L) &:= \lim_{\tau \rightarrow 0} \frac{1}{\tau} \Sigma(\tau) \\ &\geq \lim_{\tau \rightarrow 0} \frac{1}{\tau} [D(p_Z \| \pi_Z) - D(p'_Z \| \pi_Z)] \\ &= - \int \partial_t p_Z(t)(z) \ln \frac{p_Z(z)}{\pi_Z(z)} dz \geq 0, \end{aligned} \quad (\text{D15})$$

where $\partial_t p_Z(t) = \hat{L} p_Z$. Equation (D15) follows from Eqs. (A3)–(A6) above (with summations replaced by integrals). The discrete-state form of Eq. (D15), and also where p and L are explicitly time dependent, appears in the main text as Eq. (87).

-
- [1] K. Takara, H.-H. Hasegawa, and D. Driebe, *Generalization of the Second Law for a Transition between Nonequilibrium States*, *Phys. Lett. A* **375**, 88 (2010).
 - [2] J. M. R. Parrondo, J. M. Horowitz, and T. Sagawa, *Thermodynamics of Information*, *Nat. Phys.* **11**, 131 (2015).
 - [3] M. Esposito and C. Van den Broeck, *Second Law and Landauer Principle far from Equilibrium*, *Europhys. Lett.* **95**, 40004 (2011).
 - [4] We use a Brownian model of the Szilard engine, which is similar to setups commonly employed in modern non-equilibrium statistical physics [2,5–9], as shown in Fig. 1. This model can be justified by imagining a box that

- contains a large colloidal particle, as well as a medium of small solvent particles to which the vertical partition is permeable. Note that this model differs from Szilard's original proposal [10], in which the box contains a single particle in a vacuum, which is analyzed in Refs. [11–13].
- [5] A. Bérut, A. Arakelyan, A. Petrosyan, S. Ciliberto, R. Dillenschneider, and E. Lutz, *Experimental Verification of Landauer's Principle Linking Information and Thermodynamics*, *Nature (London)* **483**, 187 (2012).
 - [6] É. Roldán, I. A. Martínez, J. M. Parrondo, and D. Petrov, *Universal Features in the Energetics of Symmetry Breaking*, *Nat. Phys.* **10**, 457 (2014).
 - [7] J. V. Koski, V. F. Maisi, T. Sagawa, and J. P. Pekola, *Experimental Observation of the Role of Mutual Information in the Nonequilibrium Dynamics of a Maxwell Demon*, *Phys. Rev. Lett.* **113**, 030601 (2014).
 - [8] K. Shizume, *Heat Generation Required by Information Erasure*, *Phys. Rev. E* **52**, 3495 (1995).
 - [9] Z. Gong, Y. Lan, and H. T. Quan, *Stochastic Thermodynamics of a Particle in a Box*, *Phys. Rev. Lett.* **117**, 180603 (2016).
 - [10] L. Szilard, *Über die Entropieverminderung in Einem Thermodynamischen System Bei Eingriffen Intelligenter Wesen*, *Z. Phys.* **53**, 840 (1929).
 - [11] K. Proesmans, C. Driessen, B. Cleuren, and C. Van den Broeck, *Efficiency of Single-Particle Engines*, *Phys. Rev. E* **92**, 032105 (2015).
 - [12] T. Hondou, *Equation of State in a Small System: Violation of an Assumption of Maxwell's Demon*, *Europhys. Lett.* **80**, 50001 (2007).
 - [13] D. Bhat, S. Sabhapandit, A. Kundu, and A. Dhar, *Unusual Equilibration of a Particle in a Potential with a Thermal Wall*, *J. Stat. Mech.* (2017) 113210.
 - [14] T. Sagawa and M. Ueda, *Second Law of Thermodynamics with Discrete Quantum Feedback Control*, *Phys. Rev. Lett.* **100**, 080403 (2008).
 - [15] As is common in the literature, in Eq. (3) we consider only the work that is extractable from the system after the measurement is made. We do not account for the possible work cost of making the measurement nor any work exchanges that may be incurred by the measurement apparatus during the driving.
 - [16] P. A. Corning and S. J. Kline, *Thermodynamics, Information and Life Revisited, Part II: "Thermoeconomics" and "Control Information"*, *Syst. Res. Behav. Sci.* **15**, 453 (1998).
 - [17] A. Kolchinsky and D. H. Wolpert, *Semantic Information, Autonomous Agency and Non-equilibrium Statistical Physics*, *Interface Focus* **8**, 20180041 (2018).
 - [18] S. A. Kauffman, *Investigations* (Oxford University, New York, 2000).
 - [19] D. H. Wolpert and A. Kolchinsky, *Thermodynamics of Computing with Circuits*, *New J. Phys.* **22**, 063047 (2020).
 - [20] J. Song, S. Still, R. D. H. Rojas, I. P. Castillo, and M. Marsili, *Optimal Work Extraction and Mutual Information in a Generalized Szilárd Engine*, *Phys. Rev. E* **103**, 052121 (2021).
 - [21] F. J. Cao, L. Dinis, and J. M. R. Parrondo, *Feedback Control in a Collective Flashing Ratchet*, *Phys. Rev. Lett.* **93**, 040603 (2004).
 - [22] D. Janzing, *Quantum Thermodynamics with Missing Reference Frames: Decompositions of Free Energy into Non-increasing Components*, *J. Stat. Phys.* **125**, 761 (2006).
 - [23] J. A. Vaccaro, F. Anselmi, H. M. Wiseman, and K. Jacobs, *Tradeoff between Extractable Mechanical Work, Accessible Entanglement, and Ability to Act as a Reference System, under Arbitrary Superselection Rules*, *Phys. Rev. A* **77**, 032114 (2008).
 - [24] The assumption of unique stationary distributions can be relaxed as long as the operator ϕ (as discussed in Sec. III) satisfies the following weak technical condition: For all $p \in \mathcal{P}$ and each stationary distribution π of each $L \in \Lambda$, $D[p||\phi(\pi)] < \infty$ whenever $D(p||\pi) < \infty$. Note that $\phi(\pi)$ is also a stationary distribution of L by Lemma 1 in Appendix A, so this condition is automatically satisfied when the generators have unique stationary distributions [since in that case $\pi = \phi(\pi)$]. Note also that, if some $L \in \Lambda$ have multiple stationary distributions π , the corresponding EP rate in Eq. (11) can be equivalently defined using any π such that $D(p||\pi) < \infty$.
 - [25] M. Esposito and C. Van den Broeck, *Three Faces of the Second Law. I. Master Equation Formulation*, *Phys. Rev. E* **82**, 011143 (2010).
 - [26] N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier, New York, 1992), Vol. 1.
 - [27] H. Risken, *Fokker-Planck Equation*, in *The Fokker-Planck Equation* (Springer, New York, 1996), pp. 63–95.
 - [28] C. Van den Broeck and M. Esposito, *Three Faces of the Second Law. II. Fokker-Planck Formulation*, *Phys. Rev. E* **82**, 011144 (2010).
 - [29] D. L. Ermak and J. A. McCammon, *Brownian Dynamics with Hydrodynamic Interactions*, *J. Chem. Phys.* **69**, 1352 (1978).
 - [30] S.-I. Amari, *Information Geometry and Its Applications* (Springer, New York, 2016), Vol. 194.
 - [31] Note that $D(p||q) \geq D[p||\phi(p)]$ for any $q \in \text{img } \phi$, which follows from Eq. (14) and the non-negativity of KL divergence.
 - [32] A. Kolchinsky and D. H. Wolpert, *Entropy Production Given Constraints on the Energy Functions*, *Phys. Rev. E* **104**, 034129 (2021).
 - [33] U. Seifert, *Stochastic Thermodynamics, Fluctuation Theorems and Molecular Machines*, *Rep. Prog. Phys.* **75**, 126001 (2012).
 - [34] A. Kolchinsky and D. H. Wolpert, *The State Dependence of Integrated, Instantaneous, and Fluctuating Entropy Production in Quantum and Classical Processes*, arXiv: 2103.05734.
 - [35] H. Kwon and M. S. Kim, *Fluctuation Theorems for a Quantum Channel*, *Phys. Rev. X* **9**, 031029 (2019).
 - [36] A compact group \mathcal{G} has a measurable action over X if the action $\mathcal{G} \times X \rightarrow X$ is a measurable function, where we assume \mathcal{G} and X are endowed with their respective Borel algebras.
 - [37] Technically, the definition of the twirling operator in Eq. (42) applies only when p is a finite-valued probability density function (which excludes things such as the Dirac delta "function"). A more general formulation of our results can be developed in terms of probability measures

- rather than probability densities [see Chap. 3 in Ref. [38] for a version of Eq. (42) defined in terms of probability measures].
- [38] M. L. Eaton, *Group Invariance Applications in Statistics*, in *Regional Conference Series in Probability and Statistics* (JSTOR, New York, 1989), pp. i–133.
- [39] K. G. H. Vollbrecht and R. F. Werner, *Entanglement Measures under Symmetry*, *Phys. Rev. A* **64**, 062307 (2001).
- [40] Technically, the wall potential as defined in Eq. (45) is nondifferentiable. To be more accurate, one should imagine it in terms of the limit $V_w(|x|) = \lim_{\alpha \rightarrow \infty} |x|^\alpha$ [41].
- [41] A. Dhar, A. Kundu, S. N. Majumdar, S. Sabhapandit, and G. Schehr, *Run-and-Tumble Particle in One-Dimensional Confining Potentials: Steady-State, Relaxation, and First-Passage Properties*, *Phys. Rev. E* **99**, 032132 (2019).
- [42] S. Still and D. Daimer, *Partially Observable Szilard Engines*, arXiv:2103.15803.
- [43] P. L. Krapivsky, S. Redner, and E. Ben-Naim, *A Kinetic View of Statistical Physics* (Cambridge University Press, Cambridge, England, 2010).
- [44] J. Lekscha, H. Wilming, J. Eisert, and R. Gallego, *Quantum Thermodynamics with Local Control*, *Phys. Rev. E* **97**, 022142 (2018).
- [45] N. Gershenfeld, *Signal Entropy and the Thermodynamics of Computation*, *IBM Systems Journal* **35**, 577 (1996).
- [46] A. B. Boyd, D. Mandal, and J. P. Crutchfield, *Thermodynamics of Modularity: Structural Costs beyond the Landauer Bound*, *Phys. Rev. X* **8**, 031036 (2018).
- [47] G. Schlosser and G. P. Wagner, *Modularity in Development and Evolution* (University of Chicago, Chicago, 2004).
- [48] O. Sporns and R. F. Betzel, *Modular Brain Networks*, *Annu. Rev. Psychol.* **67**, 613 (2016).
- [49] One can also apply the results in this section to Fokker-Planck equations that can be put in the form of Eq. (62) via an appropriate change of variables; see Sec. IV.9 in Ref. [27].
- [50] The multi-information is a well-known generalization of mutual information, which is also sometimes called “total correlation” [51].
- [51] S. Watanabe, *Information Theoretical Analysis of Multivariate Correlation*, *IBM J. Res. Dev.* **4**, 66 (1960).
- [52] E. Craig, N. Kuwada, B. Lopez, and H. Linke, *Feedback Control in Flashing Ratchets*, *Ann. Phys. (Berlin)* **17**, 115 (2008).
- [53] M. Feito and F. J. Cao, *Information and Maximum Power in a Feedback Controlled Brownian Ratchet*, *Eur. Phys. J. B* **59**, 63 (2007).
- [54] B. J. Lopez, N. J. Kuwada, E. M. Craig, B. R. Long, and H. Linke, *Realization of a Feedback Controlled Flashing Ratchet*, *Phys. Rev. Lett.* **101**, 220601 (2008).
- [55] G. Nicolis, *Transformation Properties of Entropy Production*, *Phys. Rev. E* **83**, 011112 (2011).
- [56] If $\xi(x) = Wx$ and $WW^T \neq I$, one can define an equivalent, rescaled coarse-graining function $\xi'(x) = W'x$, where $W' := (WW^T)^{-1/2}W$, which obeys $W'W'^T = I$.
- [57] M. H. Duong, A. Lamacz, M. A. Peletier, A. Schlichting, and U. Sharma, *Quantification of Coarse-Graining Error in Langevin and Overdamped Langevin Dynamics, Non-linearity* **31**, 4517 (2018).
- [58] M. Esposito, *Stochastic Thermodynamics under Coarse Graining*, *Phys. Rev. E* **85**, 041125 (2012).
- [59] H. Wilming, R. Gallego, and J. Eisert, *Second Law of Thermodynamics under Control Restrictions*, *Phys. Rev. E* **93**, 042126 (2016).
- [60] M. Perarnau-Llobet, A. Riera, R. Gallego, H. Wilming, and J. Eisert, *Work and Entropy Production in Generalised Gibbs Ensembles*, *New J. Phys.* **18**, 123035 (2016).
- [61] M. Esposito, R. Kawai, K. Lindenberg, and C. Van den Broeck, *Finite-Time Thermodynamics for a Single-Level Quantum Dot*, *Europhys. Lett.* **89**, 20003 (2010).
- [62] D. A. Sivak and G. E. Crooks, *Thermodynamic Metrics and Optimal Paths*, *Phys. Rev. Lett.* **108**, 190602 (2012).
- [63] N. Shiraishi, K. Funo, and K. Saito, *Speed Limit for Classical Stochastic Processes*, *Phys. Rev. Lett.* **121**, 070601 (2018).
- [64] A. Gomez-Marín, T. Schmiedl, and U. Seifert, *Optimal Protocols for Minimal Work Processes in Underdamped Stochastic Thermodynamics*, *J. Chem. Phys.* **129**, 024114 (2008).
- [65] H. Then and A. Engel, *Computing the Optimal Protocol for Finite-Time Processes in Stochastic Thermodynamics*, *Phys. Rev. E* **77**, 041105 (2008).
- [66] P. R. Zulkowski and M. R. DeWeese, *Optimal Finite-Time Erasure of a Classical Bit*, *Phys. Rev. E* **89**, 052140 (2014).
- [67] T. Schmiedl and U. Seifert, *Optimal Finite-Time Processes in Stochastic Thermodynamics*, *Phys. Rev. Lett.* **98**, 108301 (2007).
- [68] B. B. Machta, *Dissipation Bound for Thermodynamic Control*, *Phys. Rev. Lett.* **115**, 260603 (2015).
- [69] G. Verley, C. V. d. Broeck, and M. Esposito, *Work Statistics in Stochastically Driven Systems*, *New J. Phys.* **16**, 095001 (2014).
- [70] A. E. Allahverdyan, R. Balian, and T. M. Nieuwenhuizen, *Maximal Work Extraction from Finite Quantum Systems*, *Europhys. Lett.* **67**, 565 (2004).
- [71] R. Uzdin and S. Rahav, *The Passivity Deformation Approach for the Thermodynamics of Isolated Quantum Setups*, *PRX Quantum* **2**, 010336 (2021).
- [72] A. P. Solon and J. M. Horowitz, *Phase Transition in Protocols Minimizing Work Fluctuations*, *Phys. Rev. Lett.* **120**, 180605 (2018).
- [73] T. R. Gingrich, G. M. Rotskoff, G. E. Crooks, and P. L. Geissler, *Near-Optimal Protocols in Complex Nonequilibrium Transformations*, *Proc. Natl. Acad. Sci. U.S.A.* **113**, 10263–263 (2016).
- [74] E. Aurell, C. Mejía-Monasterio, and P. Muratore-Ginanneschi, *Optimal Protocols and Optimal Transport in Stochastic Thermodynamics*, *Phys. Rev. Lett.* **106**, 250601 (2011).
- [75] S. Still, *Thermodynamic Cost and Benefit of Memory*, *Phys. Rev. Lett.* **124**, 050601 (2020).
- [76] That paper proposes to divide the system into two subsystems Y and \bar{Y} , such that the accessible information is given by $I(M; Y)$, under three assumption: (i) the system’s marginal distributions remain constant during all steps of feedback control, (ii) the conditional distribution of \bar{Y}

given the system and the measurement does not change during the driving, and (iii) all conditional information about \bar{Y} is lost by the time that driving begins. After private communication with the author of Ref. [75], we think that condition (iii) may need to be formalized as $p[\bar{y}(t_2)|y(t_2), z(t_0)] = p[\bar{y}(t_2)|y(t_2)]$, although this equation does not appear in that paper.

- [77] D. H. Wolpert, *The Stochastic Thermodynamics of Computation*, *J. Phys. A* **52**, 193001 (2019).
- [78] D. Wolpert, *Uncertainty Relations and Fluctuation Theorems for Bayes Nets*, *Phys. Rev. Lett.* **125**, 200602 (2020).
- [79] D. H. Wolpert, *Fluctuation Theorems for Multipartite Processes*, arXiv:2003.11144.
- [80] D. H. Wolpert, *Minimal Entropy Production Rate of Interacting Systems*, *New J. Phys.* **22**, 113013 (2020).
- [81] D. H. Wolpert, *Strengthened Landauer Bound for Composite Systems*, arXiv:2007.10950.
- [82] The reader should be aware that those papers use different terminology from this paper. In Refs. [79,80], each degree of freedom $v \in V$ is called a “subsystem,” and the modular decomposition \mathcal{C} is called a “unit structure,” while each $A \in \mathcal{C}$ is called a “unit.”
- [83] M. Lostaglio, D. Jennings, and T. Rudolph, *Description of Quantum Coherence in Thermodynamic Processes Requires Constraints beyond Free Energy*, *Nat. Commun.* **6**, 6383 (2015).
- [84] J. P. Santos, L. C. Céleri, G. T. Landi, and M. Paternostro, *The Role of Quantum Coherence in Non-equilibrium Entropy Production*, *npj Quantum Inf.* **5**, 23 (2019).
- [85] C. Elphick and P. Wocjan, *Spectral Lower Bounds for the Quantum Chromatic Number of a Graph*, *J. Comb. Theory Ser. A* **168**, 338 (2019).
- [86] E. Chitambar and G. Gour, *Quantum Resource Theories*, *Rev. Mod. Phys.* **91**, 025001 (2019).
- [87] R. Gallego, J. Eisert, and H. Wilming, *Thermodynamic Work from Operational Principles*, *New J. Phys.* **18**, 103017 (2016).
- [88] F. G. S. L. Brandão, M. Horodecki, J. Oppenheim, J. M. Renes, and R. W. Spekkens, *Resource Theory of Quantum States out of Thermal Equilibrium*, *Phys. Rev. Lett.* **111**, 250404 (2013).
- [89] M. Lostaglio, M. P. Müller, and M. Pastena, *Stochastic Independence as a Resource in Small-Scale Thermodynamics*, *Phys. Rev. Lett.* **115**, 150402 (2015).
- [90] P. Faist and R. Renner, *Fundamental Work Cost of Quantum Processes*, *Phys. Rev. X* **8**, 021011 (2018).
- [91] N. Y. Halpern and J. M. Renes, *Beyond Heat Baths: Generalized Resource Theories for Small-Scale Thermodynamics*, *Phys. Rev. E* **93**, 022126 (2016).
- [92] I. Marvian and R. W. Spekkens, *Extending Noether’s Theorem by Quantifying the Asymmetry of Quantum States*, *Nat. Commun.* **5**, 3821 (2014).
- [93] I. Marvian and R. W. Spekkens, *Asymmetry Properties of Pure Quantum States*, *Phys. Rev. A* **90**, 014102 (2014).
- [94] I. Marvian and R. W. Spekkens, *Modes of Asymmetry: The Application of Harmonic Analysis to Symmetric Quantum Dynamics and Quantum Reference Frames*, *Phys. Rev. A* **90**, 062110 (2014).
- [95] Z.-W. Liu, X. Hu, and S. Lloyd, *Resource Destroying Maps*, *Phys. Rev. Lett.* **118**, 060502 (2017).
- [96] H. K. Lee, C. Kwon, and H. Park, *Fluctuation Theorems and Entropy Production with Odd-Parity Variables*, *Phys. Rev. Lett.* **110**, 050602 (2013).
- [97] S. Ito, *Stochastic Thermodynamic Interpretation of Information Geometry*, *Phys. Rev. Lett.* **121**, 030605 (2018).
- [98] K. Takahashi, *Shortcuts to Adiabaticity Applied to Non-equilibrium Entropy Production: An Information Geometry Viewpoint*, *New J. Phys.* **19**, 115007 (2017).
- [99] S. Ito, M. Oizumi, and S.-I. Amari, *Unified Framework for the Entropy Production and the Stochastic Interaction Based on Information Geometry*, *Phys. Rev. Research* **2**, 033048 (2020).
- [100] S. B. Nicholson, A. del Campo, and J. R. Green, *Non-equilibrium Uncertainty Principle from Information Geometry*, *Phys. Rev. E* **98**, 032106 (2018).
- [101] S. Ito and A. Dechant, *Stochastic Time Evolution, Information Geometry, and the Cramér-Rao Bound*, *Phys. Rev. X* **10**, 021056 (2020).
- [102] T. Nakamura, H. Hasegawa, and D. Driebe, *Reconsideration of the Generalized Second Law Based on Information Geometry*, *J. Phys. Commun.* **3**, 015015 (2019).
- [103] D. Petz, *Quantum Information Theory and Quantum Statistics*, Theoretical and Mathematical Physics (Springer, Berlin, 2008).
- [104] E. T. Jaynes, *The Gibbs Paradox*, in *Maximum Entropy and Bayesian Methods* (Springer, New York, 1992), pp. 1–21.
- [105] M. Weilenmann, L. Kraemer, P. Faist, and R. Renner, *Axiomatic Relation between Thermodynamic and Information-Theoretic Entropies*, *Phys. Rev. Lett.* **117**, 260601 (2016).
- [106] I. Csiszar and J. Körner, *Information Theory: Coding Theorems for Discrete Memoryless Systems* (Cambridge University Press, Cambridge, England, 2011).
- [107] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (John Wiley & Sons, New York, 2006).
- [108] C. Van den Broeck and M. Esposito, *Ensemble and Trajectory Thermodynamics: A Brief Introduction*, *Physica (Amsterdam)* **418A**, 6 (2015).
- [109] U. Seifert, *Entropy Production along a Stochastic Trajectory and an Integral Fluctuation Theorem*, *Phys. Rev. Lett.* **95**, 040602 (2005).
- [110] M. Esposito and C. Van den Broeck, *Three Detailed Fluctuation Theorems*, *Phys. Rev. Lett.* **104**, 090601 (2010).
- [111] R. E. Spinney and I. J. Ford, *Entropy Production in Full Phase Space for Continuous Stochastic Dynamics*, *Phys. Rev. E* **85**, 051113 (2012).
- [112] Y. Murashita, K. Funo, and M. Ueda, *Nonequilibrium Equalities in Absolutely Irreversible Processes*, *Phys. Rev. E* **90**, 042110 (2014).
- [113] C. Jarzynski, *Equalities and Inequalities: Irreversibility and the Second Law of Thermodynamics at the Nanoscale*, *Annu. Rev. Condens. Matter Phys.* **2**, 329 (2011).
- [114] J. J. Benedetto and W. Czaja, *Integration and Modern Analysis* (Birkhäuser, Boston, 2009).

-
- [115] A. C. Barato and U. Seifert, *Coherence of Biochemical Oscillations Is Bounded by Driving Force and Network Topology*, *Phys. Rev. E* **95**, 062409 (2017).
- [116] C. N. Yang, *The Spontaneous Magnetization of a Two-Dimensional Ising Model*, *Phys. Rev.* **85**, 808 (1952).
- [117] K.-J. Engel and R. Nagel, *One-Parameter Semigroups for Linear Evolution Equations*, Graduate Texts in Mathematics (Springer-Verlag, New York, 2000).
- [118] A. Gomez-Marín, J. Parrondo, and C. Van den Broeck, *Lower Bounds on Dissipation upon Coarse Graining*, *Phys. Rev. E* **78**, 011107 (2008).