Entropy production given constraints on the energy functions

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

(Received 11 May 2021; revised 3 August 2021; accepted 25 August 2021; published 23 September 2021)

We consider the problem of driving a finite-state classical system from some initial distribution p to some final distribution p' with vanishing entropy production (EP), under the constraint that the driving protocols can only use some limited set of energy functions $\mathcal{E}$. Assuming no other constraints on the driving protocol, we derive a simple condition that guarantees that such a transformation can be carried out, which is stated in terms of the smallest probabilities in $\{p, p'\}$ and a graph-theoretic property defined in terms of $\mathcal{E}$. Our results imply that a surprisingly small amount of control over the energy function is sufficient (in particular, any transformation $p \to p'$ can be carried out as soon as one can control some one-dimensional parameter of the energy function, e.g., the energy of a single state). We also derive a lower bound on the EP under more general constraints on the transition rates, which is formulated in terms of a convex optimization problem.

DOI: 10.1103/PhysRevE.104.034129

I. INTRODUCTION

Entropy production (EP) refers to the total increase of the entropy of a system and its environment during a physical process. EP is the fundamental measure of thermodynamic inefficiency [1], and in particular, the amount of work that can be extracted during a transformation between two thermodynamic states decreases as EP increases [2]. For this reason, one of the central issues in thermodynamics involves characterizing the minimal amount of EP needed to transform a system between two thermodynamic states.

For concreteness, consider a system coupled to a work reservoir and a single heat bath at inverse temperature β. Suppose that one wishes to drive this system from some initial distribution p at time $t = 0$ to some final distribution p' at time $t = \tau$ while minimizing EP. In general, this can be accomplished by driving the system along a continuous trajectory of energy functions that sends

$$E(0) = -\beta^{-1} \ln p \;\to\; E(\tau) = -\beta^{-1} \ln p'. \tag{1}$$

In the quasistatic limit $\tau \to \infty$, this will carry out the transformation $p \to p'$ up to an arbitrary accuracy and for an arbitrarily small amount of EP [3,4]. However, implementing such a driving protocol may be impossible in many real-world situations, where there are often strong limitations on the ability to control the energy function.

Here we investigate which transformations can be carried out without EP, assuming that one is constrained to use some limited set of energy functions, but is otherwise free to use any thermodynamically consistent driving protocol. We consider a finite-state classical system which evolves according to a Markovian master equation over time $t \in [0, \tau]$,

$$\frac{d}{dt} p_i(t) = \sum_{j(\neq i)} \big[W_{ij}(t) p_j(t) - W_{ji}(t) p_i(t)\big], \tag{2}$$

where $p_i(t)$ is the probability of state i at time t and $W_{ji}(t)$ is the transition rate from state i to state j at time t. As standard in stochastic thermodynamics [5], we assume that the rate matrix obeys local detailed balance (LDB) at all times t, so that the transition rates can be written in the following form [6, Ch. 2]:

$$W_{ji}(t) = \psi_{ji}(t) e^{\beta[E_i(t) - E_j(t)]/2}, \tag{3}$$

where $E(t)$ is the energy function at time t and $\psi_{ji}(t) = \psi_{ij}(t) \geqslant 0$ is a symmetric positive function that controls the overall timescale of transitions between states i and j, sometimes called the *activity* [6]. Equation (3) implies that the rate matrix $W(t)$ has a Boltzmann equilibrium distribution, $\pi^{E(t)} = e^{-\beta E(t)}/Z$. We use the term *(driving) protocol* to refer to a time-dependent trajectory of LDB-obeying rate matrices and energy functions, $\Gamma = \{(E(t), W(t)) : t \in [0, \tau]\}$. (The inverse temperature β associated with the driving protocol is left implicit in our notation.)

Suppose that there is some limited set of energy functions $\mathcal{E}$ that one can impose on the system. Given some desired transformation $p \to p'$, we ask whether there is a driving protocol that carries out this transformation for a vanishing amount of EP, while obeying the constraint that $E(t) \in \mathcal{E}$ at all $t \in [0, \tau]$.

Our first result, presented in Sec. IV, is a simple sufficient condition that guarantees that there is such a protocol. This sufficient condition is stated in terms of the minimum probability values in p and p', as well as a simple graph-theoretic property defined in terms of $\mathcal{E}$.

As we show, this result implies that a surprisingly limited amount of control over the energy function is sufficient to carry out an arbitrary transformation $p \to p'$ with vanishing EP. In particular, it typically suffices to manipulate the

[*] artemyk@gmail.com
[†] Also at Complexity Science Hub, Vienna; Arizona State University, Tempe, AZ, USA, http://davidwolpert.weebly.com

energy along some arbitrary one-dimensional control parameter (for example, the magnetic field applied to an Ising model, or the energy assigned to some arbitrary single state). This can be contrasted to the driving protocol in Eq. (1), which requires access to energy functions that are tailored to both p and p', rather than being restricted to an arbitrary fixed one-dimensional set.

As a motivating example, consider a system composed of two subsystems $X \times Y$, and suppose that $\mathcal{E}$ only contains decoupled energy functions such as

$$E_{x,y} = E_x + E_y. \tag{4}$$

Imagine that one wants to bring the system from some initial equilibrium distribution $p = \pi^E$ for some $E \in \mathcal{E}$ to some final distribution p'. Suppose that the final distribution has correlations between the two subsystems, such that the associated mutual information obeys $I_{p'}(X;Y) > 0$. For any decoupled energy function as in Eq. (4), the equilibrium distribution is product distribution, with zero mutual information. Intuitively, it might seem impossible to introduce correlations between X and Y, as one must do to bring the system to the desired ending distribution p', while driving the system using only decoupled energy functions. As we show below, this intuition is wrong: there is a driving protocol that increases correlations between X and Y, incurs vanishing EP, and only uses decoupled energy functions (in fact, it suffices to only manipulate the energy of a single state of a single subsystem).

Conversely, imagine that one wants to use the same set of decoupled energy functions $\mathcal{E}$ to bring the system from some correlated initial distribution p, such that $I_p(X;Y) > 0$, to a final equilibrium distribution $p' = \pi^E$ for some $E \in \mathcal{E}$. One way to achieve this transformation would be to apply the energy function E to the system and then let it relax freely to the final equilibrium distribution π^E. However, under such a free relaxation, any mutual information in the initial distribution would be dissipated as EP [7]. Intuition might suggest that for any protocol that uses decoupled energy functions, EP can always be bounded in terms of the mutual information $I_p(X;Y)$. (In fact, such bounds have been previously derived under related but different assumptions, involving constraints not only on the energy functions but also some other parameters that determine the rate matrices [8–13].) It turns out, however, that initial correlations do not have to be dissipated as EP, even if only decoupled energy functions are available. Rather, as we show below, there is a driving protocol that carries out the transformation $p \to \pi^E$ without EP while only using decoupled energy functions.

It is important to emphasize that our first result is proved by construction, in which we assume that there are no other constraints on the protocol beyond that the energy function must belong to $\mathcal{E}$. In particular, we assume that one can impose a separation of timescales and fix any subset of the transition rates W_{ij} to zero (i.e., by setting the corresponding activity parameters to zero). For this reason, our construction can be seen as an illustration of the limits of what is allowed by the laws of stochastic thermodynamics given constraints on the energy functions (rather than an example of a protocol that is easy to implement in practice).

Of course, in many real-world scenarios, there are also other constraints on the rate matrices, and in particular it may not be possible to set certain rates to zero. In our second result, presented in Sec. V, we derive a lower bound on the minimal EP that must be incurred when carrying out the transformation $p \to p'$ given more general constraints on the available rate matrices. This bound is stated in terms of a convex optimization problem, which can sometimes be solved using standard numerical techniques.

For simplicity of presentation, we introduce our results in the context of a system coupled to a single heat bath and subject to conservative forces, as arising from some time-dependent energy potential $E(t)$. However, our results can be generalized to systems coupled to multiple reservoirs and/or subject to nonconservative forces. In Appendix B, we present a generalization of our first result for nonconservative forces, including a simple sufficient condition that guarantees that a given transformation $p \to p'$ can be carried out for a vanishing amount of EP, assuming that there are constraints on both the energy functions and the nonconservative forces that can be applied to a given system. In Appendix F, we present a generalization of our second result, involving a bound on EP under a general set of constraints on the rate matrices, for a system coupled to any number of reservoirs.

II. PRIOR WORK

This paper extends our recent work [13], in which we derived bounds on EP for systems that evolve under rate matrices that have symmetrical, modular, or coarse-grained structure.

In other relevant literature, Ref. [14] analyzed EP in the presence of constraints on the Hamiltonian for a quantum system coupled to a finite-sized heat bath. That paper derived bounds for special protocols that consist of sequences of unitary transformations of the system+bath and total relaxations of the system to an equilibrium state. In contrast, we analyze a classical system coupled to an idealized (i.e., infinite) reservoir and consider a much broader set of protocols than just sequences of the two kinds of operation considered in [14].

The derivation of our first main result, Theorem 1, is based on decomposing a general transformation $p \to p'$ into a sequence of simpler transformations over two-state subsystems. Similar constructions have been used to analyze allowed transformations and work extraction in quantum systems [15–17]. In classical systems, similar constructions have been used to study how a logical input-output map can be implemented using a continuous-time Markovian process [18–20].

Finally, at a broader level, our paper complements previous analysis of optimal heat and EP generation under realistic constraints on the driving protocols, such as finite time [16,21–30], stochastically fluctuating control [31], finite-sized work reservoirs [32], and a limited set of "free operations" as considered in quantum resource theory of thermodynamics [33].

III. NOTATION AND PRELIMINARIES

We consider a system with n states which undergoes a driving protocol. We use notation like $p, p(t), \ldots \in \mathbb{R}^n_+$ to refer to (possibly time-dependent) probability distributions over states. Similarly, we use notation like $E, E(t), \ldots \in \mathbb{R}^n$ to refer to (possibly time-dependent) energy functions over the states. The probability of a particular state i is indicated

with subscript notation [p_i, $p_i(t)$, etc.], and similarly for the energy level of a particular state [E_i, $E_i(t)$, etc.]. We use $\delta_{i,j}$ to indicate the Kronecker delta function.

We use the Kullback-Leibler (KL) divergence, an information-theoretic non-negative measure of the difference between distributions. The KL divergence from distribution p to distribution q is defined as

$$D(p\|q) = \sum_i p_i \ln \frac{p_i}{q_i}. \tag{5}$$

Given a driving protocol $\Gamma = \{(E(t), W(t)) : t \in [0, \tau]\}$, the incurred EP on initial distribution p is given by

$$\Sigma(p, \Gamma) = \int_0^\tau \dot{\Sigma}(p(t), W(t))\, dt, \tag{6}$$

where $p(t)$ is the distribution at time t, as determined by the master equation in Eq. (2) under the initial condition $p(0) = p$, and $\dot{\Sigma}[p(t), W(t)]$ is the instantaneous EP rate at time t. In stochastic thermodynamics, the instantaneous EP rate incurred by some rate matrix W and distribution p is given by [34,35]

$$\dot{\Sigma}(p, W) = \frac{1}{2} \sum_{i,j} (p_i W_{ji} - p_j W_{ij}) \ln \frac{p_i W_{ji}}{p_j W_{ij}} \geqslant 0. \tag{7}$$

IV. EP UNDER ENERGY CONSTRAINTS

Consider some driving protocol $\Gamma = \{(E(t), W(t)) : t \in [0, \tau]\}$ which obeys LDB, and suppose that only a restricted set of energy functions is available, so that $E(t) \in \mathcal{E} \subseteq \mathbb{R}^n$ at all times $t \in [0, \tau]$. Given this constraint, we investigate whether a given transformation $p \to p'$ can be carried out while achieving an arbitrarily small amount of EP. We make the weak assumption that the set $\mathcal{E}$ is closed and *path-connected*, meaning that any two elements of $\mathcal{E}$ can be connected by a continuous curve in $\mathcal{E}$.

Before proceeding, we define the concept of the *controllable energy gap*, which will play a central role in our analysis. The controllable energy gap between a pair of states i, j is

$$g_{ij}(\mathcal{E}) = \min\Big\{ \max_{E \in \mathcal{E}} (E_i - E_j), \max_{E \in \mathcal{E}} (E_j - E_i) \Big\}. \tag{8}$$

In words, $g_{ij}(\mathcal{E})$ quantifies how much the energy gap between states i and j can be varied by choosing among different $E \in \mathcal{E}$. As an illustration, if the set of available energy functions is a line segment, $\mathcal{E} = \{\lambda E^\varnothing : \lambda \in [-1, 1]\}$ for some fixed energy function $E^\varnothing$, then $g_{ij}(\mathcal{E}) = |E_j^\varnothing - E_i^\varnothing|$.

We first consider the special case where going from p to p' only involves moving Δ probability from some state a to some other state b, so that

$$p'_i = p_i + \Delta(\delta_{b,i} - \delta_{a,i}). \tag{9}$$

We then show how to construct a special driving protocol which carries out this transformation while incurring an arbitrarily small amount of EP, and while only using energy functions $E \in \mathcal{E}$ along with rate matrices that obey LDB.

To begin, recall that one can always set the transition rates going both forward and backward between any two states i, j to zero without violating LDB [i.e., by setting the corresponding activity parameters $\psi_{ij}(t) = \psi_{ji}(t) = 0$ in Eq. (3)]. This allows us to construct a driving protocol such that at all t, only transitions between states a and b are allowed, because $W_{ji} = 0$ whenever $i \notin \{a, b\}$ or $j \notin \{a, b\}$. In this case, only the pair of states a, b contributes to the EP rate in Eq. (7), allowing us to write the EP rate at time t as

$$\dot{\Sigma}(p(t), W(t)) = [p_a(t) W_{ba}(t) - p_b(t) W_{ab}(t)] \times \left[\ln \frac{p_a(t)}{p_b(t)} - \beta[E_a(t) - E_b(t)] \right], \tag{10}$$

where we used Eq. (7) as well as LDB, Eq. (3). Now suppose that the controllable energy gap between a and b obeys

$$g_{ab}(\mathcal{E}) > \beta^{-1} \max\left\{ \left| \ln \frac{p_a}{p_b} \right|, \left| \ln \frac{p'_a}{p'_b} \right| \right\}. \tag{11}$$

Given the definition of $g_{ab}(\mathcal{E})$ in Eq. (8), as well as the assumption that $\mathcal{E}$ is path-connected, there must be some $E \in \mathcal{E}$ such that $E_a - E_b$ is equal to any desired value between $\beta^{-1}|\ln(p_a/p_b)|$ and $\beta^{-1}|\ln(p'_a/p'_b)|$. Intuitively, this suggests that at any time t, one can choose the energy function within $E(t) \in \mathcal{E}$ so that the bracketed term in Eq. (10) becomes arbitrarily small. In fact, in the quasistatic limit of $\tau \to \infty$, the energy function can be varied in such a way that the total EP incurred over $t \in [0, \tau]$ becomes arbitrarily small. By formalizing this intuition, we derive the following result, which is proved in Appendix A.

Proposition 1. If Eqs. (9) and (11) hold, then for any $\epsilon > 0$, there is a protocol $\Gamma = \{(E(t), W(t)) : t \in [0, \tau]\}$ that sends $p \to p'$ while obeying $\Sigma(p, \Gamma) \leqslant \epsilon$ and $E(t) \in \mathcal{E}$ at all t.

We refer to the driving protocol constructed in the proof of Proposition 1, which moves probability between two states while incurring a vanishing amount of EP, as a *transfer*.

We are now ready to prove our first main result, which states that one can drive the system from any initial distribution p to any final distribution p' by chaining together an appropriate sequence of transfers. Moreover, because each transfer can be carried out quasistatically and thereby incur an arbitrarily small amount of EP, the overall sequence of transfers can also be made to incur an arbitrarily small amount of EP.

There are many protocols that can carry out the transformation $p \to p'$ using a sequence of transfers. One relatively simple one involves a two stage process, illustrated in Fig. 1. In the first stage, we pick one particular "buffer state" (without loss of generality, this can be state 1) to accumulate probability from states with excess probability (this accumulation is done via a sequence of transfers). In the second stage, the probability accumulated in the buffer state is distributed to states that need probability (this distribution is again done via a sequence of transfers). This procedure is outlined more formally as follows.

(Stage 1) Consider in turn each state i such that $p_i > p'_i$. For each such i, we will move $\Delta_i = p_i - p'_i$ of probability from state i to state 1 (if $i = 1$, do nothing). To do so, select some *path* (sequence of states) of length ℓ_i that starts on state i and ends on state 1,

$$\vec{x} = (x_1 = i \to x_2 \to \cdots \to x_{\ell_i - 1} \to x_{\ell_i} = 1),$$

and then run $\ell_i - 1$ transfers, each of which moves Δ_i of probability from state x_k to state x_{k+1} during a temporal interval of length τ.

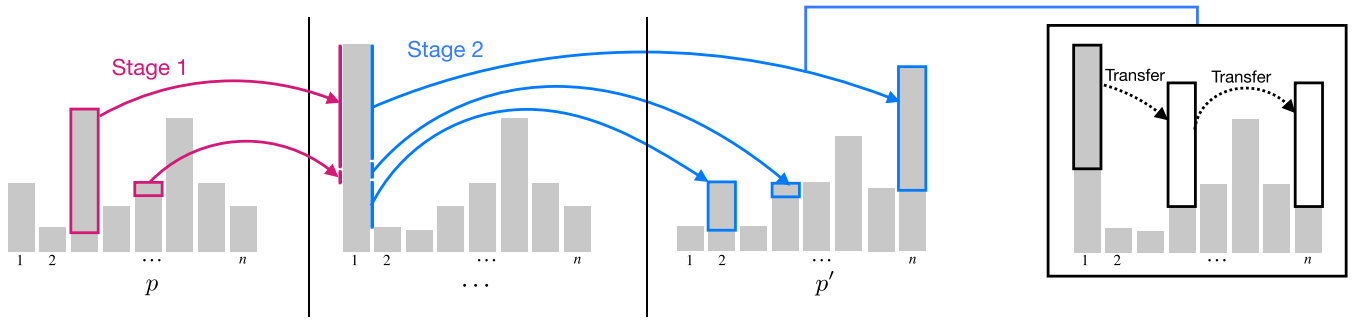

FIG. 1. A two stage construction used to carry out a transformation $p \to p'$ using a sequence of transfers. In stage 1, for all states i with excess probability, $p_i > p'_i$, $p_i - p'_i$ of probability is moved into a buffer state (state 1). In stage 2, $p'_i - p_i$ of probability from the buffer state is moved to all states i with $p_i < p'_i$. Each probability move (red and blue arrows) is done using a sequence of state-to-state transfers along a path (box on right).

(Stage 2) Consider in turn each state i such that $p_i < p'_i$. For each such i, move $\Delta_i = p'_i - p_i$ of probability from state 1 to state i (if $i = 1$, do nothing). To do so, select some path of length ℓ_i that starts on state 1 and ends on state i,

$$\vec{x} = (x_1 = 1 \to x_2 \to \cdots \to x_{\ell_i - 1} \to x_{\ell_i} = i),$$

and then run $\ell_i - 1$ transfers, each of which moves Δ_i of probability from state x_k to state x_{k+1} during a temporal interval of length τ.

It is easy to see that this construction will transform p to p'. In addition, as mentioned above, each individual transfer can be made to have an arbitrarily small amount of EP by taking τ (the time taken by each transfer) to be sufficiently large. Finally, because the construction involves only a finite number of transfers [36], the overall procedure can be made to incur an arbitrarily small amount of EP.

The construction described above can be carried out as long as the controllable energy gaps are large enough for each transfer to be feasible—in other words, as long as the inequality in Eq. (11) is satisfied for each transfer. We now derive a simple sufficient condition that guarantees that this inequality is satisfied for each transfer.

First, observe that at no step in the above construction does the probability concentrated in any state i drop below $\min\{p_i, p'_i\}$. This implies that for any time t at which a transfer begins or ends, the probability concentrated in state i obeys

$$p_i(t) \geqslant \min\{p_i, p'_i\} \geqslant \min_j \min\{p_j, p'_j\}. \tag{12}$$

Next, observe that in both stages of the above construction, one has the freedom to choose the specific path to or from state 1 and state i. In fact, one can choose paths optimally so that the controllable energy gap $g_{x_k x_{k+1}}(\mathcal{E})$ across each transfer is as large possible, since that increases the set of transformations allowed by Eq. (11). Let $\mathcal{P}(i, j)$ be the set of all paths from state i to state j, and define the *capacity* of any such path $\vec{x} = (x_1 = 1, x_2, \ldots, x_m = j) \in \mathcal{P}(i, j)$ as the minimum controllable energy gap in the path, $\mathsf{c}(\vec{x}) := \min_k g_{x_k x_{k+1}}(\mathcal{E})$. The capacity of the optimal path between the worst-case pair of states is then given by

$$\mathsf{C}(\mathcal{E}) := \min_{i \neq j} \max_{\vec{x} \in \mathcal{P}(i,j)} \mathsf{c}(\vec{x}). \tag{13}$$

In graph theory, $\mathsf{C}(\mathcal{E})$ is known as the *maximum capacity* of the undirected graph with n vertices the edge weights of which are specified by $g_{ij}(\mathcal{E})$ [37]. A classical result states that if paths are chosen from a *maximum spanning tree* of the graph, then no path will have capacity less than $\mathsf{C}(\mathcal{E})$ [38]. This means that in practice $\mathsf{C}(\mathcal{E})$ can be calculated as

$$\mathsf{C}(\mathcal{E}) := \min_{(i,j) \in T} g_{ij}(\mathcal{E}) \tag{14}$$

where T is any maximum spanning tree.

To summarize, suppose that the two stage construction described above chooses paths between state i and state 1 that lie along a maximum spanning tree of the graph with edge weights $g_{ij}(\mathcal{E})$, as illustrated in Fig. 2. Then, by Eq. (14), the controllable energy gap involved in any transfer will obey

$$g_{x_k x_{k+1}}(\mathcal{E}) \geqslant \mathsf{C}(\mathcal{E}). \tag{15}$$

Finally, suppose that the probabilities in p_i and p'_i satisfy the following condition:

$$\min_i \min\{p_i, p'_i\} > e^{-\beta \mathsf{C}(\mathcal{E})}. \tag{16}$$

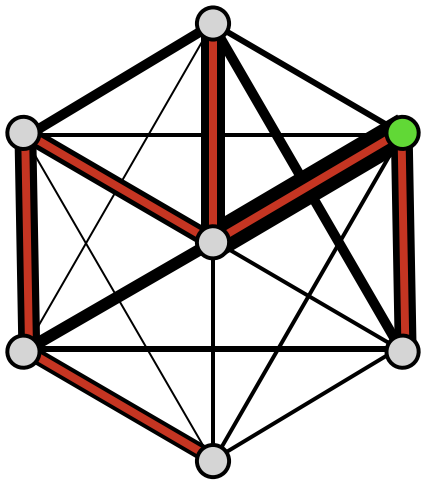

FIG. 2. Consider a weighted graph with n vertices, where each vertex represents a state of the system and edge weights are given by the controllable energy gaps $g_{ij}(\mathcal{E})$. If the two stage construction shown in Fig. 1 chooses paths to or from the buffer state (green node) that lie within the maximum spanning tree of the graph (red edges), then the controllable energy gap across each transfer will be lower bounded by the maximum capacity of the graph, as stated in Eq. (15).

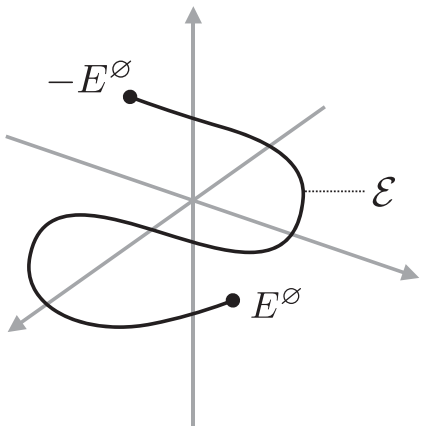

FIG. 3. Suppose that the set of available energy functions $\mathcal{E} \subseteq \mathbb{R}^n$ contains some fixed one-dimensional curve of energy functions from $E^{\varnothing}$ to $-E^{\varnothing}$. Then any transformation $p \to p'$ can be carried out for vanishing EP, as long as p and p' have full support and $E^{\varnothing}$ has a large enough gap between highest and lowest energy levels, Eq. (17).

Combining this inequality with Eqs. (12) and (15) and rearranging implies $g_{x_k x_{k+1}}(\mathcal{E}) > -\beta^{-1} \ln p_i(t)$ for any state i and any time point t at which a transfer begins or ends. Since $p_i(t) \in [0, 1]$ for all i and t, this in turn implies that Eq. (11) holds for each transfer. This leads to our first main result, which is proved informally using the construction and arguments outlined above.

Theorem 1. If Eq. (16) is satisfied, then for any $\epsilon > 0$ there is a protocol $\Gamma = \{(E(t), W(t)) : t \in [0, \tau]\}$ that carries out $p \to p'$ while obeying $\Sigma(p, \Gamma) \leqslant \epsilon$ and $E(t) \in \mathcal{E}$ at all t.

The graph-theoretic quantity $\mathsf{C}(\mathcal{E})$ measures the ability of the set of available energy functions $\mathcal{E}$ to implement arbitrary transformations. As expected, this quantity is invariant if any $E \in \mathcal{E}$ is shifted by a constant [since $g_{ij}(\mathcal{E})$ is invariant under shifts such as $E \to E + \lambda$] and scales multiplicatively with multiplicative scaling of the energy functions [$\mathsf{C}(\mathcal{E}) \to \lambda \mathsf{C}(\mathcal{E})$ when $E \to \lambda E$ for all $E \in \mathcal{E}$].

Theorem 1 implies that many transformations can be carried out with vanishing EP even if only a very restricted set of energy functions is available. In particular, suppose that $\mathcal{E}$ contains a continuous curve that connects some fixed energy function $E^{\varnothing}$ and its negation $-E^{\varnothing}$, as shown in Fig. 3. This curve might involve separately varying the energy level of an individual state, or it might involve varying the entire energy function along a one-dimensional manifold without being able to vary the energy of individual states. In Appendix C, we show that for any such $\mathcal{E}$,

$$\mathsf{C}(\mathcal{E}) \geqslant \min_i \max_j |E_i^{\varnothing} - E_j^{\varnothing}| \geqslant \frac{\max_i E_i^{\varnothing} - \min_i E_i^{\varnothing}}{2}. \quad (17)$$

We also show that the first inequality is tight whenever $\mathcal{E}$ is a line segment ($\mathcal{E} = \{\lambda E^{\varnothing} : \lambda \in [-1, 1]\}$). Equation (17) implies that even a single dimension of control over the energy function suffices to carry out any transformation, as long as $E^{\varnothing}$ has a sufficiently large range of energy values. In particular, $\mathsf{C}(\mathcal{E}) \to \infty$ as the energetic gap between the highest and lowest energy state increases ($\max_i E_i^{\varnothing} - \min_i E_i^{\varnothing} \to \infty$). Thus, with a sufficiently large gap, any transformation $p \to p'$ can be carried out (as long as p and p' have full support).

Finally, as mentioned in the Introduction, one can generalize the results in the section to account for constraints on both the energy functions and the nonconservative forces which can be applied to a system. For details, see Appendix B.

A. Example

We now use Theorem 1 to analyze the examples mentioned in the introduction, which involve bringing a system containing two subsystems $X \times Y$ from an initial correlated distribution to a final uncorrelated product distribution, or vice versa, while using decoupled energy functions as in Eq. (4).

Suppose that the available energy functions only allow the manipulation of the energy of the up state of spin x from -10 to 10, while other states are at energy zero. Formally, this corresponds to a line segment of energy functions,

$$\mathcal{E} = \{\lambda E^{\varnothing} : \lambda \in [-1, 1]\} \quad \text{for} \quad E^{\varnothing}_{x,y} = 10 \cdot \delta_{x,1},$$

which is a special case of Eq. (4). Plugging into Eq. (17) and rearranging gives

$$\mathsf{C}(\mathcal{E}) = 10 \times \min_{x,y} \max_{x',y'} |\delta_{x,1} - \delta_{x',1}|. \quad (18)$$

For any x, y, we have $\max_{x',y'} |\delta_{x,1} - \delta_{x',1}| = |\delta_{x,1} - \delta_{-x,1}| = 1$, so Eq. (18) implies that $\mathsf{C}(\mathcal{E}) = 10$. Then, by Eq. (16) and Theorem 1, any transformation where all initial and final probability values are greater than $e^{-10\beta}$ is possible.

Imagine that one wanted to transform a correlated initial distribution p to an uncorrelated final distribution p'. For concreteness, assume that the inverse temperature is $\beta = 1$, that the initial distribution is

$$p_{x,y} = \begin{cases} 0.4 & \text{if } (x, y) \in \{(-1, -1), (1, 1)\} \\ 0.1 & \text{if } (x, y) \in \{(1, -1), (-1, 1)\} \end{cases}, \quad (19)$$

and that the final distribution is the uniform $p'_{x,y} = 1/4$. Note that the initial and final mutual information terms obey $I_p(X; Y) \approx 0.18$ (in units of nats) and $I_{p'}(X; Y) = 0$. It is also straightforward to verify that the entries of p and p' are greater than e^{-10}, so Eq. (16) is satisfied. Then, the transformation $p \to p'$ can be implemented using the following procedure.

(1) Choose state $(-1, -1)$ as the "buffer state." Then, transfer 0.15 probability from state (1,1) to state $(-1, -1)$ over time $t \in [0, \tau]$.

(2) Move 0.15 probability from state $(-1, -1)$ to state $(-1, 1)$ via two transfers: first from state $(-1, -1)$ to state $(1, -1)$ over time $t \in [\tau, 2\tau]$, then from state $(1, -1)$ to state $(-1, 1)$ over time $t \in [2\tau, 3\tau]$.

(3) Transfer 0.15 probability from state $(-1, -1)$ to state $(1, -1)$ over time $t \in [3\tau, 4\tau]$.

After running this procedure on the initial distribution p, the system will be left in the uniform distribution p'. Moreover, by taking τ large enough and using the construction that appears in Appendix A, each transfer can be done for an arbitrarily small amount of EP.

Conversely, imagine that one wanted to increase mutual information by transforming an uncorrelated uniform initial distribution p to a correlated final distribution p'. Suppose again that $\beta = 1$, that p is the uniform distribution, and that p' is given by the right hand side of Eq. (19). Then, this

transformation can be accomplished by running the same sequence of transfers as described above, but now in reverse: first transfer 0.15 probability $(1,-1) \to (-1,-1)$, then move 0.15 probability using two transfers $(-1,1) \to (1,-1)$ and $(1,-1) \to (-1,-1)$, and finally transfer 0.15 probability $(-1,-1) \to (1,1)$.

It is important to note that each of the above transfers involves a pair of states which differ in the state of spin X, and therefore have a controllable energy gap equal to 10. On the other hand, it is impossible to use a transfer to move probability between a pair of states which only differ in the state of spin Y, since the corresponding controllable energy gap will be zero and Eq. (11) cannot be satisfied.

B. Physical assumptions behind Theorem 1

We briefly discuss three important physical assumptions that underlie the results described above.

First, we assume that one has access to arbitrarily long driving protocols, i.e., the *quasistatic limit*. This assumption is required for any protocol that transforms some initial distribution p to a different final distribution p' while achieving arbitrarily small EP (as known from research on finite-time thermodynamics [23,29]). Interestingly, if one does not require that EP is arbitrarily small, then arbitrarily long timescales are not necessary for carrying out the transformation $p \to p'$ with complete accuracy while using a limited set of energy functions. Instead, as can be seen from the proof of Proposition 1 in Appendix A, as long as Eqs. (9) and (11) hold, then there is a protocol $\Gamma = \{(E(t), W(t)) : t \in [0,\tau]\}$ with a finite τ that carries out the transfer $p \to p'$ exactly, while obeying $E(t) \in \mathcal{E}$ at all t.

Second, we assume that one can impose a *separation of timescales* in which the transition rates between all states are set to zero, except for a given pair of states involved in a transfer. To understand the physical meaning of this assumption, note that setting a transition rate W_{ji} to zero is equivalent to setting the corresponding activity parameter ψ_{ji} in Eq. (3) to zero. Furthermore, in statistical physics, discrete-state master equations are typically (though often implicitly) derived by coarse-graining a stochastic dynamical system over a microscopic continuous phase space, as might be described by Fokker-Planck dynamics [39–41]. Under this scheme, each coarse-grained state i represents a free energy well in the underlying phase space, the corresponding energy level E_i reflects the depth of the free energy well, and the activity parameters ψ_{ji} reflect the height of the free energy barriers separating wells i and j. For instance, transition rates are often expressed using the Arrhenius form $W_{ji} = e^{-\beta[B_{ji}-E_i]}$ [40,42,43], where B_{ji} is the absolute height of the free energy barrier separating coarse-grained states i and j. This expression can be put in the form of Eq. (3) by defining $\psi_{ji} = e^{\beta[(E_i+E_j)/2-B_{ji}]}$, so that the activity parameter ψ_{ji} is determined by the relative height of the barrier separating coarse-grained states i and j. From this point of view, our second assumption means that one can impose infinite-sized barriers between all pairs of coarse-grained states, except for a given pair of states involved in a transfer.

Because we allow certain transition rates to be set to zero, the rate matrices $W(t)$ involved in our driving protocols (as constructed in Appendix A) will generally be reducible. For this reason, these rate matrices will have multiple equilibrium distributions (i.e., multiple distributions which are stationary and incur zero EP rate), beyond the unique Boltzmann equilibrium distribution prescribed by the energy function $E(t)$. In fact, by using such reducible rate matrices, our construction carries out a transformation $p \to p'$ while keeping the system (arbitrary close to) equilibrium throughout, even though those equilibria may be outside the set of the Boltzmann equilibrium distributions that correspond to the energy functions in $\mathcal{E}$.

Our third assumption is that any pair of states can be connected by a nonzero transition rate even while the other rates are set to zero (which is a kind of "converse" of our second assumption). If, as described above, the discrete-state master equation is derived by coarse-graining a microscopic continuous phase space, then (as above) the validity of this assumption depends on one's ability to manipulate free energy barriers at the microscopic level. In addition, it also depends on the geometric properties of the embedding of the coarse-grained states within the microscopic space (e.g., if the underlying microscopic continuous space is one-dimensional, then not all patterns of connectivity between states can be achieved by manipulating free energy barriers). It is possible to generalize our treatment to account for constraints on which states can or cannot be connected via nonzero transition rates [e.g., by defining $g_{ij}(\mathcal{E}) = 0$ for any pair of states i, j that cannot be connected by a nonzero transition rate], though for simplicity we do not consider this generalization in the current paper.

V. EP UNDER BOTH CONSTRAINTS ON ENERGY FUNCTIONS AND RATE MATRICES

In real-world scenarios, there are often additional restrictions on the transition rates, not only on the energy function, which can preclude the use of the protocols constructed in the previous section. For instance, such restrictions might arise because there are constraints on how the heights of the free energy barriers between coarse-grained states can be manipulated, meaning that some transition rates cannot be set to zero. Accordingly, in this section we derive our second main result, which is a bound on the EP that arises under more general constraints on the rate matrices.

As before, we consider a system coupled to a single heat bath at inverse temperature β. Suppose that one drives the system from some initial distribution p to some final distribution p', while only using energy functions in some restricted set $\mathcal{E}$ and rate matrices in some restricted set $\mathcal{W}$. We assume that there exists *some* driving protocol $\Gamma = \{(E(t), W(t)) : t \in [0,\tau]\}$ that implements the desired transformation $p \to p'$, while obeying $E(t) \in \mathcal{E}$ and $W(t) \in \mathcal{W}$ at all t. We then analyze the minimal EP that must be incurred by this protocol in terms of the properties of $\mathcal{E}$ and $\mathcal{W}$. Note that in this section, we do not assume that either $\mathcal{E}$ or $\mathcal{W}$ is path-connected (in fact, we will sometimes assume that both $\mathcal{E}$ and $\mathcal{W}$ are finite sets).

Before presenting our result, we introduce the notion of a "KL projection" of a distribution p onto the set of available equilibrium distributions, which we indicate as $\Pi(p)$ [44]. The KL projection of p refers to the Boltzmann equilibrium distribution that is closest to p in terms of KL divergence, as

defined in Eq. (5), among the set of all Boltzmann distributions allowed by the available energy functions:

$$\Pi(p) := \arg\min_{E\in\mathcal{E}} D(p\|\pi^E). \tag{20}$$

(As above, we use the notation $\pi^E = e^{-\beta E}/Z$ to indicate the Boltzmann distribution corresponding to energy function E.) Note that if p is a Boltzmann distribution for some $E \in \mathcal{E}$, then $\Pi(p) = p$.

Next, suppose that there is some number $\eta \in [0, 1]$ such that

$$\dot{\Sigma}(p, W) \geqslant -\eta \sum_{i,j} p_i W_{ji} \ln \frac{p_j}{[\Pi(p)]_j} \quad \forall p, W \in \mathcal{W}, \tag{21}$$

where $\dot{\Sigma}(p, W)$ is the rate of EP incurred by rate matrix W on distribution p, as defined in Eq. (7). This is an implicit inequality that bounds η in terms of the set of allowed rate matrices $\mathcal{W}$ as well as the KL projection $\Pi(p)$, which in turn is a function of the allowed energy functions $\mathcal{E}$. This inequality is always satisfied for $\eta = 0$, since it is then equivalent to the statement that the EP rate is non-negative. However, we will also consider sets $\mathcal{E}$ and $\mathcal{W}$ where Eq. (21) holds for some $\eta > 0$. Importantly, as long as the set of allowed rate matrices $\mathcal{W}$ is finite, it is possible to find the largest value of η that satisfies Eq. (21) via numerical convex optimization techniques. The procedure for doing this is described in detail in Appendix E, and code is provided in [45]. The right hand side of Eq. (21) is η times the rate at which the distribution p approaches the closest equilibrium distribution $\Pi(p)$ under the dynamics generated by W [see Eq. (D1) in Appendix D]. Thus, the multiplier η bounds how much faster any distribution p can approach the closest equilibrium $\Pi(p)$, relative to the rate at which that distribution incurs EP. Note also that η is independent of the overall timescale of the rate matrices $W \in \mathcal{W}$: if Eq. (21) holds for some η, some distribution p, and some rate matrix W, then it will also hold when that rate matrix is rescaled as $W \to \lambda W$ (since this is equivalent to multiplying both sides by λ).

Now consider the total EP incurred by the protocol Γ on initial distribution p, as defined in Eq. (6). Using Eq. (21) and some simple rearrangement, the EP rate at time t can be bound in terms of the time derivative of the KL divergence between $p(t)$ (the distribution at time t) and its KL projection $\Pi(p(t))$:

$$\dot{\Sigma}(p(t), W(t)) \geqslant -\eta \frac{d}{dt} D[p(t)\|\Pi(p(t))]. \tag{22}$$

(See Appendix D for details.) This leads to our second main result, which follows from integrating both sides of Eq. (22) and using the fundamental theorem of calculus.

Theorem 2. Let $\Gamma = \{(E(t), W(t)) : t \in [0, \tau]\}$ be a protocol such that $E(t) \in \mathcal{E}$ and $W(t) \in \mathcal{W}$ at all t. Then,

$$\Sigma(p, \Gamma) \geqslant \eta\{D[p\|\Pi(p)] - D[p'\|\Pi(p')]\}$$

for any $\eta \in [0, 1]$ where Eq. (21) holds, and $\Pi(\cdot)$ as in Eq. (20).

The EP bound in Theorem 2 is the product of two terms: the drop of the KL divergence $D[p\|\Pi(p)]$ and a "multiplier" η. The drop of KL divergence reflects the contribution to EP arising from constraints on the energy functions, as determined by the set $\mathcal{E}$. In particular, if all energy functions are available, then $\Pi(p) = p$ and this KL divergence term vanishes, leading to a trivial bound $\Sigma \geqslant 0$. The multiplier η reflects the contribution to the EP bound that arises from constraints on the rate matrices, as determined both by $\mathcal{E}$ and $\mathcal{W}$. Note that Theorem 2 only gives nontrivial bounds on EP when the final distribution is closer to the set of equilibrium distributions than the initial one, so that $D[p\|\Pi(p)] > D[p'\|\Pi(p')]$.

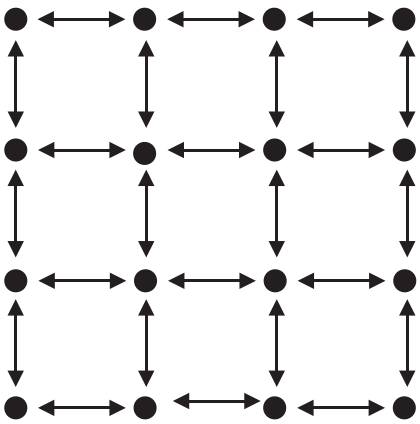

FIG. 4. A discrete-state system with 16 states arranged in a 4×4 lattice, with only nearest-neighbor transitions allowed. $\mathcal{W}$ contains 16 rate matrices, each one corresponding to an equilibrium distribution in which the energy of a particular state is increased.

Interestingly, our result holds not only when $\Pi(p)$ is defined as the KL projection to the set of equilibrium distributions, as in Eq. (20), but more generally when $\Pi(p)$ is defined as the KL projection to *any* arbitrary set of distributions Ω. In other words, if one defines $\Pi(p) = \arg\min_{q\in\Omega} D(p\|q)$ for any Ω and then finds a corresponding $\eta \in [0, 1]$ such that Eq. (21) holds, then Theorem 2 will still apply. Each set of distributions Ω will have its own maximal value of η which satisfies Eq. (21) and will therefore induce its bound on EP. It may be possible to derive tighter bounds on EP in Theorem 2 by varying the choice of the distributions Ω in the definition of $\Pi(p)$, though we leave exploration of this choice for future work.

In fact, by exploiting the freedom in how the KL projection $\Pi(p)$ can be defined, it is possible to generalize Theorem 2 to derive bounds on EP in the presence of nonconservative forces and multiple thermodynamic reservoirs (rather than a single heat bath as considered above). Details of this generalization are discussed in Appendix F.

Example

We demonstrate Theorem 2 with a discretized model of a Brownian particle on a two-dimensional lattice, where we imagine that one can only increase the energy of a single lattice site at any one time. Assume that the system is coupled to a heat bath at inverse temperature β, and that each system state i is identified with a location on a two-dimensional $N \times N$ lattice, as in Fig. 4. Suppose that there are N^2 rate matrices in $\mathcal{W}$ and N^2 energy functions in $\mathcal{E}$, one for each location. The energy function corresponding to location i assigns energy 1 to state i and energy zero to all other states, $E_j^{(i)} = \delta_{i,j}$. The rate matrix corresponding to state i allows only nearest-neighbor transitions: the off-diagonal entries obey $W_{kj}^{(i)} = e^{\beta E_j^{(i)}}$ when

states j and k are nearest neighbors on the two-dimensional lattice, and otherwise $W^{(i)}_{kj} = 0$.

We now use Theorem 2 to derive a bound on EP for this system. For concreteness, we take $\beta = 1$ and $N = 4$. Using the convex optimization technique described in Appendix E, we find that the tightest bound in Eq. (21) is given by $\eta \approx 0.6$. Thus, for any transformation $p \to p'$,

$$\Sigma \geqslant 0.6\{D[p\|\Pi(p)] - D[p'\|\Pi(p')]\}. \tag{23}$$

Given the structure of the available energy functions, the KL divergence to any equilibrium distribution $\pi^{E^{(i)}} = e^{-E^{(i)}}/Z$ can be written as

$$D(p\|\pi^{E^{(i)}}) = p_i + \ln Z - S(p),$$

where $Z = \sum_j e^{-\delta_{i,j}} = (N^2 - 1) + e^{-1}$ is a normalization constant. Thus, the KL divergence to the closest equilibrium distribution can be written as $D[p\|\Pi(p)] = \min_i p_i + \ln Z - S(p)$. Plugged into Eq. (23), this gives the bound

$$\Sigma \geqslant 0.6[S(p') - S(p) + (\min_i p_i - \min_i p'_i)].$$

In words, for this set of constrained driving protocols, at least 0.6 of the increase in the entropy of the system, plus the drop in the minimum probability value in going from p to p', must be dissipated as EP.

VI. DISCUSSION AND FUTURE WORK

In this paper, we considered a finite-state classical system that undergoes a driving protocol, represented as a time-dependent master equation. We derived two results concerning the feasibility of transforming such a system from some initial distribution p to a final distribution p' given constraints on the driving protocol, and the associated EP.

Our first result, presented in Sec. IV, concerned the case where the driving protocol can only access a restricted set of energy functions $\mathcal{E}$, but can otherwise make use of arbitrary rate matrices (as long as they obey local detailed balance). We showed that any given transformation $p \to p'$ can be implemented for a vanishing amount of EP, so long as the probabilities in p and p' are bounded away from zero by an amount which depends on $\mathcal{E}$. We also demonstrated that this condition is quite weak, since even a limited amount of control over the energy function (such as the ability to manipulate the energy of a single state) suffices to implement an arbitrary transformation with vanishing EP. This result is derived under the assumption that one can control all other aspects of the rate matrices beyond the energy function, including the ability to set arbitrary transition rates to zero or nonzero values.

In our second result, presented in Sec. V, we derived a lower bound on the EP involved in carrying out some transformation $p \to p'$ in the presence of constraints on both the energy functions and the rate matrices. We show that in some cases, this bound can be determined numerically using standard convex optimization techniques.

We briefly mention some possible directions for future work. First, in deriving our first result, we do not prove that Eq. (16) is a *necessary* condition for carrying out the transformation $p \to p'$ with vanishing EP while using energy functions in $\mathcal{E}$, only that it is a *sufficient* condition. An interesting research direction for future work would investigate sufficient and necessary conditions for carrying out such a transformation.

Second, in deriving our first result, we constructed a protocol that carries out $p \to p'$ without restricting how the protocol behaves on other initial distributions besides p. Future work may consider the related, but more difficult, problem of carrying out a *logical map* (a conditional probability distribution) T_{ij} from initial states j to final states i while using a limited set of energy functions. In other words, this research direction would analyze driving protocols which (1) implement some desired logical map T_{ij}, (2) achieve vanishing EP on some particular initial distribution p [46], and (3) only use energy functions from some limited set $\mathcal{E}$. It is possible that this problem can be tackled by combining the "transfer" protocol used in Proposition 1 with the constructions developed in [18,19], which show how to implement a given logical map with continuous-time master equations (while possibly using some number of auxiliary "hidden" states).

Third, in this paper we considered finite-state master equation dynamics. Future work may investigate whether our results can be extended to continuous-state dynamics, such as Fokker-Planck dynamics over probability densities. Such an extension would be far from trivial, as our first result relies on the ability to set arbitrary transition rates to zero or nonzero values (which is inappropriate for Fokker-Planck dynamics, even when discretized). On the other hand, our second result can be formally generalized to continuous-state dynamics, but it is not clear when the corresponding infinite-dimensional optimization problem (as defined in Appendix E) can be solved in practice. We note that in previous work [13], we investigated bounds on EP for Fokker-Planck dynamics under highly structured constraints on the energy functions, such as symmetry, modularity, and coarse-graining constraints.

Finally, future work may consider whether the methods developed here can be extended to Markovian quantum systems. This research direction would analyze an open quantum system evolving according to time-inhomogeneous Lindbladian dynamics [47], and investigate how constraints on the available Hamiltonians and Lindbladian operators translate into bounds on the quantum EP involved in bringing the system from some initial mixed state ρ to some final mixed state ρ'.

ACKNOWLEDGMENTS

We thank M. Esposito, A. Boyd, and the anonymous reviewers for helpful discussions and suggestions. 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: PROOF OF PROPOSITION 1

We prove the result by construction. In particular, we first construct a family of protocols $\Gamma_\tau = \{(E^\tau(t), W^\tau(t)) : t \in [0, \tau]\}$ parametrized by temporal duration τ (where $\tau \to \infty$ corresponds to the quasistatic limit). We then show that if Eqs.

(9) and (11) hold, then for any $\epsilon > 0$, there is some τ such that the protocol Γ_τ satisfies the following three conditions.

(I) The protocol maps the initial distribution p at time $t = 0$ to the final distribution p' at time $t = \tau$, where

$$p'_i = p_i + \Delta(\delta_{b,i} - \delta_{a,i}). \tag{A1}$$

(II) The EP obeys $\Sigma(p, \Gamma_\tau) \leqslant \epsilon$.

(III) The energy functions obey $E^\tau(t) \in \mathcal{E}$ at all t.

We assume below that $\Delta \geqslant 0$ in Eq. (A1), so that $p'_a \leqslant p_a$. Note that this is done without loss of generality (if $\Delta < 0$, one can swap the roles of states a and b).

We begin with a few definitions. First, for any $\tau > 0$, define the function $f_\tau : [0, \tau] \to \mathbb{R}$:

$$f_\tau(t) := p_a + \frac{t}{\tau - 1 + e^{-\tau}}(p'_a - p_a). \tag{A2}$$

Second, let $C = p_a + p_b = p'_a + p'_b$ refer to the sum of the probability in states a and b. Third, define a time-dependent rate matrix W^τ with the following entries for $i \neq j$:

$$W^\tau_{ij}(t) = \begin{cases} f_\tau(t)/C & \text{if } i = a,\ j = b \\ 1 - f_\tau(t)/C & \text{if } i = b,\ j = a \\ 0 & \text{otherwise} \end{cases}. \tag{A3}$$

Given Eq. (A2) and the inequality $\tau - 1 + e^{-\tau} > 0$, as well as the assumption that $p'_a \leqslant p_a$, it can be verified that $f_\tau(t)$ is monotonically decreasing in t, so

$$p_a + \frac{\tau}{\tau - 1 + e^{-\tau}}(p'_a - p_a) = f_\tau(\tau) \leqslant f_\tau(t) \leqslant f_\tau(0) = p_a. \tag{A4}$$

Observe that the lower bound above converges to p'_a as $\tau \to \infty$ and that $p'_a > 0$, which follows Eq. (11). In the following, we will assume that τ is large enough so that

$$f_\tau(\tau) \geqslant p'_a/2 > 0 \quad \forall t \in [0, \tau]. \tag{A5}$$

Note that given Eqs. (A3) and (A4), this implies that the transition rates in W^τ are always bounded between 0 and 1.

Let $p^\tau(t)$ indicate a time-dependent probability which undergoes the Markovian evolution $\frac{d}{dt}p^\tau(t) = W^\tau(t)p^\tau(t)$ starting from the initial condition $p^\tau(0) = p$. It is clear from Eq. (A3) that $\frac{d}{dt}p^\tau_i(t) = 0$ for any $i \notin \{a, b\}$, therefore the sum of the probability of states a, b is conserved over time:

$$p^\tau_a(t) + p^\tau_b(t) = C \quad \forall t \in [0, \tau]. \tag{A6}$$

Next, write the temporal derivative of $p^\tau_a(t)$ as

$$\begin{aligned} \frac{d}{dt}p^\tau_a(t) &= W^\tau_{ab}(t)p^\tau_b(t) - W^\tau_{ba}(t)p^\tau_a(t) \\ &= [f_\tau(t)/C]\big[C - p^\tau_a(t)\big] - [1 - f_\tau(t)/C]p^\tau_a(t) \\ &= f_\tau(t) - p^\tau_a(t), \end{aligned} \tag{A7}$$

where we used Eqs. (A3) and (A6). The differential equation in Eq. (A7) can be solved using some calculus (or a symbolic computation package such as MATHEMATICA) to give

$$p^\tau_a(t) = \left(1 - \frac{t - 1 + e^{-t}}{\tau - 1 + e^{-\tau}}\right)p_a + \frac{t - 1 + e^{-t}}{\tau - 1 + e^{-\tau}}p'_a. \tag{A8}$$

We now demonstrate that the protocol constructed above satisfies conditions I–III.

We begin with condition I. From Eq. (A8), it can be verified by inspection that $p^\tau_a(t)$ evolves from the initial probability $p^\tau_a(0) = p_a$ to the final probability $p^\tau_a(\tau) = p'_a$. In combination with Eqs. (A6) and (A1), this means that $p^\tau_b(t)$ evolves from the initial probability $p^\tau_b(0) = p_b$ to the final probability $p^\tau_b(\tau) = p_b + (p_a - p'_a) = p'_b$. Meanwhile, the probability of all states $i \notin \{a, b\}$ stays constant.

Next, we prove that condition II is satisfied by showing that

$$\lim_{\tau\to\infty} \Sigma(p, \Gamma_\tau) = 0, \tag{A9}$$

which means that EP can be made arbitrarily small by choosing a sufficiently large τ. First, plug Eqs. (A2) and (A8) into Eq. (A7) and simplify to give

$$\begin{aligned} \frac{d}{dt}p^\tau_a(t) &= W^\tau_{ab}(t)p^\tau_b(t) - W^\tau_{ba}(t)p^\tau_a(t) \\ &= \frac{1 - e^{-t}}{\tau - 1 + e^{-\tau}}(p'_a - p_a) \leqslant 0. \end{aligned} \tag{A10}$$

where the last inequality follows from $p'_a \leqslant p_a$. Next, given Eq. (7), the EP rate incurred by protocol Γ_τ at time t is

$$\begin{aligned} \dot\Sigma(p^\tau(t), W^\tau(t)) = \big[W^\tau_{ba}(t)p^\tau_a(t) - W^\tau_{ab}(t)p^\tau_b(t)\big] \\ \times \ln \frac{W^\tau_{ba}(t)p^\tau_a(t)}{W^\tau_{ab}(t)p^\tau_b(t)}. \end{aligned} \tag{A11}$$

Since $W^\tau_{ba}(t)p^\tau_a(t) \geqslant W^\tau_{ab}(t)p^\tau_b(t)$ from Eq. (A10), Eq. (A11) is the product of two positive terms. We then use the inequality $\ln x \leqslant (x - 1)$ to bound the EP rate as

$$\begin{aligned} \dot\Sigma(p^\tau(t), W^\tau(t)) &\leqslant \frac{\big[W^\tau_{ba}(t)p^\tau_a(t) - W^\tau_{ab}(t)p^\tau_b(t)\big]^2}{W^\tau_{ab}(t)p^\tau_b(t)} \\ &= \frac{(p'_a - p_a)^2}{W^\tau_{ab}(t)p^\tau_b(t)}\Big(\frac{1 - e^{-t}}{\tau - 1 + e^{-\tau}}\Big)^2, \end{aligned} \tag{A12}$$

where in the second line we used Eq. (A10). Observe that

$$W^\tau_{ab}(t) = f_\tau(t)/C \geqslant p'_a/(2C)$$

from Eqs. (A3)–(A5), and that $p^\tau_b(t)$ is increasing in t, since $\frac{d}{dt}p^\tau_b(t) = -\frac{d}{dt}p^\tau_a(t) \geqslant 0$ from Eq. (A10). Thus,

$$p^\tau_b(t) \geqslant p^\tau_b(0) = p_b > 0,$$

where the last inequality is implied by Eq. (11). Plugging these bounds into Eq. (A12) gives

$$\begin{aligned} \dot\Sigma(p^\tau(t), W^\tau(t)) &\leqslant \frac{2C(p'_a - p_a)^2}{p'_a p_b}\Big(\frac{1 - e^{-t}}{\tau - 1 + e^{-\tau}}\Big)^2 \\ &\leqslant \frac{2C(p'_a - p_a)^2}{p'_a p_b}\frac{1}{(\tau - 1 + e^{-\tau})^2}. \end{aligned}$$

Using Eq. (6) and integrating, we can bound total EP as

$$\Sigma(p, \Gamma_\tau) \leqslant \frac{2C(p'_a - p_a)^2}{p'_a p_b}\frac{\tau}{(\tau - 1 + e^{-\tau})^2}. \tag{A13}$$

Since $\lim_{\tau\to\infty} \tau/(\tau - 1 + e^{-\tau})^2 = 0$, Eq. (A13) implies that $\lim_{\tau\to\infty} \Sigma(p, \Gamma_\tau) \leqslant 0$. Since $\Sigma(p, \Gamma_\tau) \geqslant 0$ by the nonnegativity of the EP rate, this proves condition II.

Finally, we show that condition III is satisfied. In particular, we show that at all $t \in [0, \tau]$, W^τ satisfies LDB, Eq. (3), for

some energy function $E^\tau(t) \in \mathcal{E}$ [and some arbitrary symmetric activity function $\psi(t)$]. Clearly, W^τ satisfies LDB for all $i, j \notin \{a, b\}$, which follows by taking $\psi_{ij}(t) = 0$. It remains to show that the following conditions hold at all times for some $E^\tau(t) \in \mathcal{E}$:

$$\begin{aligned} W^\tau_{ab}(t) = f_\tau(t)/C &= \psi_{ab}(t) e^{\beta[E^\tau_a(t) - E^\tau_b(t)]/2}, \\ W^\tau_{ba}(t) = 1 - f_\tau(t)/C &= \psi_{ba}(t) e^{\beta[E^\tau_b(t) - E^\tau_a(t)]/2} \\ &= \psi_{ab}(t) e^{\beta[E^\tau_b(t) - E^\tau_a(t)]/2}, \end{aligned}$$

where we used Eq. (A3) as well as the symmetry $\psi_{ab}(t) = \psi_{ba}(t)$. Since $\psi(t)$ can be chosen arbitrarily, the above conditions can be restated in terms of the following single equality:

$$\beta^{-1} \ln \frac{f_\tau(t)}{C - f_\tau(t)} = E^\tau_b(t) - E^\tau_a(t). \tag{A14}$$

We now demonstrate that it is possible to choose a sufficiently large τ such that Eq. (A14) holds at all times $t \in [0, \tau]$.

Recall that $f_\tau(t)$ is monotonically decreasing in t, so

$$\begin{aligned} \beta^{-1} \ln \frac{f_\tau(t)}{C - f_\tau(t)} &\leqslant \beta^{-1} \ln \frac{f_\tau(0)}{C - f_\tau(0)} \\ &= \beta^{-1} \ln \frac{p_a}{p_b} < g_{ab}(\mathcal{E}), \end{aligned} \tag{A15}$$

where the last inequality follows from Eq. (11). Similarly,

$$\beta^{-1} \ln \frac{f_\tau(t)}{C - f_\tau(t)} \geqslant \beta^{-1} \ln \frac{f_\tau(\tau)}{C - f_\tau(\tau)}. \tag{A16}$$

Since $\lim_{\tau\to\infty} f_\tau(t) = p_a$, it follows that

$$\lim_{\tau\to\infty} \beta^{-1} \ln \frac{f_\tau(\tau)}{C - f_\tau(\tau)} = \beta^{-1} \ln \frac{p'_a}{p'_b} > -g_{ab}(\mathcal{E}), \tag{A17}$$

where the last inequality follows from Eq. (11). Combining Eqs. (A15)–(A17) implies that for some sufficiently large τ,

$$-g_{ab}(\mathcal{E}) < \beta^{-1} \ln \frac{f_\tau(t)}{C - f_\tau(t)} < g_{ab}(\mathcal{E}) \tag{A18}$$

for all $t \in [0, \tau]$. Recall from the definition of $g_{ab}(\mathcal{E})$ in Eq. (8) that there must be some $E^{(0)}, E^{(1)} \in \mathcal{E}$ such that

$$g_{ab}(\mathcal{E}) \leqslant E^{(0)}_b - E^{(0)}_a, \quad g_{ab}(\mathcal{E}) \leqslant E^{(1)}_a - E^{(1)}_b. \tag{A19}$$

Combined with Eq. (A18), this means that for all $t \in [0, \tau]$,

$$E^{(1)}_b - E^{(1)}_a < \beta^{-1} \ln \frac{f_\tau(t)}{C - f_\tau(t)} < E^{(0)}_b - E^{(0)}_a. \tag{A20}$$

Finally, since $\mathcal{E}$ is a path-connected set, there is a continuous curve of energy functions that connects $E^{(0)}$ and $E^{(1)}$. Given Eq. (A20) and the intermediate value theorem, there is some $E^\tau(t) \in \mathcal{E}$ such that Eq. (A14) is satisfied for every $t \in [0, \tau]$.

APPENDIX B: GENERALIZATION OF PROPOSITION 1 AND THEOREM 1 TO NONCONSERVATIVE FORCES

In this Appendix, we discuss how Proposition 1 and Theorem 1 can be generalized to the case when there are constraints both on the energy functions and the nonconservative forces that can be applied to a system.

In the presence of nonconservative forces, transitions rates can be parametrized via a generalized version of Eq. (3),

$$W_{ji} = \psi_{ji} e^{\beta[E_i - E_j + N_{ji}]/2}, \tag{B1}$$

where $N_{ji} = -N_{ij}$ is an antisymmetric function that reflects a nonconservative force that biases transitions from state i to state j. (As above, E is the energy function and $\psi_{ji} = \psi_{ij} \geqslant 0$ is a symmetric positive function that controls the overall timescale of transitions between states i and j.) Now suppose there are constraints on *both* the energy functions and the nonconservative forces that can be applied to a system. For notational convenience, we define an antisymmetric matrix F such that $F_{ji} = -F_{ij}$ is the drop of the potential energy plus the nonconservative bias in going from i to j, $F_{ji} := E_i - E_j + N_{ji}$. Equation (B1) can then be written as

$$W_{ji} = \psi_{ji} e^{\beta[F_{ji}]/2}. \tag{B2}$$

The presence of constraints on the energy functions and nonconservative driving forces can be stated formally as follows: there is some set $\mathcal{F} \subset \mathbb{R}^{n\times n}$ of antisymmetric matrices such that for any possible driving protocol $\Gamma = \{(E(t), W(t)) : t \in [0, \tau]\}$, the rate matrix $W(t)$ at all $t \in [0, \tau]$ satisfies Eq. (B2) for *some* $F \in \mathcal{F}$.

We now consider how Proposition 1, which constructs a protocol that we call a *transfer*, can be generalized to this more general type of constraint. To begin, define the controllable energy gap in terms of $\mathcal{F}$ [rather than in terms of $\mathcal{E}$, as in Eq. (8)] in the following manner:

$$g'_{ij}(\mathcal{F}) := \min\Big\{ \max_{F\in\mathcal{F}} F_{ij}, \max_{F\in\mathcal{F}} F_{ji} \Big\}. \tag{B3}$$

Then, in the presence of constraints on both the energy functions and the nonconservative driving forces, Proposition 1 continues to hold under this new definition. In particular, consider any initial distribution p and final distribution p' which involve transferring Δ probability from state a to state b, as in Eq. (9). Assume that Eq. (11) holds for this pair of distributions, where $g_{ab}(\mathcal{E})$ is replaced by $g'_{ab}(\mathcal{F})$ as defined in Eq. (B3). Then, for any $\epsilon > 0$, there is a protocol Γ with a time-dependent rate matrix $W(t)$ over $t \in [0, \tau]$ such that (I) p is transformed to p', (II) $\Sigma(p, \Gamma) \leqslant \epsilon$, and (III) $W(t)$ satisfies Eq. (B2) for some $F \in \mathcal{F}$ at all t. The proof of statements I and II is exactly the same as appears in Appendix A. The proof of statement III is also the same as appears in Appendix A, as long as the following replacements are made.

(1) $g_{ij}(\mathcal{E})$ should be replaced by $g'_{ij}(\mathcal{F})$.

(2) Statements like "$E(t) \in \mathcal{E}$" should be replaced by statements like "$W(t)$ obeys Eq. (B2) for some $F \in \mathcal{F}$."

(3) Expressions like $E_b - E_a$ in Eqs. (A14), (A19), and (A20) should be replaced by corresponding versions with F_{ba}.

Finally, note that Theorem 2 is proved (informally in the main text) by an explicit construction that shows how the initial distribution p can be transformed to a final distribution p' via an appropriate sequence of transfers. The same construction also works for transfers defined under constraints on both energy functions and nonconservative driving forces (as discussed above), as long as the capacity term C which appears in Eqs. (13), (14), and (16) is defined in terms of controllable energy gaps $g'_{ij}(\mathcal{F})$ from Eq. (B3), rather than $g_{ij}(\mathcal{E})$.

APPENDIX C: DERIVATION OF EQ. (17)

Suppose that $\mathcal{E}$ contains a one-dimensional curve of energy functions that connects $E^{\varnothing}$ and $-E^{\varnothing}$. Consider some pair of states i, j, and first assume that $E_i^{\varnothing} \geqslant E_j^{\varnothing}$. Then, using the definition of $g_{ij}(\mathcal{E})$ in Eq. (8),

$$g_{ij}(\mathcal{E}) \geqslant \min\left\{E_i^{\varnothing} - E_j^{\varnothing}, \left(-E_j^{\varnothing}\right) - \left(-E_i^{\varnothing}\right)\right\} = E_i^{\varnothing} - E_j^{\varnothing}.$$

Conversely, if $E_i^{\varnothing} \leqslant E_j^{\varnothing}$, then $g_{ij}(\mathcal{E}) \geqslant E_j^{\varnothing} - E_i^{\varnothing}$. Combining these results implies

$$g_{ij}(\mathcal{E}) \geqslant \left|E_i^{\varnothing} - E_j^{\varnothing}\right|. \tag{C1}$$

Next, for any state i, let $m(i)$ indicate the state with the largest energy difference from i under $E^{\varnothing}$:

$$m(i) \in \arg\max_k \left|E_i^{\varnothing} - E_k^{\varnothing}\right|. \tag{C2}$$

Given a pair of states i, j, we consider two possibilities. Under the first possibility, $E_{m(i)}^{\varnothing} = E_{m(j)}^{\varnothing}$ and the path $\vec{x} = (i \to m(i) \to j)$ obeys

$$\begin{aligned}\mathsf{c}(\vec{x}) &= \min\{g_{im(i)}(\mathcal{E}), g_{m(i)j}(\mathcal{E})\} \\ &\geqslant \min\left\{\left|E_i^{\varnothing} - E_{m(i)}^{\varnothing}\right|, \left|E_{m(i)}^{\varnothing} - E_j^{\varnothing}\right|\right\} \\ &= \min\left\{\left|E_i^{\varnothing} - E_{m(i)}^{\varnothing}\right|, \left|E_{m(j)}^{\varnothing} - E_j^{\varnothing}\right|\right\},\end{aligned} \tag{C3}$$

where in the second line we used Eq. (C1). Under the second possibility, $E_{m(i)}^{\varnothing} \neq E_{m(j)}^{\varnothing}$ and capacity of the path $\vec{x} = (i \to m(i) \to m(j) \to j)$ obeys

$$\begin{aligned}\mathsf{c}(\vec{x}) &= \min\{g_{im(i)}(\mathcal{E}), g_{m(i)m(j)}(\mathcal{E}), g_{m(j)j}(\mathcal{E})\} \\ &\geqslant \min\left\{\left|E_i^{\varnothing} - E_{m(i)}^{\varnothing}\right|, \left|E_{m(i)}^{\varnothing} - E_{m(j)}^{\varnothing}\right|, \left|E_{m(j)}^{\varnothing} - E_j^{\varnothing}\right|\right\}.\end{aligned} \tag{C4}$$

Observe that for any i, the state $m(i)$ defined in Eq. (C2) either obeys $m(i) \in \arg\min_j E_j^{\varnothing}$ or $m(i) \in \arg\max_j E_j^{\varnothing}$. Since we assumed that $E_{m(i)}^{\varnothing} \neq E_{m(j)}^{\varnothing}$, it must be that the set $\{m(i), m(j)\}$ includes both lowest and highest energy states. This in turn implies that $|E_{m(i)}^{\varnothing} - E_{m(j)}^{\varnothing}|$ is larger than both $|E_{m(i)}^{\varnothing} - E_i^{\varnothing}|$ and $|E_{m(j)}^{\varnothing} - E_j^{\varnothing}|$, allowing us to rewrite Eq. (C4) as

$$\mathsf{c}(\vec{x}) \geqslant \min\left\{\left|E_i^{\varnothing} - E_{m(i)}^{\varnothing}\right|, \left|E_{m(j)}^{\varnothing} - E_j^{\varnothing}\right|\right\}, \tag{C5}$$

which is the same as Eq. (C3).

To summarize, above we showed that for any pair of states i, j, there exists a path $\vec{x}$ from i to j such that

$$\begin{aligned}\mathsf{c}(\vec{x}) &\geqslant \min\left\{\left|E_i^{\varnothing} - E_{m(i)}^{\varnothing}\right|, \left|E_{m(j)}^{\varnothing} - E_j^{\varnothing}\right|\right\} \\ &= \min\left\{\max_k \left|E_i^{\varnothing} - E_k^{\varnothing}\right|, \max_k \left|E_j^{\varnothing} - E_k^{\varnothing}\right|\right\} \\ &\geqslant \min_l \max_k \left|E_l^{\varnothing} - E_k^{\varnothing}\right|,\end{aligned} \tag{C6}$$

where we used Eqs. (C3) and (C5) and the definition in Eq. (C2). The first inequality in Eq. (17) then follows immediately from Eq. (C6) and the definition of $\mathsf{C}(\mathcal{E})$ in Eq. (13):

$$\mathsf{C}(\mathcal{E}) = \min_{i \neq j} \max_{\vec{x} \in \mathcal{P}(i,j)} \mathsf{c}(\vec{x}) \tag{C7}$$

$$\geqslant \min_i \max_j \left|E_i^{\varnothing} - E_j^{\varnothing}\right|. \tag{C8}$$

To derive the second inequality in Eq. (17), we consider two cases. First, for any i such that $E_i^{\varnothing} \leqslant \frac{\max_j E_j^{\varnothing} + \min_j E_j^{\varnothing}}{2}$,

$$\max_j \left|E_i^{\varnothing} - E_j^{\varnothing}\right| = \max_j E_j^{\varnothing} - E_i^{\varnothing} \geqslant \frac{\max_j E_j^{\varnothing} - \min_j E_j^{\varnothing}}{2}.$$

Second, for any i such that $E_i^{\varnothing} > \frac{\max_j E_j^{\varnothing} + \min_j E_j^{\varnothing}}{2}$,

$$\max_j \left|E_i^{\varnothing} - E_j^{\varnothing}\right| = E_i^{\varnothing} - \min_j E_j^{\varnothing} \geqslant \frac{\max_j E_j^{\varnothing} - \min_j E_j^{\varnothing}}{2}.$$

We finish by showing that equality is achieved in Eq. (C8) if the set of energy functions is a one-dimensional line segment:

$$\mathcal{E} = \{\lambda E^{\varnothing} : \lambda \in [-1, 1]\}.$$

In that case, it is easy to verify from the definition of $g_{ij}(\mathcal{E})$ that $g_{ij}(\mathcal{E}) = |E_i^{\varnothing} - E_j^{\varnothing}|$, so equality is achieved in Eq. (C1). Next, note that for any path $\vec{x}$ that starts or ends on state i, it must be that

$$\mathsf{c}(\vec{x}) \leqslant \max_j g_{ij}(\mathcal{E}) = \max_j \left|E_i^{\varnothing} - E_j^{\varnothing}\right|.$$

Plugging into the definition of $\mathsf{C}(\mathcal{E})$ [see Eq. (C7)] and simplifying gives

$$\mathsf{C}(\mathcal{E}) \leqslant \min_i \max_j \left|E_i^{\varnothing} - E_j^{\varnothing}\right|,$$

which implies equality in Eq. (C8).

APPENDIX D: DERIVATION OF EQ. (22)

Assume that Eq. (21) holds for all p and $W \in \mathcal{W}$. Then, it must hold for the distribution and rate matrix at all times $t \in [0, \tau]$, $p(t)$ and $W(t)$. The right side of Eq. (21) can be written in terms of the following time derivative:

$$\begin{aligned}&-\eta \sum_{i,j} p_i(t) W_{ji}(t) \ln \frac{p_j(t)}{\{\Pi(p(t))\}_j} \\ &\quad = -\eta \frac{d}{dt} D[p(t)\|q]\big|_{q=\Pi(p(t))}.\end{aligned} \tag{D1}$$

The total derivative rule from calculus then gives

$$\begin{aligned}-\frac{d}{dt} D[p(t)\|q]|_{q=\Pi(p(t))} &= -\frac{d}{dt} D[p(t)\|\Pi(p(t))] \\ &\quad + \frac{d}{dt} D[q\|\Pi(p(t))]|_{q=p(t)}.\end{aligned} \tag{D2}$$

The second term above can be bounded as

$$\begin{aligned}&\frac{d}{dt} D[q\|\Pi(p(t))]|_{q=p(t)} \\ &\quad = \lim_{s \to 0} \frac{1}{s} \left(D[p(t)\|\Pi(e^{sW(t)} p(t))] - D[p(t)\|\Pi(p(t))]\right) \geqslant 0,\end{aligned}$$

where the inequality follows from the definition of the KL projection in Eq. (20). Combining with Eq. (D2) gives

$$-\frac{d}{dt} D[p(t)\|q]|_{q=\Pi(p(t))} \geqslant -\frac{d}{dt} D[p(t)\|\Pi(p(t))],$$

which can be combined with Eqs. (21) and (D1) to give Eq. (22).

APPENDIX E: FINDING η IN THEOREM 2 WITH CONVEX OPTIMIZATION

In this Appendix, we show how to identify the optimal value of η for the EP bound in Theorem 2 using convex optimization techniques. We use $\mathcal{P}$ to indicate the set of all probability distributions over the state space of the system. In addition, we use the general definition of the KL projection to some set of distributions Ω:

$$\Pi(p) := \arg\min_{q\in\Omega} D(p\|q). \tag{E1}$$

Here we assume that Ω is a finite set.

For notational convenience, define the following function:

$$f(p,q,W,\eta) = (1-\eta)\left[-\sum_{i,j} p_i W_{ji} \ln p_j\right] + \sum_{i,j} p_i W_{ji}\left[\ln\frac{W_{ji}}{W_{ij}} - \eta \ln q_j\right], \tag{E2}$$

where $p \in \mathcal{P}$ and $q \in \mathcal{P}$ are probability distributions, W is a rate matrix, and $\eta \in [0,1]$ is some scalar. Note that the first bracketed term on the right hand side of Eq. (E2) is the rate of increase of the Shannon entropy of distribution p under rate matrix W. This term is convex in p, since it differs from the EP rate by a linear function of p [see Eq. (E3) below], and the EP rate is convex in p [49]. Thus, for a fixed W, E, and η, f is a weighted sum of a convex function of p and a linear function of p, meaning that it is a convex function of p.

For any distribution p and rate matrix W, rewrite the left hand side of Eq. (21) (the EP rate incurred by distribution p under rate matrix W) as

$$-\sum_{i,j} p_i W_{ji} \ln p_j + \sum_{i,j} p_i W_{ji} \ln\frac{W_{ji}}{W_{ij}}, \tag{E3}$$

where we used

$$\sum_{i,j} p_i W_{ji} \ln p_i = \sum_i p_i \left(\sum_j W_{ji}\right) \ln p_i = 0.$$

Next, rewrite the right hand side of Eq. (21) as

$$-\eta \sum_{i,j} p_i W_{ji} \ln p_j + \eta \sum_{i,j} p_i W_{ji} \ln[\Pi(p)]_j. \tag{E4}$$

We can now use Eqs. (E2)–(E4) to restate Eq. (21) in the following way: Eq. (21) holds for a given $\eta \in [0,1]$ if for all distributions p, $W \in \mathcal{W}$, and $q \in \Omega$ such that $\Pi(p) = q$,

$$f(p,q,W,\eta) \geqslant 0. \tag{E5}$$

Observe that $\Pi(p) = q$ is equivalent to the condition

$$D(p\|q) \leqslant D(p\|q') \quad \forall q' \in \Omega, \tag{E6}$$

which, with some simple rearranging, is equivalent to the following set of linear constraints on p:

$$\forall q' \in \Omega : \sum_i p_i \ln\frac{q_i}{q'_i} \geqslant 0.$$

To summarize, Eq. (21) holds for a given η when the following inequality is satisfied for each $W \in \mathcal{W}$ and $q \in \Omega$:

$$0 \leqslant \min_{p\in\mathcal{P}} f(p,q,W,\eta) \quad \text{s.t. } \forall q' \in \Omega : \sum_i p_i \ln\frac{q_i}{q'_i} \geqslant 0. \tag{E7}$$

For each W and q, Eq. (E7) can be verified for a given η by solving a convex minimization problem subject to linear constraints, which can be done efficiently using standard numerical techniques. For a finite set of rate matrices $\mathcal{W}$ and distributions Ω, one can solve $|\mathcal{W}| \times |\Omega|$ such problems to verify whether Eq. (21) holds for a given η. Finally, to find the largest such $\eta \in [0,1]$ [thereby making Eq. (21) as tight as possible], one can use the bisection method on the interval $\eta \in [0,1]$. Code for doing this optimization is available in [45].

APPENDIX F: GENERALIZATION OF THEOREM 2 TO MULTIPLE RESERVOIRS AND/OR NONCONSERVATIVE FORCES

Our derivation Theorem 2 makes no assumptions about the structure of the rate matrices in $\mathcal{W}$, including whether or not local detailed balance is obeyed. For this reason, the results derived in Sec. V can be generalized to consider systems coupled to multiple reservoirs and/or subject to nonconservative forces. In this more general situation, $\Pi(\cdot)$ should be defined as the projection to the set of nonequilibrium stationary distributions of the rate matrices in $\mathcal{W}$ (or alternatively, as mentioned in Sec. V, to any other set of distributions). Theorem 2 then applies for any $\eta \in [0,1]$ which satisfies Eq. (21) for this $\Pi(\cdot)$.

In addition, the bound in Theorem 2 can be strengthened for systems coupled to multiple reservoirs, as long as each rate matrix in $W \in \mathcal{W}$ can be decomposed into separate contributions from each reservoir r as $W = \sum_r W^{(r)}$ [35]. In that case, the EP rate [which appears on the left hand side of Eq. (21)] can defined as [35]

$$\dot{\Sigma}(p,W) = \frac{1}{2}\sum_{i,j,r} \left(p_i W^{(r)}_{ji} - p_j W^{(r)}_{ij}\right) \ln\frac{p_i W^{(r)}_{ji}}{p_j W^{(r)}_{ij}}. \tag{F1}$$

[Compare to Eq. (7) in the main text, which applies in the presence of a single reservoir.] The expression in Eq. (F1) is an upper bound on the expression in Eq. (7) [35], so the largest η which satisfies the implicit inequality Eq. (21) for $\dot{\Sigma}$ as defined in Eq. (F1) will be no smaller than (and possibly larger than) the largest η which satisfies that inequality for $\dot{\Sigma}$ as defined in Eq. (7).

[1] U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines, Rep. Prog. Phys. **75**, 126001 (2012).

[2] M. Esposito and C. Van den Broeck, Second law and Landauer principle far from equilibrium, Europhysics Letters **95**, 40004 (2011).

[3] 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).

[4] J. M. Parrondo, J. M. Horowitz, and T. Sagawa, Thermodynamics of information, Nat. Phys. **11**, 131 (2015).

[5] C. Van den Broeck, Stochastic thermodynamics: A brief introduction, Phys. Complex Colloids **184**, 155 (2013).

[6] C. Maes, *Non-Dissipative Effects in Nonequilibrium Systems* (Springer, New York, 2018).

[7] It is known that for a free relaxation from an initial distribution p to a final equilibrium distribution $p' = \pi^E$, the EP obeys $\Sigma = D(p\|\pi^E)$ [48]. In addition, we can rewrite $D(p\|\pi^E) = I_p(X;Y) + D(p_X p_Y \|\pi_X^E \pi_Y^E) \geqslant I_p(X;Y)$. Combining implies $\Sigma \geqslant I_p(X;Y)$.

[8] D. H. Wolpert, Uncertainty Relations and Fluctuation Theorems for Bayes Nets, Phys. Rev. Lett. **125**, 200602 (2020).

[9] D. Wolpert, Minimal entropy production rate of interacting systems, New J. Phys. **22**, 113013 (2020).

[10] D. H. Wolpert, The stochastic thermodynamics of computation, J. Phys. A: Math. Theor. **52**, 193001 (2019).

[11] D. H. Wolpert and A. Kolchinsky, Thermodynamics of computing with circuits, New J. Phys. **22**, 063047 (2020).

[12] A. B. Boyd, D. Mandal, and J. P. Crutchfield, Thermodynamics of Modularity: Structural Costs Beyond the Landauer Bound, Phys. Rev. X **8**, 031036 (2018).

[13] A. Kolchinsky and D. H. Wolpert, Work, Entropy Production, and Thermodynamics of Information under Protocol Constraints, Phys. Rev. X arXiv:2008.10764 [Phys. Rev. X (to be published)].

[14] H. Wilming, R. Gallego, and J. Eisert, Second law of thermodynamics under control restrictions, Phys. Rev. E **93**, 042126 (2016).

[15] M. Lostaglio, Á. M. Alhambra, and C. Perry, Elementary thermal operations, Quantum **2**, 52 (2018).

[16] E. Bäumer, M. Perarnau-Llobet, P. Kammerlander, H. Wilming, and R. Renner, Imperfect thermalizations allow for optimal thermodynamic processes, Quantum **3**, 153 (2019).

[17] C. Perry, P. Ćwikliński, J. Anders, M. Horodecki, and J. Oppenheim, A Sufficient set of Experimentally Implementable Thermal Operations for Small Systems, Phys. Rev. X **8**, 041049 (2018).

[18] J. A. Owen, A. Kolchinsky, and D. H. Wolpert, Number of hidden states needed to physically implement a given conditional distribution, New J. Phys. **21**, 013022 (2019).

[19] D. H. Wolpert, A. Kolchinsky, and J. A. Owen, A space–time tradeoff for implementing a function with master equation dynamics, Nat. Commun. **10**, 1727 (2019).

[20] K. Korzekwa and M. Lostaglio, Quantum Advantage in Simulating Stochastic Processes, Phys. Rev. X **11**, 021019 (2021).

[21] M. Esposito, R. Kawai, K. Lindenberg, and C. Van den Broeck, Finite-time thermodynamics for a single-level quantum dot, Europhysics Letters **89**, 20003 (2010).

[22] D. A. Sivak and G. E. Crooks, Thermodynamic Metrics and Optimal Paths, Phys. Rev. Lett. **108**, 190602 (2012).

[23] N. Shiraishi, K. Funo, and K. Saito, Speed Limit for Classical Stochastic Processes, Phys. Rev. Lett. **121**, 070601 (2018).

[24] A. Gomez-Marin, T. Schmiedl, and U. Seifert, Optimal protocols for minimal work processes in underdamped stochastic thermodynamics, J. Chem. Phys. **129**, 024114 (2008).

[25] H. Then and A. Engel, Computing the optimal protocol for finite-time processes in stochastic thermodynamics, Phys. Rev. E **77**, 041105 (2008).

[26] P. R. Zulkowski and M. R. DeWeese, Optimal finite-time erasure of a classical bit, Phys. Rev. E **89**, 052140 (2014).

[27] T. Schmiedl and U. Seifert, Optimal Finite-Time Processes in Stochastic Thermodynamics, Phys. Rev. Lett. **98**, 108301 (2007).

[28] B. Remlein and U. Seifert, Optimality of nonconservative driving for finite-time processes with discrete states, Phys. Rev. E **103**, L050105 (2021).

[29] E. Aurell, C. Mejía-Monasterio, and P. Muratore-Ginanneschi, Optimal Protocols and Optimal Transport in Stochastic Thermodynamics, Phys. Rev. Lett. **106**, 250601 (2011).

[30] E. Aurell, K. Gawędzki, C. Mejía-Monasterio, R. Mohayaee, and P. Muratore-Ginanneschi, Refined second law of thermodynamics for fast random processes, J. Stat. Phys. **147**, 487 (2012).

[31] B. B. Machta, Dissipation Bound for Thermodynamic Control, Phys. Rev. Lett. **115**, 260603 (2015).

[32] G. Verley, C. V. d. Broeck, and M. Esposito, Work statistics in stochastically driven systems, New J. Phys. **16**, 095001 (2014).

[33] G. Gour, M. P. Müller, V. Narasimhachar, R. W. Spekkens, and N. Y. Halpern, The resource theory of informational nonequilibrium in thermodynamics, Phys. Rep. **583**, 1 (2015).

[34] J. Schnakenberg, Network theory of microscopic and macroscopic behavior of master equation systems, Rev. Mod. Phys. **48**, 571 (1976).

[35] M. Esposito and C. Van den Broeck, Three faces of the second law. I. Master equation formulation, Phys. Rev. E **82**, 011143 (2010).

[36] There are at most $n-1$ states that have either $p_i < p_i'$ or $p_i > p_i'$ and are not state 1. Furthermore, each path can have at most $\ell = n-1$ steps. Thus, the number of transfers involved in the two stage construction is no greater than $(n-1)^2$.

[37] M. Pollack, Letter to the editor: The maximum capacity through a network, Oper. Res. **8**, 733 (1960).

[38] T. Hu, Letter to the editor: The maximum capacity route problem, Oper. Res. **9**, 898 (1961).

[39] P. Talkner and J. Łuczka, Rate description of Fokker-Planck processes with time-dependent parameters, Phys. Rev. E **69**, 046109 (2004).

[40] N. G. Van Kampen, *Stochastic Processes in Physics and Chemistry* (Elsevier, Amsterdam, 1992), Vol. 1.

[41] P. Hänggi, P. Talkner, and M. Borkovec, Reaction-rate theory: Fifty years after Kramers, Rev. Mod. Phys. **62**, 251 (1990).

[42] R. D. Astumian, Adiabatic operation of a molecular machine, Proc. Natl. Acad. Sci. USA **104**, 19715 (2007).

[43] S. Rahav, J. Horowitz, and C. Jarzynski, Directed Flow in Nonadiabatic Stochastic Pumps, Phys. Rev. Lett. **101**, 140602 (2008).

[44] I. Csiszár and F. Matus, Information projections revisited, IEEE Trans. Inf. Theory **49**, 1474 (2003).

[45] See, https://github.com/artemyk/epbound.

[46] A. Kolchinsky and D. H. Wolpert, Dependence of dissipation on the initial distribution over states, J. Stat. Mech.: Theory Exp. (2017) 083202.

[47] H. Spohn, Entropy production for quantum dynamical semigroups, J. Math. Phys. **19**, 1227 (1978).

[48] F. Schlögl, On the statistical foundation of the thermodynamic evolution criterion of Glansdorff and Prigogine, Ann. Phys. (NY) **45**, 155 (1967).

[49] C. Maes and K. Netocny, Minimum entropy production principle, Scholarpedia **8**, 9664 (2013).