Generalized free energy and excess/housekeeping decomposition in nonequilibrium systems: From large deviations to thermodynamic speed limits

Artemy Kolchinsky [1,2,*] Andreas Dechant [3] Kohei Yoshimura [4] and Sosuke Ito [2,4]

[1] *ICREA-Complex Systems Lab, Universitat Pompeu Fabra, 08003 Barcelona, Spain*

[2] *Universal Biology Institute, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan*

[3] *Department of Physics, No. 1 Graduate School of Science, Kyoto University, Kyoto 606-8502, Japan*

[4] *Department of Physics, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan*

(Received 11 December 2024; accepted 23 February 2026; published 8 April 2026)

In genuine nonequilibrium systems under continuous driving, thermodynamic forces are nonconservative and cannot be described by any free energy potential. Nonetheless, we show that such systems can be associated with a *generalized free energy* derived from a large-deviation variational principle. This variational principle yields a decomposition of fluxes, forces, and entropy production into a conservative *excess* part and a nonconservative *housekeeping* part, exemplifying an information-geometric Pythagorean theorem. The decomposition is broadly applicable—including to stochastic master equations as well as closed and open deterministic chemical reaction networks—and accessible to thermodynamic inference from short-time trajectory data. We also show that the excess entropy production obeys a thermodynamic speed limit bounding the rate of state evolution and external fluxes. We illustrate the framework on driven Markov jump processes, nonlinear chemical oscillators, and real-world metabolic networks, where we obtain tight dissipation bounds and identify futile metabolic cycles. Connections are drawn to large deviations, Onsager theory, and previous excess/housekeeping decompositions.

DOI: 10.1103/r48t-dghl

I. INTRODUCTION

Thermodynamics is one of the most far reaching and useful theoretical frameworks in the physical sciences. Traditionally, this framework was mainly used to study "passive" systems, which tend toward equilibrium when allowed to relax freely. When a passive system is brought out of equilibrium, many of its thermodynamic, dynamical, and statistical properties are governed by the *nonequilibrium free energy* $\mathcal{F}$. As we discuss below, the nonequilibrium free energy controls relaxation dynamics [1,2], equilibrium fluctuations [3,4], and extractable work [5,6]. In addition, the thermodynamic forces that describe a passive system are "conservative," being determined by the gradient of the nonequilibrium free energy. For this reason, we refer to passive systems as *conservative systems*.

In nonequilibrium thermodynamics, *entropy production rate* (EPR) refers to the increase of the entropy of a system and its environment, and it serves as the fundamental measure of thermodynamic dissipation. In conservative systems, EPR is proportional to the loss of the nonequilibrium free energy over time, vanishing at equilibrium, where the state remains constant. Recently, the connection between entropy production and nonstationarity has been used to derive *thermodynamic speed limits* (TSLs) [7–11]. TSLs bound the time and dissipation needed to transform a system from one state to another, and they quantify the thermodynamic efficiency of finite-time transformations [12]. Moreover, the minimal dissipation required to go between two states in a finite time defines a thermodynamic distance between states. In the near-equilibrium regime, this distance is called "thermodynamic length" [13–16]. Far from equilibrium, it is the distance based on thermodynamic optimal transport [17–24].

Nowadays, there is growing interest in "genuine nonequilibrium" systems that undergo continuous driving. Such driving can arise from nonequilibrium boundary conditions, applied mechanical forces, or internal fuel sources (as in active matter [25]). When allowed to relax freely, such systems tend toward nonequilibrium steady states, or possibly limit cycles and chaos in the case of nonlinear dynamics. Many examples are found in molecular biology, where driving is provided by nonequilibrium concentrations of adenosine triphosphate (ATP), adenosine diphosphate (ADP), and phosphate (Pi).

We refer to continuously driven systems as *nonconservative*, because they exhibit nonconservative thermodynamic forces that cannot be expressed as the gradient of any free energy potential. Nonconservative systems exhibit cyclic fluxes that contribute to dissipation but do not influence state evolution. For example, a nonconservative system may remain near a steady state while parameters slowly change, incurring arbitrarily large entropy production despite negligible state change. In nonconservative systems, vanishing speed of evolution does not imply small EPR. Therefore, TSLs that relate EPR and speed may become arbitrarily loose in non-

*Contact author: artemyk@gmail.com

TABLE I. Summary of our results and comparison with the steady-state/Hatano-Sasa approach (discussed in more detail in Secs. III B and IX B). Top row: $\boldsymbol{j}$ indicates the vector of one-way fluxes, ∇ the stoichiometric matrix, and $\boldsymbol{x}$ the concentration vector or probability distribution.

	Our generalized potential $\boldsymbol{\phi}^*$	Section	Steady-state potential $\boldsymbol{\phi}^{\rm ss}$
Definition	$\nabla^\top \boldsymbol{j} = -\nabla^\top(\boldsymbol{j} \circ e^{\nabla \boldsymbol{\phi}^*})$	Secs. IV and V	$\boldsymbol{\phi}^{\rm ss} = \ln \frac{\boldsymbol{x}}{\boldsymbol{x}^{\rm ss}}$
Excess/housekeeping decomposition	Information-geometric decomposition	Sec. V E	Hatano-Sasa, also called adiabatic/nonadiabatic
Excess part	Conservative		Nonstationary
Housekeeping part	Nonconservative		Stationary
Linear-response coefficients	Short-time diffusion coefficients	Sec. V G	Steady-state diffusion coefficients
Large deviations	Dynamical fluctuations	Sec. VI	Steady-state fluctuations
Thermodynamic speed limit	Wasserstein speed	Sec. VII	Total variation speed (MJPs only)

conservative systems, and they typically do not provide useful measures of efficiency for finite-time transformations.

In this paper, we consider the problem of defining a "generalized" free energy for nonconservative systems. We also consider the related problem of decomposing nonconservative systems into *excess* and *housekeeping* parts. The excess part is the conservative contribution associated with the generalized free energy, and it vanishes once the potential becomes "equilibrated" and the internal dynamics become stationary. The excess part should discount the contribution of cyclic fluxes to dissipation, thus giving rise to meaningful TSLs in nonconservative systems. The housekeeping part is the nonconservative contribution that cannot be associated with the generalized free energy, and it does not vanish even when the internal dynamics are stationary. As we show, this decomposition may be considered at the level of thermodynamic forces (conservative/nonconservative forces), fluxes (gradient/cyclic fluxes), or dissipation (excess/housekeeping EPR).

There has been extensive work on generalizations of the free energy [26–33] and excess/housekeeping decompositions [22,34–42] in nonconservative systems. The best-known existing approach is based on nonequilibrium steady states [33,35,43,44]. Here, the generalized free energy is defined as the large-deviation rate function of steady-state fluctuations, which in many cases is the relative entropy to the steady-state distribution, sometimes called the "quasipotential" [44]. In this approach, the excess EPR is defined as the decrease of the quasipotential over time, and it captures the nonstationary component of dissipation [36,38,45–47]. The housekeeping EPR is defined as the remainder, and it captures the stationary component of dissipation. In the literature, this excess/housekeeping decomposition is termed the "Hatano-Sasa" (HS) [36] or the "nonadiabatic/adiabatic" decomposition [38,45–47].

The quasipotential is very useful for studying stability and fluctuations of stationary systems. However, as we discuss below, it is not as useful for studying transient systems away from steady state, nor for deriving transient relations such as TSLs. Fundamentally, this is because the quasipotential is defined in terms of steady-state statistics, rather than "local-in-time" statistics of a transient system.

Our definition of the generalized free energy and the excess/housekeeping decomposition is based on a variational principle, and it makes no explicit reference to steady states. This principle is universally applicable to stochastic and deterministic discrete systems, including stochastic master equations as well as closed and open deterministic chemical reaction networks (CRNs) with arbitrary kinetics. As we show, this variational principle can be understood in terms of dynamical large deviations of local-in-time flux fluctuations. From this perspective, the excess EPR quantifies the statistical irreversibility of state dynamics, while the generalized free energy is the "most irreversible" observable—the least likely to evolve backward in time due to a stochastic fluctuation. The connection to large deviations means that our generalized free energy and excess EPR are directly related to fluctuation statistics, allowing for practical thermodynamic inference [48]. A brief summary of our results and comparison with the steady-state approach is provided in Table I.

In addition, we derive a tight TSL that relates our excess EPR to dynamical activity and speed of evolution. Importantly, speed is measured in terms of optimal-transport distance (1-Wasserstein) [11,19,49]. This distance is sensitive to the system's internal topology, thus capturing important dynamical constraints. In open systems, including stationary ones, our TSL bounds internal dissipation in terms of activity and the velocity of the external fluxes.

Along the way, we highlight fundamental connections to several theoretical frameworks. In particular, we relate our excess/housekeeping decomposition to the celebrated Pythagorean theorem from information geometry [50], and we relate the linear-response regime of our approach to Onsager theory and thermodynamic length [15].

Our results are illustrated on three examples. The first is a simple unicyclic master equation, where we demonstrate important differences between our approach and the steady-state HS approach. The second example is the Brusselator, a nonlinear chemical oscillator, where we derive our decomposition in closed form and demonstrate a general TSL for limit-cycle oscillators. The third example is a set of real-world metabolic networks, modeled as open CRNs in steady state. Using our TSL, we quantify the efficiency of several metabolic pathways. We also use our excess/housekeeping decomposition to identify futile metabolic cycles and quantify their associated dissipation.

The paper is laid out as follows. In the next section, we describe our physical setup and formalism. Section III A provides background on nonequilibrium free energy in conservative systems and its generalization based on nonequilibrium steady states. Section IV introduces our variational principle

for conservative systems. Section V extends this variational principle to nonconservative systems, introduces the excess/housekeeping decomposition, and discusses its behavior under coarse graining. It also considers the linear-response regime and makes connections to information geometry, Onsager theory, and thermodynamic length. Section VI relates our approach to dynamical large deviations and derives a thermodynamic uncertainty relation (TUR). Section VII derives a TSL for excess entropy production based on Wasserstein distance. Section VIII illustrates our approach on several examples, including real-world metabolic networks. Section IX compares our proposal to several previous approaches, including ones based on Euclidean-Onsager geometry [22], Hessian geometry [51,52], and the steady-state approach. We finish in Sec. X with a discussion and suggestions for future work.

Appendixes A and B discuss nonlinear master equations and systems with odd variables. Supporting derivations and additional numerical comparisons are found in the Supplemental Material [53].

II. SETUP AND PRELIMINARIES

Before proceeding, we fix notation. Vectors are written in bold, $\boldsymbol{a} = (a_1, a_2, \ldots) \in \mathbb{R}^d$, including the special vectors $\boldsymbol{0} = (0, 0, 0, \ldots)$ and $\boldsymbol{1} = (1, 1, 1, \ldots)$. We use $e^{\boldsymbol{a}}, \boldsymbol{a}^2, |\boldsymbol{a}|, \ldots$ to indicate elementwise operations, for example, $e^{\boldsymbol{a}} := (e^{a_1}, e^{a_2}, \ldots)$. The notation $\boldsymbol{a} \circ \boldsymbol{b} := (a_1 b_1, a_2 b_2, \ldots)$ indicates elementwise multiplication and $\boldsymbol{a}/\boldsymbol{b} := (a_1/b_1, a_2/b_2, \ldots)$ indicates elementwise division. As discussed below, ∇ refers to the stoichiometric matrix, and it acts like the discrete gradient operator. To avoid confusion, we write the usual gradient of a function $f : \mathbb{R}^d \to \mathbb{R}$ as $\mathrm{grad}_{\boldsymbol{x}} f = (\partial_{x_1} f, \ldots, \partial_{x_d} f)$.

A. States, reactions, and dynamics

We focus on discrete Markovian systems with general (linear or nonlinear) dynamics. As special cases, this includes Markov jump processes (MJPs) that represent the transport of probability between microstates in stochastic systems. It also includes deterministic chemical reaction networks used to model chemical systems in the large-volume limit [46,54,55]. Continuous-state systems are briefly mentioned in Sec. IX A and non-Markovian systems in Sec. X.

The system's thermodynamic *state* at time t is specified by a non-negative vector $\boldsymbol{x}(t) = (x_1(t), \ldots, x_d(t)) \in \mathbb{R}^d_+$. We always use the term *state* to refer to the system's thermodynamic state $\boldsymbol{x}(t)$. When referring to the microstates, we do so explicitly. The term *state function* refers to functions $f : \mathbb{R}^d_+ \to \mathbb{R}$. The term *state observable* refers to functions over microstates/species, represented by vectors $\boldsymbol{\phi} \in \mathbb{R}^d$.

In an MJP, the state is a probability distribution over d microstates (or coarse-grained "mesostates"). For notational clarity, when referring to probability distributions, we sometimes write the state $\boldsymbol{x}$ as $\boldsymbol{p}$. In a deterministic chemical reaction network, the state is an unnormalized vector of concentrations of d chemical species. For notational clarity, when referring to a deterministic concentration vector, we sometimes write the state $\boldsymbol{x}$ as $\boldsymbol{c}$.

The system is associated with m one-way *transitions* or *reactions* indexed by $\rho \in \{1, \ldots, m\}$, corresponding to jumps between microstates in an MJP or one-way chemical reactions in a CRN. We use the term *reaction observable* to refer to functions of individual reactions, represented by a vector $\boldsymbol{\theta} \in \mathbb{R}^m$. As an example, the reaction observable $\boldsymbol{\theta}$ may indicate heat exchanges; then, θ_ρ indicates the number of joules released to (or absorbed from) a heat bath during reaction ρ. Except where otherwise noted, we make no assumptions about the nature of reaction observables $\boldsymbol{\theta}$ (like antisymmetry).

The one-way *flux* of reaction ρ is the expected number of reactions per unit time and volume, written as $j_\rho(\boldsymbol{x}(t), t)$. The flux depends on state $\boldsymbol{x}(t)$ through system-specific kinetics, and on time t due to external driving (i.e., changing control parameters). Except where otherwise noted, we make no assumptions about the kinetics or external driving. The entire set of one-way fluxes is represented by the vector $\boldsymbol{j} = (j_1, \ldots, j_m) \in \mathbb{R}^m_+$.

The system may also be open to exchange matter with the environment via inflows and outflows, as in a flow reactor. The exchanges are encoded in the vector $\boldsymbol{I}(\boldsymbol{x}(t), t) \in \mathbb{R}^d$, where $I_i(\boldsymbol{x}(t), t)$ is the net outflow (outflow minus inflow) of species i for state $\boldsymbol{x}(t)$ and time t. We will refer to exchanges encoded by $\boldsymbol{I}$ as the *external fluxes*. As we discuss below, chemical systems may also be open to exchanges with external "chemostated" species that are kept at constant concentrations.

The system's state evolves according to

$$\dot{\boldsymbol{x}} = \nabla^\top \boldsymbol{j} - \boldsymbol{I}, \tag{1}$$

where $\nabla \in \mathbb{Z}^{m \times d}$ is a matrix whose entries $\nabla_{\rho i}$ indicate the amount of i created or destroyed by reaction ρ. We often leave dependence on time t and/or state $\boldsymbol{x}$ implicit, as in Eq. (1). The quantity $\nabla^\top \boldsymbol{j}$, which we sometimes call the *net production*, is the change of species due to reaction fluxes. Equation (1) expresses the state evolution as a combination of net production and external fluxes. In closed systems, the state evolution is equal to net production.

The matrix $\nabla \in \mathbb{Z}^{m \times d}$ is called the *incidence matrix* in MJPs and the *stoichiometric matrix* in CRNs. (Note that our convention differs from some of the literature, where "stoichiometric matrix" often refers to the transposed matrix $\nabla^\top \in \mathbb{Z}^{d \times m}$.) Equation (1) can be interpreted as a discrete continuity equation, where the matrix $\nabla^\top$ acts as the negative discrete divergence operator. In turn, the stoichiometric matrix ∇ acts like the discrete gradient operator: The net increase of any state observable $\boldsymbol{\phi}$ due to reaction ρ is $[\nabla \boldsymbol{\phi}]_\rho = \sum_i \nabla_{\rho i} \phi_i$.

B. Thermodynamic forces and entropy production

Each one-way reaction ρ is associated with a *reverse flux* $\tilde{j}_\rho$, where $\tilde{\boldsymbol{j}} = (\tilde{j}_1, \ldots, \tilde{j}_m)$ indicates the vector of all reverse fluxes. We sometimes refer to $\boldsymbol{j}$ as the *forward fluxes* to distinguish them from the reverse fluxes $\tilde{\boldsymbol{j}}$. Each reaction ρ is associated with a *(thermodynamic) force*, sometimes called "reaction affinity" or "affinity" in the literature, defined as the log ratio of the forward and reverse fluxes,

$$f_\rho := \ln \frac{j_\rho}{\tilde{j}_\rho}. \tag{2}$$

We represent the forces across all reactions using the vector

$$\boldsymbol{f} = \ln \frac{\boldsymbol{j}}{\tilde{\boldsymbol{j}}} = (f_1, \dots, f_m) \in \mathbb{R}^m.$$

Like the fluxes, the forces generally depend on the state $\boldsymbol{x}$ and time t, although we leave this implicit in our notation.

For simplicity, in the main text, we focus on systems without odd variables (such as velocities or momenta). However, many of our results generalize to systems with odd variables; see Appendix B for more details. In systems without odd variables, each reaction $\rho \in \{1, \dots, m\}$ is associated one to one with a reverse reaction $\tilde{\rho} \in \{1, \dots, m\}$ such that

$$\nabla_{\rho i} = -\nabla_{\tilde{\rho} i}, \quad \tilde{j}_\rho = j_{\tilde{\rho}}. \tag{3}$$

Observe that $\sum_\rho j_\rho = \sum_\rho j_{\tilde{\rho}}$ since reversal is one to one.

Our main thermodynamic quantity of interest is the entropy production rate. For systems without odd variables, the EPR is written as

$$\sigma = \sum_\rho j_\rho \ln \frac{j_\rho}{\tilde{j}_\rho} = \boldsymbol{j}^\top \boldsymbol{f}. \tag{4}$$

Using Eq. (3), we may also write it as

$$\sigma = \frac{1}{2} \sum_\rho (j_\rho - j_{\tilde{\rho}}) \ln \frac{j_\rho}{j_{\tilde{\rho}}} \geqslant 0. \tag{5}$$

The total entropy production (EP) over time $0 \leqslant t \leqslant T$ is

$$\Sigma = \int_0^T \sigma(t)\, dt, \tag{6}$$

where $\sigma(t)$ is the EPR incurred by the fluxes $\boldsymbol{j}(t)$ at time t. We always write thermodynamic quantities (free energy, forces, chemical potentials, etc.) in dimensionless units.

EPR vanishes only when the reverse and forward fluxes are equal ($\boldsymbol{j} = \tilde{\boldsymbol{j}}$), equivalently when the forces vanish ($\boldsymbol{f} = \boldsymbol{0}$); thus, it quantifies the dynamical irreversibility of the system's fluxes. EPR acquires additional thermodynamic meaning when the condition of *local detailed balance* (LDB) holds, which says that the force f_ρ associated with each reaction ρ is equal to the increase of thermodynamic entropy of the system and its environment due to that reaction [46,55,56]. Assuming LDB, EPR is the expected rate of increase of the thermodynamic entropy of the system and its environment.

C. Example: Markov jump process

To illustrate our formalism, we consider an MJP that represents a stochastic system with d microstates, coupled to one or more thermodynamic reservoirs indexed by α. The probability distribution $\boldsymbol{p}$ evolves according to a master equation,

$$\dot{p}_i = \sum_\alpha \sum_{j=1}^d \left(p_j R^\alpha_{ij} - p_i R^\alpha_{ji} \right), \tag{7}$$

where R^α_{ji} is the transition rate from microstate i to microstate j mediated by reservoir α [57].

In our formalism, there is a one-way reaction ρ for each transition $i \to j$ mediated by reservoir α. It has incidence matrix entries

$$\nabla_{\rho k} = \delta_{kj} - \delta_{ki} \tag{8}$$

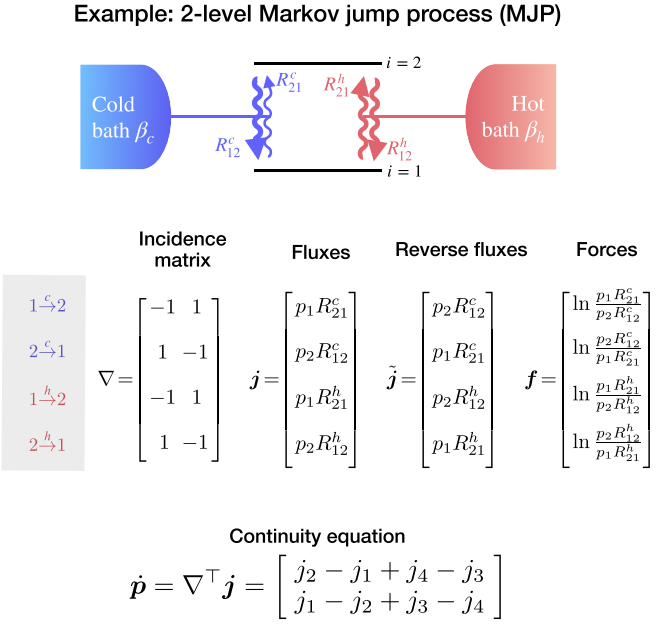

FIG. 1. Formalism illustrated on a two-level MJP coupled to a pair of heat baths, with distribution $p = (p_1, p_2)$. Transitions occur with rates R^c_{21} and R^c_{12} when exchanging energy with the cold bath at inverse temperature β_c, and with rates R^h_{21} and R^h_{12} when exchanging energy with the hot bath at inverse temperature β_h. The four one-way transitions are characterized by the incidence matrix, forward and reverse flux vectors, and force vector.

and flux $j_\rho = p_i R^\alpha_{ji}$. The reverse flux $\tilde{j}_\rho = p_j R^\alpha_{ij}$ is the forward flux across the reverse reaction $\tilde{\rho}$ (transition $j \to i$ mediated by α). In a standard MJP, there are no inflows or outflows, so $\boldsymbol{I} = \boldsymbol{0}$. Given these definitions, the continuity equation (1) is equivalent to the master equation (7). The force across each reaction is

$$f_\rho = \ln \frac{j_\rho}{\tilde{j}_\rho} = \ln \frac{R^\alpha_{ji} p_i}{R^\alpha_{ij} p_j}. \tag{9}$$

In Fig. 1, we illustrate our formalism on a simple MJP that represents a two-level system coupled to a pair of heat baths. We consider this example again in Sec. V B.

As discussed in Appendix A, our formalism can also be applied to "nonlinear MJPs," sometimes used to model many-body systems with mean-field interactions.

D. Example: Chemical reaction network

We illustrate our formalism using a CRN with d chemical species and k reversible reactions,

$$\sum_{i=1}^d \nu_{ri} X_i \rightleftharpoons \sum_{i=1}^d \kappa_{ri} X_i, \quad \forall r \in \{1, \dots, k\}, \tag{10}$$

where ν_{ri} and κ_{ri} specify the stoichiometry of species i as reactant and product in reaction r. The forward and reverse fluxes across each reversible reaction r are indicated as j^+_r and j^-_r. These fluxes generally depend on the concentration vector $\boldsymbol{c}$, possibly in a nonlinear manner. For instance, for

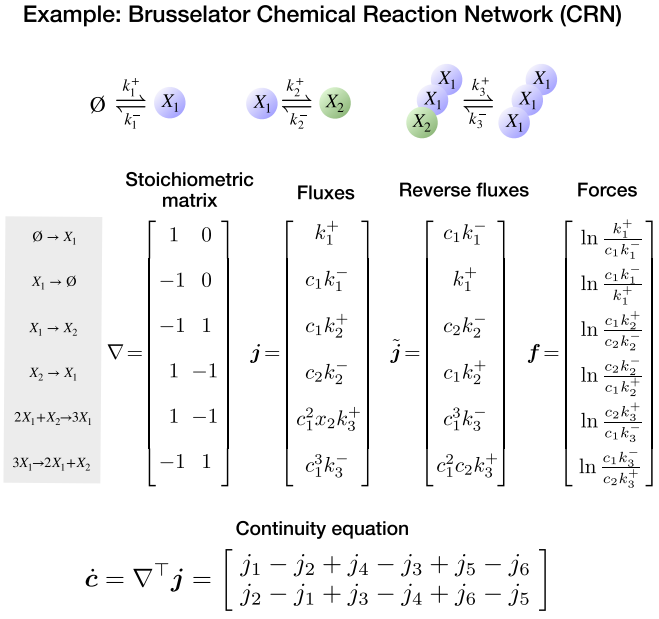

FIG. 2. Formalism illustrated on the Brusselator CRN, a nonlinear chemical oscillator, with concentrations $\boldsymbol{c} = (c_1, c_2)$. There are two species X_1, X_2 and three reversible reactions: $\varnothing \rightleftharpoons X_1$ (inflow), $X_1 \rightleftharpoons X_2$ (conversion), and $2X_1 + X_2 \rightleftharpoons 3X_1$ (second-order autocatalysis). The six one-way reactions are characterized by the stoichiometric matrix, forward and reverse flux vectors, and thermodynamic forces.

mass-action kinetics, the forward flux for reaction r is

$$j_r^+ = k_r^+ \prod_{i=1}^{d} c_i^{\nu_{ri}}, \tag{11}$$

where k_r^+ is the forward rate constant of reaction r.

In our formalism, each reversible reaction $r \in \{1, \dots, k\}$ is treated as two one-way reactions $(\rho, \tilde{\rho})$ with fluxes $j_\rho = j_r^+$, $\tilde{j}_\rho = j_r^-$ and stoichiometry $\nabla_{\rho i} = \kappa_{ri} - \nu_{ri} = -\nabla_{\tilde{\rho} i}$. As mentioned above, a CRN can be open due to inflows and outflows, as specified by the flow vector $\boldsymbol{I}$. The concentrations $\boldsymbol{c}$ evolve according to Eq. (1), sometimes called the "reaction rate equation" in the CRN literature.

In Fig. 2, we illustrate our formalism on the Brusselator [58], a simple chemical oscillator. We also consider this example in Sec. VIII B.

In addition to inflows and outflows $\boldsymbol{I}$, a CRN may be open due to presence of "external" species. Such species are not included in the concentration vector $\boldsymbol{c}$ or the stoichiometric matrix ∇, although their concentrations may still affect kinetics of reactions (via pseudo-rate constants). This is often used to represent "chemostatting," where the concentrations of some species are externally controlled [46]. For example, Prigogine's original model of the Brusselator [58] has four additional species coupled to synthesis/degradation of X_1 and conversion $X_1 \rightleftharpoons X_2$. The simplified Brusselator model in Fig. 2 is derived by assuming that these four species have constant concentrations and thus can be eliminated from the model.

Our general formalism applies to arbitrary CRNs, but the condition of LDB requires further physical assumptions. LDB is valid for reversible elementary reactions, including some nonideal systems with non-mass-action kinetics [59]. LDB is also valid for some types of nonelementary reactions, including reversible Michaelis-Menten kinetics [60], under appropriate definitions of forward and reverse fluxes [61–63].

E. Relative entropy

Relative entropy is an information-theoretic measure of divergence that plays a central role in our work. The relative entropy between a pair of states $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}_+^d$ is defined as

$$D(\boldsymbol{x}\|\boldsymbol{y}) := \sum_{i=1}^{d} \left(x_i \ln \frac{x_i}{y_i} - x_i + y_i \right). \tag{12}$$

It is always non-negative and vanishes only when $\boldsymbol{x} = \boldsymbol{y}$. We use a generalized version of relative entropy, appropriate for unnormalized states that may not sum to unity (such as concentration vectors in CRNs). For normalized probability distributions, it reduces to the well-known Kullback-Leibler divergence, $D(\boldsymbol{p}\|\boldsymbol{q}) = \sum_i p_i \ln(p_i/q_i)$.

We also consider the relative entropy between pairs of flux vectors $\boldsymbol{j}, \boldsymbol{j}' \in \mathbb{R}_+^m$, written as

$$\mathcal{D}(\boldsymbol{j}\|\boldsymbol{j}') := \sum_{\rho=1}^{m} \left(j_\rho \ln \frac{j_\rho}{j'_\rho} - j_\rho + j'_\rho \right). \tag{13}$$

We use the calligraphic $\mathcal{D}$ (rather than D) to distinguish the relative entropy between fluxes.

Importantly, the EPR can be written as the relative entropy between the forward and reverse fluxes,

$$\sigma = \mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}}) = \sum_{\rho} \left(j_\rho \ln \frac{j_\rho}{\tilde{j}_\rho} - j_\rho + \tilde{j}_\rho \right) \geqslant 0. \tag{14}$$

For systems without odd variables, this expression follows from Eq. (4) and $\sum j_\rho = \sum j_{\tilde{\rho}} = \sum \tilde{j}_\rho$. However, Eq. (14) also correctly captures EPR for systems with odd variables, where Eq. (4) may not apply (see Appendix B).

The generalized relative entropy (14) induces a decomposition of EPR into non-negative contributions from individual one-way reactions:

$$\sigma = \sum_{\rho} \sigma^{(\rho)}, \quad \sigma^{(\rho)} := j_\rho \ln \frac{j_\rho}{\tilde{j}_\rho} - j_\rho + \tilde{j}_\rho. \tag{15}$$

The contribution from reaction ρ can also be written as

$$\sigma^{(\rho)} = j_\rho (f_\rho - 1 + e^{-f_\rho}) \geqslant 0. \tag{16}$$

Our EPR decomposition (15) is finer grained than the usual one, which considers reversible reaction pairs [44, Eq. (68)]:

$$\sigma = \frac{1}{2} \sum_{\rho} \sigma_{\text{rev}}^{(\rho)}, \quad \sigma_{\text{rev}}^{(\rho)} := \sigma^{(\rho)} + \sigma^{(\tilde{\rho})} = (j_\rho - \tilde{j}_\rho) f_\rho.$$

III. BACKGROUND ON NONEQUILIBRIUM FREE ENERGY

In this section, we review the concept of nonequilibrium free energy in conservative systems. We then discuss a previously proposed generalization to nonconservative systems based on steady states.

A. Nonequilibrium free energy in conservative systems

We use the term *conservative system* to refer to a system governed by a nonequilibrium free energy $\mathcal{F}(\boldsymbol{x}, t)$. In such systems, the thermodynamic forces can be expressed in terms of the chemical potential $\boldsymbol{\mu} = \mathrm{grad}_{\boldsymbol{x}}\mathcal{F}(\boldsymbol{x}, t)$ as

$$\boldsymbol{f} = -\nabla\boldsymbol{\mu}, \tag{17}$$

where $\nabla \in \mathbb{R}^{m\times d}$ is the discrete gradient operator defined above (i.e., stoichiometric matrix). Forces that have the gradient form (17) are termed conservative forces.

The internal dynamics of a conservative system can be expressed as a gradient flow for $\mathcal{F}$. To show this, we write the continuity equation (1) as

$$\dot{\boldsymbol{x}} = \tfrac{1}{2}\nabla^\top(\boldsymbol{j} - \tilde{\boldsymbol{j}}) - \boldsymbol{I} = \nabla^\top L\boldsymbol{f} - \boldsymbol{I}. \tag{18}$$

Here, we used $\nabla^\top\boldsymbol{j} = \frac{1}{2}\nabla^\top(\boldsymbol{j} - \tilde{\boldsymbol{j}})$, since reverse reactions have antisymmetric stoichiometry, and defined L as a diagonal matrix with entries $L_{\rho\rho} = \frac{1}{2}j_\rho(1 - e^{-f_\rho})/f_\rho \geqslant 0$. Using Eq. (17), we can further write

$$\dot{\boldsymbol{x}} = -(\nabla^\top L\nabla)\mathrm{grad}_{\boldsymbol{x}}\mathcal{F} - \boldsymbol{I}. \tag{19}$$

The (first) net production term is a gradient flow for $\mathcal{F}$ (the positive-semidefinite matrix $\nabla^\top L\nabla$ defines a metric). The relationship between conservative systems and gradient flows was first derived for MJPs in Ref. [1] and CRNs with mass-action kinetics in Ref. [2]. Since then, the connection between gradient flows and large deviations has inspired an active research program [64–66].

Consider a conservative system that relaxes without time-dependent driving ($\mathcal{F}$ does not depend explicitly on time) or external fluxes ($\boldsymbol{I} = \boldsymbol{0}$). In this case, $\mathcal{F}$ decreases monotonically over time until the system reaches a stationary equilibrium state $\boldsymbol{x}^{\mathrm{eq}}$ of minimal nonequilibrium free energy. At this state, the chemical potential reaches its equilibrium value and the thermodynamic forces vanish:

$$\boldsymbol{\mu}^{\mathrm{eq}} = \mathrm{grad}_{\boldsymbol{x}}\mathcal{F}(\boldsymbol{x}^{\mathrm{eq}}), \quad \boldsymbol{f} = -\nabla\boldsymbol{\mu}^{\mathrm{eq}} = \boldsymbol{0}.$$

By the definition of the thermodynamic forces, forward and reverse fluxes in equilibrium must obey the condition of *detailed balance*,

$$\boldsymbol{j}(\boldsymbol{x}^{\mathrm{eq}}) = \tilde{\boldsymbol{j}}(\boldsymbol{x}^{\mathrm{eq}}). \tag{20}$$

The condition of conservative forces (17) is invariant if $\boldsymbol{\mu}$ is shifted by any null vector of ∇. Such null vectors represent conserved quantities that do not affect thermodynamic forces. For our purposes, it will be convenient to consider the equilibrium chemical potential as a reference null vector (recall $-\nabla\boldsymbol{\mu}^{\mathrm{eq}} = \boldsymbol{0}$) and define $\boldsymbol{\phi}^{\mathrm{eq}} = \boldsymbol{\mu} - \boldsymbol{\mu}^{\mathrm{eq}}$ as the purely nonequilibrium contribution to the chemical potential,

$$\boldsymbol{f} = -\nabla\boldsymbol{\phi}^{\mathrm{eq}}. \tag{21}$$

We usually refer to $\boldsymbol{\phi}^{\mathrm{eq}}$ as the "free energy potential."

The preceding statements hold for any conservative system. However, things become less abstract for MJPs [6] and ideal CRNs with mass-action kinetics [46]. In these cases, $\boldsymbol{\phi}^{\mathrm{eq}}$ has an explicit form given by the gradient of the relative entropy between the current state and the equilibrium state:

$$\boldsymbol{\phi}^{\mathrm{eq}} = \mathrm{grad}_{\boldsymbol{x}}D(\boldsymbol{x}\|\boldsymbol{x}^{\mathrm{eq}}), \quad \phi_i^{\mathrm{eq}} = \ln\frac{x_i}{x_i^{\mathrm{eq}}}. \tag{22}$$

As a concrete example, consider an MJP that evolves according to Eq. (7). Suppose there is a single reservoir and an equilibrium distribution $\boldsymbol{p}^{\mathrm{eq}}$ that obeys detailed balance, $p_i^{\mathrm{eq}}R_{ji} = p_j^{\mathrm{eq}}R_{ij}$. Then, given any other distribution $\boldsymbol{p}$, the forces $f_\rho := \ln(p_iR_{ji}/p_jR_{ij}) = \phi_i^{\mathrm{eq}} - \phi_j^{\mathrm{eq}}$ are conservative for the potential $\phi_i^{\mathrm{eq}} = \ln(p_i/p_i^{\mathrm{eq}})$.

The relative entropy $D(\boldsymbol{x}\|\boldsymbol{x}^{\mathrm{eq}})$ in Eq. (22) quantifies the purely nonequilibrium contribution to $\mathcal{F}$ [6],

$$D(\boldsymbol{x}\|\boldsymbol{x}^{\mathrm{eq}}) = \mathcal{F}(\boldsymbol{x}) - \mathcal{F}(\boldsymbol{x}^{\mathrm{eq}}). \tag{23}$$

It appears under many names in the literature, including "pseudo-Helmholtz function" [46], "shear Lyapunov function" [46], negative "nonequilibrium Massieu potential" [67], or simply "free energy" [4,28,68,69]. It plays a central role in the thermodynamics of MJPs and closed ideal CRNs with conservative forces. For instance, it is proportional to the maximal work that can be extracted while bringing such systems from $\boldsymbol{x}$ to $\boldsymbol{x}^{\mathrm{eq}}$ [5,6,46]. Conversely, its decrease under free relaxation is equal to EPR [70],

$$\sigma(t) = -\partial_t D(\boldsymbol{x}(t)\|\boldsymbol{x}^{\mathrm{eq}}) = -\dot{\boldsymbol{x}}^\top\boldsymbol{\phi}^{\mathrm{eq}} \geqslant 0. \tag{24}$$

The relative entropy has an important statistical interpretation in large-deviation theory. Consider an ensemble of n independent and identical fluctuating systems, and let $\boldsymbol{P}$ be the empirical distribution of their microstates at a given point in time. Due to stochastic fluctuations, $\boldsymbol{P}$ will itself be a random variable. Under the equilibrium distribution $\boldsymbol{p}^{\mathrm{eq}}$, the probability of observing the empirical distribution $\boldsymbol{P} \approx \boldsymbol{p}$ due to a fluctuation scales as [3,4,71]

$$\Pr[\boldsymbol{P} \approx \boldsymbol{p}] \asymp e^{-n\,D(\boldsymbol{p}\|\boldsymbol{p}^{\mathrm{eq}})}, \tag{25}$$

where $\asymp$ indicates equality up to subexponential factors in the scale parameter n. This type of expression is known as a large-deviation principle (LDP), with the relative entropy $D(\boldsymbol{x}\|\boldsymbol{x}^{\mathrm{eq}})$ playing the role of the "rate function" [72]. Equation (25) implies that distributions with larger nonequilibrium free energy are exponentially less likely to emerge from equilibrium fluctuations. A similar result can be derived for stochastic chemical systems. Consider an ideal well-mixed system in a large reactor volume V, and let the random variable $\boldsymbol{C}$ indicate the empirical concentration in an average subvolume. The probability that concentration vector $\boldsymbol{c}$ emerges as a fluctuation, given true equilibrium concentrations $\boldsymbol{c}^{\mathrm{eq}}$, scales as [73]

$$\Pr[\boldsymbol{C} \approx \boldsymbol{c}] \asymp e^{-V\,D(\boldsymbol{c}\|\boldsymbol{c}^{\mathrm{eq}})}. \tag{26}$$

B. The steady-state approach to nonconservative systems

Recent work in thermodynamics has focused on *nonconservative systems*, whose forces cannot be expressed in the form of Eq. (21) for any $\boldsymbol{\phi}$. Such systems are not governed by a nonequilibrium free energy and they have nonequilibrium steady states that do not satisfy detailed balance.

Nonetheless, it has been suggested that it may be possible to define a "generalized" free energy for nonconservative systems. A related idea is that the fluxes, forces, and EPR may be decomposed into excess and housekeeping contributions. In particular, EPR may be decomposed as

$$\sigma = \sigma_{\mathrm{ex}} + \sigma_{\mathrm{hk}}. \tag{27}$$

The excess EPR is associated with the generalized free energy and vanishes in steady state, while the housekeeping EPR is the genuine nonequilibrium contribution arising from cyclic fluxes.

The best-known generalized free energy and excess/housekeeping decomposition is based on nonequilibrium steady states [30,33,44]. In this approach, the generalized free energy is defined as the large-deviation rate function of steady-state fluctuations, written as $\Phi_{\rm ss}$, sometimes called the *quasipotential*. Let random variable $\boldsymbol{X}$ indicate the empirical state averaged across a stochastic system of size V [e.g., number of independent copies as in Eq. (25) or reactor volume as in Eq. (26)]. Then, the quasipotential governs the probability that some state $\boldsymbol{x}$ (distribution or concentration vector) emerges as a steady-state fluctuation,

$$\Pr[\boldsymbol{X} \approx \boldsymbol{x}] \asymp e^{-V\,\Phi_{\rm ss}(\boldsymbol{x})}. \tag{28}$$

Importantly, the quasipotential decreases monotonically under free relaxation; thus, it serves as a Lyapunov function [44]. However, the dynamics of nonconservative systems are not a gradient flow for the quasipotential, nor any other state function.

For systems described at the microscopic fluctuating level, the quasipotential is the relative entropy between the system's actual distribution and the steady-state distribution, $\Phi_{\rm ss}(\boldsymbol{p}) = D(\boldsymbol{p}\|\boldsymbol{p}^{\rm ss})$. This quasipotential is also used to define an excess/housekeeping decomposition of EPR called the HS [36] or the adiabatic/nonadiabatic [57] decomposition. In analogy to Eq. (24), the HS excess EPR is the decrease of the quasipotential due to relaxation,

$$\sigma_{\rm ex}^{\rm HS}(t) := -\partial_t D(\boldsymbol{p}(t)\|\boldsymbol{p}^{\rm ss}) \geqslant 0, \tag{29}$$

where $\boldsymbol{p}^{\rm ss}$ is the steady-state distribution specified by the system's parameters at time t. It can be written in slightly more explicit form as

$$\sigma_{\rm ex}^{\rm HS} = -\boldsymbol{j}^\top \nabla \boldsymbol{\phi}^{\rm ss} = -\dot{\boldsymbol{p}}^\top \boldsymbol{\phi}^{\rm ss}. \tag{30}$$

Here, $\boldsymbol{\phi}^{\rm ss} := \mathrm{grad}_{\boldsymbol{p}} D(\boldsymbol{p}\|\boldsymbol{p}^{\rm ss}) = \ln(\boldsymbol{p}/\boldsymbol{p}^{\rm ss})$ is the generalization of the free energy potential, which we sometimes refer to as the *steady-state potential*. [This terminology should not be confused with the term "quasipotential," which refers to a function $\Phi_{\rm ss}(\boldsymbol{x})$ over thermodynamic states.] The HS housekeeping EPR is the remainder EPR,

$$\sigma_{\rm hk}^{\rm HS} := \sigma - \sigma_{\rm ex}^{\rm HS}. \tag{31}$$

This housekeeping contribution is non-negative for MJPs without odd variables. (With odd variables, it can take on unphysical negative values; see Appendix B.)

In systems described at the macroscopic level, including deterministic CRNs, the situation is more complicated. For CRNs that obey *complex balance*, the quasipotential is simply the relative entropy $\Phi_{\rm ss}(\boldsymbol{c}) = D(\boldsymbol{c}\|\boldsymbol{c}^{\rm ss})$, where $\boldsymbol{c}^{\rm ss}$ is the steady-state concentration of the deterministic dynamics (1) [73,74]. Excess and housekeeping EPR are defined in direct analogy to Eqs. (29) and (31) [71,74]. However, complex balance is a highly restrictive condition that excludes most CRNs, including those that exhibit rich nonlinear phenomena like oscillations and chaos [47].

Without complex balance, the quasipotential $\Phi_{\rm ss}(\boldsymbol{c})$ is no longer the relative entropy $D(\boldsymbol{c}\|\boldsymbol{c}^{\rm ss})$ [44,75]. In fact, there may not even exist a stable fixed point $\boldsymbol{c}^{\rm ss}$ for the deterministic dynamics (1), because the large-volume limit (used to derive the deterministic CRN description) may not commute with the long-time limit (used to derive steady-state behavior). This phenomenon, known as Keizer's paradox [76,77], highlights the implicit choice of timescales involved in the large-volume and long-time limits. In general deterministic CRNs, computing $\Phi_{\rm ss}$ is challenging, although sophisticated numerical methods have been developed [78–82]. However, once it is found, macroscopic HS excess and housekeeping EPR can be defined as in Eqs. (29) and (31), with the relative entropy replaced by $\Phi_{\rm ss}$ [44].

The steady-state approach has found numerous applications in physics, chemistry, and biology. For example, it has been used to study transition rates between stochastic attractors, quantify stability and fluctuations in stationarity, and visualize potential landscapes that govern long-term relaxations [30,33,44,83,84].

On the other hand, the steady-state approach is less useful for studying transient systems that never approach stationarity, including systems observed over short timescales and systems driven by changing external parameters and/or flows. Because the quasipotential is defined via steady-state statistics, rather than "local-in-time" properties such as instantaneous fluxes or forces, its physical meaning is unclear in systems far from steady state [41,85,86]. Importantly, the situation is different in conservative systems. There, the quasipotential reduces to the nonequilibrium free energy, which is directly related to instantaneous forces by the conservative force expression (21).

In the following, we propose an alternative, local-in-time definition of the generalized free energy and excess/housekeeping decomposition. Our generalized free energy is defined as the conservative part of the thermodynamic forces, and the excess/housekeeping EPRs quantify the conservative/nonconservative contributions to dissipation. Using large-deviation theory, we will show that our generalized free energy governs short-time current fluctuations in the transient regime. This may be contrasted to the quasipotential $\Phi_{\rm ss}$, which governs state fluctuations in the stationary regime. Our definition is well suited to study dynamical properties of transient systems, but less useful for studying stability and fluctuations in stationarity. In this sense, our approach is different and complementary to the steady-state approach.

IV. VARIATIONAL APPROACH: CONSERVATIVE SYSTEMS

We now introduce our first set of results. Specifically, we show that in conservative systems, the nonequilibrium free energy can be derived from a variational principle that makes no explicit reference to equilibrium states. We extend our principle to nonconservative systems in Sec. V.

Before proceeding, recall from Eq. (14) that the EPR can be expressed as the relative entropy between forward and backward fluxes, $\sigma = \mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}})$. This allows us to express the EPR in a variational way,

$$\sigma = \max_{\boldsymbol{\theta}\in\mathbb{R}^m} [\boldsymbol{j}^\top \boldsymbol{\theta} - \tilde{\boldsymbol{j}}^\top (e^{\boldsymbol{\theta}} - \mathbf{1})], \tag{32}$$

where we maximize over all possible reaction observables. To derive this expression, observe that $\min_{\boldsymbol{\theta}} \mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}} \circ e^{\boldsymbol{\theta}}) = 0$, since $\mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}} \circ e^{\boldsymbol{\theta}})|_{\boldsymbol{\theta}=f} = \mathcal{D}(\boldsymbol{j}\|\boldsymbol{j}) = 0$. Thus, we can write the EPR as

$$\begin{aligned}\sigma &= \mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}}) - \min_{\boldsymbol{\theta}} \mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}} \circ e^{\boldsymbol{\theta}}) \\ &= \max_{\boldsymbol{\theta}} [\mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}}) - \mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}} \circ e^{\boldsymbol{\theta}})].\end{aligned}$$

Equation (32) follows by expanding the relative entropy terms and simplifying. The optimal observable $\boldsymbol{\theta}^\star = \boldsymbol{f}$ is unique due to the strict concavity of the objective.

The variational expression (32) is the Legendre transform of the function $\boldsymbol{\theta} \mapsto \tilde{\boldsymbol{j}}^\top (e^{\boldsymbol{\theta}} - \mathbf{1})$ evaluated at $\boldsymbol{j}$. As we will see below, this function is the cumulant generating function (CGF) of dynamical fluctuations. In large deviations, the variational expression (32) is often called the "Donsker-Varadhan" form.

Equation (32) provides a family of lower bounds on the EPR, one for each reaction observable. The EPR is achieved by maximizing over the choice of observable, and the thermodynamic forces are recovered as the optimal observables. In the setting of stochastic master equations, a related variational principle was recently proposed as a technique for thermodynamic inference [87,88]. Moreover, it applies to both conservative and nonconservative systems. However, in conservative systems, the forces have the conservative form $\boldsymbol{f} = -\nabla \boldsymbol{\phi}^{\mathrm{eq}}$. Therefore, we may restrict the optimization to reaction observables having the form $\boldsymbol{\theta} = -\nabla \boldsymbol{\phi}$ for some state observable $\boldsymbol{\phi} \in \mathbb{R}^d$,

$$\sigma = \max_{\boldsymbol{\phi} \in \mathbb{R}^d} [-\boldsymbol{j}^\top \nabla \boldsymbol{\phi} - \tilde{\boldsymbol{j}}^\top (e^{-\nabla \boldsymbol{\phi}} - \mathbf{1})]. \tag{33}$$

This expression may be written using only the forward fluxes,

$$\sigma = \max_{\boldsymbol{\phi} \in \mathbb{R}^d} [-\boldsymbol{j}^\top \nabla \boldsymbol{\phi} - \boldsymbol{j}^\top (e^{\nabla \boldsymbol{\phi}} - \mathbf{1})], \tag{34}$$

where we used $\tilde{\boldsymbol{j}}^\top e^{-\nabla \boldsymbol{\phi}} = \boldsymbol{j}^\top e^{\nabla \boldsymbol{\phi}}$ from Eq. (3) and antisymmetry of the stoichiometric matrix.

Equation (34) is a variational expression of the EPR in conservative systems. Taking derivatives shows that the optimal $\boldsymbol{\phi}^\star$ obeys

$$\nabla^\top (\boldsymbol{j} \circ e^{\nabla \boldsymbol{\phi}^\star}) = -\nabla^\top \boldsymbol{j}. \tag{35}$$

In conservative systems, the free energy potential satisfies this relation, as can be verified using Eq. (21), thus $\boldsymbol{\phi}^\star = \boldsymbol{\phi}^{\mathrm{eq}}$. Moreover, due to strict concavity of the objective, $\boldsymbol{\phi}^{\mathrm{eq}}$ is the unique optimum, up to the choice of a null vector of ∇ that represents a conserved quantity. If desired, it is possible to identify $\boldsymbol{\phi}^{\mathrm{eq}}$ that is consistent with the conserved quantities encoded in $\boldsymbol{x}$ (see SM1 in the Supplemental Material [53]).

We may consider our variational principle (34) from perspective of convex duality [89]. Observe that Eq. (34) involves the maximization of a concave objective over state observables. By duality, it has an equivalent formulation as the minimization of a convex objective over flux vectors:

$$\sigma = \min_{\boldsymbol{j}' \in \mathbb{R}^m_+} \mathcal{D}(\boldsymbol{j}'\|\tilde{\boldsymbol{j}}), \quad \text{where } \nabla^\top \boldsymbol{j}' = \nabla^\top \boldsymbol{j}. \tag{36}$$

To show the equivalence, we write Eq. (36) in its Lagrangian dual form [89, Chap. 5],

$$\sigma = \max_{\boldsymbol{\phi} \in \mathbb{R}^d} \min_{\boldsymbol{j}' \in \mathbb{R}^m_+} [\mathcal{D}(\boldsymbol{j}'\|\tilde{\boldsymbol{j}}) + \boldsymbol{\phi}^\top \nabla^\top (\boldsymbol{j}' - \boldsymbol{j})], \tag{37}$$

where $\boldsymbol{\phi} \in \mathbb{R}^d$ indicates the Lagrangian multipliers. The inner optimization can be solved by taking derivatives. After a bit of algebra, we find that the optimal fluxes have the form $\tilde{\boldsymbol{j}} \circ e^{-\nabla \boldsymbol{\phi}}$. Plugging back into Eq. (37) and simplifying shows equivalence to Eq. (34). The optimal Lagrange multipliers are equal to $\boldsymbol{\phi}^{\mathrm{eq}}$, while the optimal fluxes are equal to the actual forward fluxes $\boldsymbol{j}^\star = \tilde{\boldsymbol{j}} \circ e^{-\nabla \boldsymbol{\phi}^{\mathrm{eq}}} = \boldsymbol{j}$.

The variational expressions [Eqs. (34)–(36)] comprise our first set of results. They demonstrate that $\boldsymbol{\phi}^{\mathrm{eq}}$ may be derived in a simple way from the fluxes, without any explicit reference to equilibrium states.

V. VARIATIONAL APPROACH: NONCONSERVATIVE SYSTEMS

This section contains most of our main results. We extend our variational principle to nonconservative systems, which leads to our definitions of the generalized free energy and the excess/housekeeping decomposition. We also show that our variational principle satisfies an important consistency condition under coarse graining. We discuss connections to information geometry and discuss the linear-response regime of slow evolution.

A. Overview

As we show above, in conservative systems, the variational principle for EPR (32) can be restricted to observables having the gradient form $\boldsymbol{\theta} = -\nabla \boldsymbol{\phi}$. This leads to a variational expression of EPR and $\boldsymbol{\phi}^{\mathrm{eq}}$ [Eq. (34)].

In nonconservative systems, restricting Eq. (32) to gradient observables gives the following variational expression:

$$\sigma_{\mathrm{ex}} = \max_{\boldsymbol{\phi} \in \mathbb{R}^d} [-\boldsymbol{j}^\top \nabla \boldsymbol{\phi} - \boldsymbol{j}^\top (e^{\nabla \boldsymbol{\phi}} - \mathbf{1})]. \tag{38}$$

This does not recover the full EPR unless all forces are conservative. Instead, the value of σ_{ex} defines the excess EPR, the conservative contribution to dissipation. For a temporally extended process, the time-integrated excess EP is $\Sigma_{\mathrm{ex}} = \int \sigma_{\mathrm{ex}}(t)\, dt$.

We have the lower bound $\sigma_{\mathrm{ex}} \geqslant 0$, since the objective in Eq. (38) vanishes when $\boldsymbol{\phi} = \mathbf{0}$. We also have the upper bound $\sigma_{\mathrm{ex}} \leqslant \sigma$, since Eq. (38) is a restriction of the maximization (32) to $\boldsymbol{\theta} \in \mathrm{Im}\nabla$.

The optimal state observable $\boldsymbol{\phi}^\star$ in Eq. (38) defines our generalized free energy potential, and $-\nabla \boldsymbol{\phi}^\star$ represents the conservative contribution to the thermodynamic forces. In the following, we refer to $\boldsymbol{\phi}^\star$ as the *generalized potential*. The generalized potential satisfies the optimality condition,

$$\nabla^\top (\boldsymbol{j} \circ e^{\nabla \boldsymbol{\phi}^\star}) = -\nabla^\top \boldsymbol{j}, \tag{39}$$

which is the analog of Eq. (35). Recall that $\nabla^\top \boldsymbol{j}$ is the net production of species due to reaction fluxes $\boldsymbol{j}$. Equation (39) defines $\boldsymbol{\phi}^\star$ by the property that net production is reversed when the fluxes are exponentially tilted by $\nabla \boldsymbol{\phi}^\star$. The optimality condition (39) determines $\boldsymbol{\phi}^\star$ up to the null-space of

∇, representing conserved quantities. Although not necessary for most of our results, $\boldsymbol{\phi}^\star$ can always be chosen within this null-space to satisfy the system's conservation laws (see SM1 in the Supplemental Material [53]).

Equation (38) can also be rearranged as

$$\sigma_{\text{ex}} = \max_{\boldsymbol{\phi}\in\mathbb{R}^d}[-2\boldsymbol{j}^\top\nabla\boldsymbol{\phi} - \boldsymbol{j}^\top(e^{\nabla\boldsymbol{\phi}} - \nabla\boldsymbol{\phi} - \mathbf{1})]. \tag{40}$$

The first term is twice the average change of $\boldsymbol{\phi}$ due to reaction fluxes. As we show below, the second term quantifies the fluctuations of $\boldsymbol{\phi}$, and it is always non-negative (since $e^x - x - 1 \geqslant 0$ for all x). Since excess EPR is non-negative, the generalized potential obeys

$$-\boldsymbol{j}^\top\nabla\boldsymbol{\phi}^\star \geqslant 0, \tag{41}$$

which reduces to $-\dot{\boldsymbol{x}}^\top\boldsymbol{\phi}^\star \geqslant 0$ in a closed system without external fluxes. This is analogous to Eq. (24) in conservative systems, $-\dot{\boldsymbol{x}}^\top\boldsymbol{\phi}^{\text{eq}} \geqslant 0$. Equation (41) implies that the net production points downhill in the generalized potential. However, Eq. (41) does not imply Lyapunov stability in the same way as Eq. (24), because $\boldsymbol{\phi}^\star$ need not be the gradient of any state function (like $\mathcal{F}$).

The variational expression (38) is the central expression of this paper. Its physical meaning is explored further below in this section and in Sec. VI, where we relate it to large-deviation theory and thermodynamic speed limits. To our knowledge, this expression is largely unknown in stochastic thermodynamics, either in conservative or nonconservative systems. One exception is a paper by Shiraishi and Saito [90], who proposed a different variational expression for EPR in conservative MJPs. In Sec. VII D, we show that the result of Ref. [90] is a special case of Eq. (34), allowing us to extend the main result of Ref. [90] to nonconservative MJPs.

The numerical values of excess EPR and the generalized potential can be found by solving the convex optimization (38) using standard numerical algorithms. Also, as shown in Sec. V D, these quantities can also be found in closed form for a special class of systems (those with conservative forces after coarse graining). Finally, for a time-extended process characterized by a flux trajectory $\boldsymbol{j}(t)$, the generalized potential can be found by solving an ordinary differential equation. The time derivative of $\boldsymbol{\phi}^\star(t)$ can be found by differentiating both sides of Eq. (39) and rearranging. After simplifying, this gives

$$\dot{\boldsymbol{\phi}}^\star(t) = -\tfrac{1}{2}\mathcal{H}_t^+\nabla^\top(\dot{\boldsymbol{j}}(t) \circ (e^{\nabla\boldsymbol{\phi}^\star(t)} + \mathbf{1})), \tag{42}$$

where $\mathcal{H}_t^+$ is the pseudoinverse of the positive-semidefinite matrix $\mathcal{H}_t := \frac{1}{2}\nabla^\top\text{diag}(\boldsymbol{j}(t) \circ e^{\nabla\boldsymbol{\phi}^\star(t)})\nabla$. Given a known initial $\boldsymbol{\phi}^\star(0)$, Eq. (42) can be solved to find $\boldsymbol{\phi}^\star(t)$ at all t.

B. Example: Two-level MJP

To make things concrete, we illustrate our approach on a minimal example, the two-level MJP from Fig. 1. For this system, the dynamics obey

$$\dot{p}_2 = j_{21}^c + j_{21}^h - j_{12}^c - j_{12}^h,$$

where $\dot{p}_1 = -\dot{p}_2$ by conservation of probability. The optimality condition (39) reduces to

$$\dot{p}_2 = -\big(j_{21}^c + j_{21}^h\big)e^{\phi_2^\star - \phi_1^\star} + \big(j_{12}^c + j_{12}^h\big)e^{\phi_1^\star - \phi_2^\star}.$$

This condition is satisfied by the potential

$$\boldsymbol{\phi}^\star = \left(0, \ln\frac{j_{12}^c + j_{12}^h}{j_{21}^c + j_{21}^h}\right), \tag{43}$$

which is unique up to a scalar constant (the null-space of ∇ is one dimensional, representing conservation of probability). Plugging $\boldsymbol{\phi}^\star$ into Eq. (38) and simplifying gives

$$\sigma_{\text{ex}} = \big(j_{21}^c + j_{21}^h - j_{12}^c - j_{12}^h\big)\ln\frac{j_{21}^c + j_{21}^h}{j_{12}^c + j_{12}^h}. \tag{44}$$

In this example, the excess EPR is the same as the EPR incurred by a system with two one-way fluxes: the forward flux $j_{21}^c + j_{21}^h$ and the reverse flux $j_{12}^c + j_{12}^h$. (We will see below that this exemplifies a more general "coarse-graining" principle.) Equation (44) suggests an "effective" conservative force across the transition $1 \to 2$,

$$\ln\frac{j_{21}^c + j_{21}^h}{j_{12}^c + j_{12}^h} = (\ln p_1 + \beta^{\text{eff}}E_1) - (\ln p_2 + \beta^{\text{eff}}E_2),$$

where $\beta^{\text{eff}} = (\ln\frac{R_{21}^h + R_{21}^c}{R_{12}^h + R_{12}^c})/(E_1 - E_2)$ is an effective inverse temperature and E_1, E_2 refer to microstate energies. Note that β^{eff} does not depend on the probability distribution $\boldsymbol{p}$, only on the energy gap and the kinetics.

C. Excess EPR and nonstationarity

As mentioned, excess EPR captures the conservative contribution to EPR. Here, we use duality to show that it can also be understood as quantifying the nonstationarity of the dynamics.

The variational expression (38) has the dual form,

$$\sigma_{\text{ex}} = \min_{\boldsymbol{j}'\in\mathbb{R}_+^m}\mathcal{D}(\boldsymbol{j}'\|\tilde{\boldsymbol{j}}), \quad \text{where } \nabla^\top\boldsymbol{j}' = \nabla^\top\boldsymbol{j}, \tag{45}$$

derived in the same way as Eq. (36). The generalized potential $\boldsymbol{\phi}^\star$ specifies the optimal Lagrange multipliers for the constraints. The optimal fluxes, written as

$$\boldsymbol{j}^\star = \tilde{\boldsymbol{j}} \circ e^{-\nabla\boldsymbol{\phi}^\star}, \tag{46}$$

are generally not equal to the forward fluxes $\boldsymbol{j}$, except in conservative systems. The excess EPR can be expressed using these optimal fluxes as

$$\sigma_{\text{ex}} = \mathcal{D}(\boldsymbol{j}^\star\|\tilde{\boldsymbol{j}}). \tag{47}$$

Excess EPR vanishes when the reverse fluxes satisfy the constraints $\nabla^\top\tilde{\boldsymbol{j}} = \nabla^\top\boldsymbol{j}$, which is equivalent to $\nabla^\top\boldsymbol{j} = \mathbf{0}$ due to antisymmetry. The optimality condition (39) implies that the generalized potential also vanishes when $\nabla^\top\boldsymbol{j} = \mathbf{0} \Rightarrow \boldsymbol{\phi}^\star = \mathbf{0}$. In other words, the excess EPR and generalized potential vanish when net production vanishes, $\nabla^\top\boldsymbol{j} = \mathbf{0}$. In a closed system without external fluxes, $\dot{\boldsymbol{x}} = \nabla^\top\boldsymbol{j}$; therefore, they vanish precisely in steady state.

Equation (45) allows us to interpret our excess EPR as an information-theoretic optimal-transport cost [91]. In general, "optimal transport" studies the minimal cost of transforming a system between two states, thus giving operational definitions of distance and speed [92]. Equation (45) implies that the excess EPR is the minimal cost necessary to achieve the same dynamical evolution as that induced by the true fluxes, $\nabla^\top\boldsymbol{j}$.

The information-theoretic cost function $\boldsymbol{j}' \mapsto \mathcal{D}(\boldsymbol{j}'\|\tilde{\boldsymbol{j}})$ quantifies the breaking of time-reversal symmetry relative to the reverse fluxes $\tilde{\boldsymbol{j}}$. We discuss the relation to another optimal-transport distance (Wasserstein distance) in Sec. VII.

D. Consistency under coarse graining

Our generalized potential and excess EPR satisfy an important consistency condition under coarse graining. To introduce this idea, observe that a system may have multiple different reactions with the same stoichiometry. In fact, this is a common way to drive nonconservative systems. For example, in the two-level MJP discussed above (Fig. 1), the two reversible transitions between levels 1 and 2 have identical stoichiometry. As a result, the incidence matrix ∇ has duplicate rows (1, 3 and 2, 4). As another example, in the Brusselator CRN (Fig. 2), the reversible reactions $X_1 \rightleftharpoons X_2$ and $2X_1 + X_2 \rightleftharpoons 3X_1$ have opposite stoichiometry. As a result, ∇ has duplicate rows (3, 6 and 4, 5).

We now introduce a coarse-graining procedure that combines reactions with the same stoichiometry. Given a stoichiometric matrix ∇, the coarse-grained stoichiometric matrix is defined as $\overline{\nabla} \in \mathbb{Z}^{m' \times d}$ $(m' \leqslant m)$, where any duplicate rows (reactions with the same stoichiometry) are merged. Coarse-grained forward fluxes $\overline{\boldsymbol{j}} \in \mathbb{R}_+^{m'}$ are defined by summing fluxes j_ρ from merged reactions, and similarly for the coarse-grained reverse fluxes $\overline{\tilde{\boldsymbol{j}}} \in \mathbb{R}_+^{m'}$. In an MJP where the same transitions are mediated by different reservoirs α, the coarse-graining procedure will sum rate matrices induced by different reservoirs, $\bar{R} = \sum_\alpha R^\alpha$.

It is clear that this coarse graining preserves net production, $\nabla^\top \boldsymbol{j} = \overline{\nabla}^\top \overline{\boldsymbol{j}}$. The quantity $\boldsymbol{j}^\top(e^{\nabla\boldsymbol{\phi}} - \mathbf{1})$, which appears in our variational principle (38), is also invariant under coarse graining:

$$\boldsymbol{j}^\top(e^{\nabla\boldsymbol{\phi}} - \mathbf{1}) = \overline{\boldsymbol{j}}^\top(e^{\overline{\nabla}\boldsymbol{\phi}} - \mathbf{1}). \tag{48}$$

In Sec. VI, we show that this is the cumulant generating function of dynamical fluctuations of $\boldsymbol{\phi}$, so Eq. (48) implies that these fluctuations are invariant under coarse graining. Finally, our variational principle (38) is also invariant; therefore, our excess EPR σ_{ex} and generalized potential $\boldsymbol{\phi}^\star$ are the same, regardless of whether they are defined using $(\boldsymbol{j}, \nabla)$ or $(\overline{\boldsymbol{j}}, \overline{\nabla})$.

This coarse graining has practical implications, as it sometimes allows the generalized potential and excess EPR to be found in closed form. Consider the class of systems that have conservative coarse-grained forces:

$$\overline{\boldsymbol{f}} = \ln(\overline{\boldsymbol{j}}/\overline{\tilde{\boldsymbol{j}}}) = -\overline{\nabla}\boldsymbol{\phi} \tag{49}$$

for some $\boldsymbol{\phi}$. Then, the excess EPR is given by the regular EPR at the level of the coarse-grained fluxes,

$$\sigma_{\mathrm{ex}}(\boldsymbol{j}) = \bar{\sigma}_{\mathrm{ex}}(\overline{\boldsymbol{j}}) = \sigma(\overline{\boldsymbol{j}}), \tag{50}$$

and $\boldsymbol{\phi}^\star$ is equal to $\boldsymbol{\phi}$ from Eq. (49). In such cases, $\boldsymbol{\phi}^\star$ can be found directly by inspecting the linear system of equations (49), eliminating the need for numerical optimization. This technique was implicitly used to find $\boldsymbol{\phi}^\star$ for the two-level MJP in the previous example, as in Eq. (43). The same technique will be used in the Brusselator example in Sec. VIII B.

E. Housekeeping entropy production and the excess/housekeeping decomposition

We now consider the *housekeeping EPR*, the nonconservative part of the dissipation, defined as

$$\sigma_{\mathrm{hk}} := \sigma - \sigma_{\mathrm{ex}}. \tag{51}$$

Since the excess EPR satisfies $0 \leqslant \sigma_{\mathrm{ex}} \leqslant \sigma$, the housekeeping EPR also obeys the bounds $0 \leqslant \sigma_{\mathrm{hk}} \leqslant \sigma$. For a temporally extended process, the time-integrated housekeeping EP is $\Sigma_{\mathrm{hk}} = \int \sigma_{\mathrm{hk}}(t)\, dt$.

The housekeeping EPR has a variational expression:

$$\sigma_{\mathrm{hk}} = \min_{\boldsymbol{\phi}\in\mathbb{R}^d} \mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}} \circ e^{-\nabla\boldsymbol{\phi}}) = \mathcal{D}(\boldsymbol{j}\|\boldsymbol{j}^\star), \tag{52}$$

where the optimal fluxes are $\boldsymbol{j}^\star = \tilde{\boldsymbol{j}} \circ e^{-\nabla\boldsymbol{\phi}^\star}$ as in Eq. (46). This result is derived by plugging Eq. (38) into $\sigma_{\mathrm{hk}} = \mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}}) - \sigma_{\mathrm{ex}}$, expanding the definition of $\mathcal{D}$, and rearranging.

To clarify the meaning of Eq. (52), we introduce the notation

$$\mathscr{D}(\boldsymbol{\theta}\|\boldsymbol{\theta}') := \mathcal{D}\big(\tilde{\boldsymbol{j}} \circ e^{\boldsymbol{\theta}}\|\tilde{\boldsymbol{j}} \circ e^{\boldsymbol{\theta}'}\big) \tag{53}$$

to indicate the relative entropy between pairs of flux vectors in the exponentially tilted parametric family

$$\boldsymbol{\theta} \mapsto \tilde{\boldsymbol{j}} \circ e^{\boldsymbol{\theta}}. \tag{54}$$

Using this notation, the housekeeping EPR may be written as

$$\sigma_{\mathrm{hk}} = \min_{\boldsymbol{\phi}\in\mathbb{R}^d} \mathscr{D}(\boldsymbol{f}\|-\nabla\boldsymbol{\phi}) = \mathscr{D}(\boldsymbol{f}\|-\nabla\boldsymbol{\phi}^\star). \tag{55}$$

In this form, we see that the housekeeping EPR quantifies the information-geometric distance between the actual forces $\boldsymbol{f}$ and the closest conservative forces $-\nabla\boldsymbol{\phi}$ (for some $\boldsymbol{\phi}$). It vanishes when the forces are conservative, and otherwise quantifies their "nonconservative-ness." From this perspective, the generalized potential $\boldsymbol{\phi}^\star$, which achieves the minimum in Eq. (55), provides the "best conservative approximation" of the forces.

Equation (55) implies that housekeeping EPR vanishes whenever $\boldsymbol{f}$ belongs to the image of the stoichiometric matrix ∇. This leads to a simple necessary condition for nonvanishing housekeeping EPR, expressed as an algebraic property of ∇:

$$\sigma_{\mathrm{hk}} > 0 \quad \text{only if} \quad \operatorname{rank}\nabla < m/2, \tag{56}$$

since otherwise it is always possible to express $\boldsymbol{f} = -\nabla\boldsymbol{\phi}$ for some $\boldsymbol{\phi}$. The relevant dimensionality is $m/2$ (rather than m) due to the antisymmetry of $\boldsymbol{f}$ and ∇ [93].

Interestingly, the excess and housekeeping decomposition can be defined at the level of the forces, fluxes, and dissipation of individual reactions. This fine-grained decomposition allows us to classify the contribution of individual reactions to state evolution versus cyclic fluxes. The forces are decomposed into conservative and nonconservative parts as $\boldsymbol{f} = \boldsymbol{f}_{\mathrm{ex}} + \boldsymbol{f}_{\mathrm{hk}}$, where

$$\boldsymbol{f}_{\mathrm{ex}} := -\nabla\boldsymbol{\phi}^\star = \ln\frac{\boldsymbol{j}^\star}{\tilde{\boldsymbol{j}}}, \quad \boldsymbol{f}_{\mathrm{hk}} := \boldsymbol{f} + \nabla\boldsymbol{\phi}^\star = \ln\frac{\boldsymbol{j}}{\boldsymbol{j}^\star}. \tag{57}$$

Similarly, the fluxes are decomposed as $\boldsymbol{j} = \boldsymbol{j}_{\mathrm{ex}} + \boldsymbol{j}_{\mathrm{hk}}$, where

$$\boldsymbol{j}_{\mathrm{ex}} := \boldsymbol{j}^\star, \quad \boldsymbol{j}_{\mathrm{hk}} := \boldsymbol{j} - \boldsymbol{j}^\star. \tag{58}$$

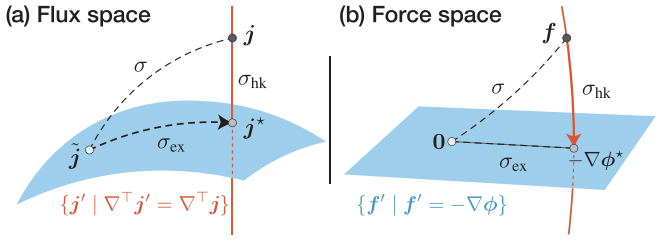

FIG. 3. Information-geometric interpretation of excess/housekeeping decomposition. (a) The EPR $\sigma = \mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}})$ is the relative entropy from the forward fluxes $\boldsymbol{j}$ to the reverse fluxes $\tilde{\boldsymbol{j}}$. Excess EPR $\sigma_{\rm ex}$ is defined by the projection of $\tilde{\boldsymbol{j}}$ onto the set of fluxes that give the actual time evolution (orange line) [Eq. (45)]. Excess/housekeeping provides an orthogonal decomposition of the EPR in flux space [Eq. (60)]. (b) In the space of forces, the EPR $\sigma = \mathscr{D}(\boldsymbol{f}\|\boldsymbol{0})$ is the relative entropy from the actual forces $\boldsymbol{f}$ to the origin $\boldsymbol{0}$. Housekeeping EPR is defined by the (dual) projection of the forces onto the set of conservative forces (blue plane) [Eq. (55)]. Excess/housekeeping provides an orthogonal decomposition of the EPR in force space [Eq. (61)]. The two projections meet at the point corresponding to the generalized potential $\boldsymbol{\phi}^\star$.

We draw attention to two subtle points. First, if one defines the "reverse" excess fluxes as $[\tilde{\boldsymbol{j}}_{\rm ex}]_\rho = [\boldsymbol{j}_{\rm ex}]_{\tilde{\rho}}$, then the excess fluxes have thermodynamic forces $\ln(\boldsymbol{j}_{\rm ex}/\tilde{\boldsymbol{j}}_{\rm ex}) = \boldsymbol{f}_{\rm ex} - \boldsymbol{f}_{\rm hk} = -2\nabla\boldsymbol{\phi}^\star - \boldsymbol{f}$; thus, one should not think of the excess fluxes as being conservative for the potential $\boldsymbol{\phi}^\star$. Second, the housekeeping fluxes $\boldsymbol{j}_{\rm hk}$ should be interpreted as net fluxes (rather than one way), because their entries may be positive or negative. These are cyclic fluxes ($\nabla^\top \boldsymbol{j}_{\rm hk} = \boldsymbol{0}$) that contribute to dissipation without changing the state, and they vanish if and only if the forces are conservative.

Finally, we decompose the dissipation incurred by individual reactions into excess and housekeeping contributions. In the same way that we decomposed total EPR in Eq. (16), we write $\sigma_{\rm ex} = \sum_\rho \sigma_{\rm ex}^{(\rho)}$ and $\sigma_{\rm hk} = \sum_\rho \sigma_{\rm hk}^{(\rho)}$, where the nonnegative contributions from reaction ρ are

$$\sigma_{\rm ex}^{(\rho)} := j_\rho^\star(f_{{\rm ex},\rho} - 1 + e^{-f_{{\rm ex},\rho}}) \geqslant 0,$$
$$\sigma_{\rm hk}^{(\rho)} := j_\rho(f_{{\rm hk},\rho} - 1 + e^{-f_{{\rm hk},\rho}}) \geqslant 0, \tag{59}$$

which may be derived from Eqs. (47) and (52), respectively. $\sigma_{\rm ex}^{(\rho)}$ and $\sigma_{\rm hk}^{(\rho)}$ vanish only when $f_{{\rm ex},\rho} = 0$ and $f_{{\rm hk},\rho} = 0$, respectively. We note that the excess/housekeeping decomposition does not always commute with the reaction decomposition, so in general $\sigma^{(\rho)} \neq \sigma_{\rm ex}^{(\rho)} + \sigma_{\rm hk}^{(\rho)}$.

F. Information-geometric interpretation

Our excess/housekeeping decomposition can be interpreted in terms of information geometry, as shown in Fig. 3. In the space of fluxes, we write our decomposition as

$$\underbrace{\mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}})}_{\sigma} = \underbrace{\mathcal{D}(\boldsymbol{j}^\star\|\tilde{\boldsymbol{j}})}_{\sigma_{\rm ex}} + \underbrace{\mathcal{D}(\boldsymbol{j}\|\boldsymbol{j}^\star)}_{\sigma_{\rm hk}}. \tag{60}$$

This expression is an instance of the Pythagorean relation from information geometry [50,94]. It is analogous to the Pythagorean relation from Euclidean geometry, except that squared Euclidean distance is replaced by relative entropy. Excess and housekeeping provide an orthogonal decomposition of EPR, as illustrated in Fig. 3(a). Orthogonality reflects the fact that the excess EPR (45) is defined by an (information-geometric) projection of $\tilde{\boldsymbol{j}}$ onto a linear subspace [94]: the subspace of fluxes that induce the same time evolution as the actual fluxes, $\nabla^\top \boldsymbol{j}' = \nabla^\top \boldsymbol{j}$.

The same decomposition can be expressed in the dual space of forces [see Fig. 3(b)]. Using notation (53), we write

$$\underbrace{\mathscr{D}(\boldsymbol{f}\|\boldsymbol{0})}_{\sigma} = \underbrace{\mathscr{D}(-\nabla\boldsymbol{\phi}^\star\|\boldsymbol{0})}_{\sigma_{\rm ex}} + \underbrace{\mathscr{D}(\boldsymbol{f}\|-\nabla\boldsymbol{\phi}^\star)}_{\sigma_{\rm hk}}. \tag{61}$$

Here, orthogonality arises because the housekeeping EPR (55) is defined by the dual projection of the forces $\boldsymbol{f}$ onto the flat manifold [50]: the manifold of conservative forces, $\boldsymbol{f}' = -\nabla\boldsymbol{\phi}$, for some $\boldsymbol{\phi}$. These two projections recover the same point, corresponding to the generalized potential $\boldsymbol{\phi}^\star$ [95].

Surprisingly, there has been almost no work on decomposing entropy production using the information-geometric Pythagorean relation at the level of fluxes or trajectories. One exception is Ref. [96], where entropy production was decomposed into partial contributions from individual subsystems. Also, Ref. [51] explored information-geometric decompositions using a different divergence (not relative entropy).

G. Linear response

Here, we consider our variational principle in the linear-response regime $\nabla^\top \boldsymbol{j} \approx \boldsymbol{0}$ around vanishing net production. For closed systems, $\dot{\boldsymbol{x}} = \nabla^\top \boldsymbol{j}$, so this is the regime of slow evolution. Below, we show that our generalized potential obeys Onsager relations, with transport coefficients given by short-time diffusion coefficients. In this sense, our linear-response regime generalizes Onsager theory to far-from-equilibrium and nonconservative systems.

Recall that $\boldsymbol{\phi}^\star = \boldsymbol{0}$ when $\nabla^\top \boldsymbol{j} = \boldsymbol{0}$. To consider the linear-response regime, we expand the variational principle (40) to second order around $\boldsymbol{\phi} \approx \boldsymbol{0}$, giving

$$\sigma_{\rm ex} \approx \max_{\boldsymbol{\phi}\in\mathbb{R}^d}[-2\boldsymbol{\phi}^\top\nabla^\top\boldsymbol{j} - \boldsymbol{\phi}^\top H\boldsymbol{\phi}], \tag{62}$$

where we introduced the matrix

$$H := \tfrac{1}{2}\nabla^\top \mathrm{diag}(\boldsymbol{j})\nabla. \tag{63}$$

For the special case of an MJP with rate matrix R, it has a more explicit expression as

$$H_{ji} = \frac{1}{2}\begin{cases} -p_i R_{ji} - p_j R_{ij}, & i \neq j, \\ \sum_{k(\neq i)}(p_i R_{ki} + p_k R_{ik}), & i = j. \end{cases} \tag{64}$$

Equation (62) can be solved to write the excess EPR as

$$\sigma_{\rm ex} \approx \boldsymbol{j}^\top \nabla H^+ \nabla^\top \boldsymbol{j}, \tag{65}$$

where H^+ is the pseudoinverse.

The optimal observable in Eq. (62) satisfies the linear-response version of Eq. (39) [97],

$$\nabla^\top \boldsymbol{j} \approx -H\boldsymbol{\phi}^\star, \quad \boldsymbol{\phi}^\star \approx -H^+\nabla^\top \boldsymbol{j}. \tag{66}$$

This is a linear Onsager phenomenological relation [98, Chap. 7] that connects the generalized potential $\boldsymbol{\phi}^\star$ to dynamics. The matrices H and H^+ are positive semidefinite

and symmetric; thus, Onsager's reciprocal relations are satisfied. The matrix H specifies the mobility coefficients, while H^+ specifies the friction coefficients. In the next section, we will see that H governs the variance of short-time dynamical fluctuations. The linear-response approximation is valid when net production ($\nabla^\top \boldsymbol{j}$) is small, relative to the scale of dynamical fluctuations (H).

Next, we consider a closed system, where $\dot{\boldsymbol{x}} = \nabla^\top \boldsymbol{j}$. Then, we can express the excess EPR as

$$\sigma_{\mathrm{ex}} \approx \dot{\boldsymbol{x}}^\top H^+ \dot{\boldsymbol{x}}. \tag{67}$$

This defines a Riemannian geometry over thermodynamic states: σ_{ex} acts as the square of the line element and the friction tensor H^+ acts as the Riemannian metric [note that it implicitly depends on the state and time via the fluxes $\boldsymbol{j}(\boldsymbol{x}(t), t)$]. This can be used to define thermodynamic length for slow-evolving nonconservative systems, thereby generalizing the thermodynamic geometry originally developed for conservative systems [16].

Equation (67) does not assume proximity to or existence of the steady state. Nonetheless, we may also consider a closed system that remains close to steady state, $\boldsymbol{x} \approx \boldsymbol{x}^{\mathrm{ss}}$. We can then further approximate the excess EPR as

$$\sigma_{\mathrm{ex}} \approx \dot{\boldsymbol{x}}^{\mathrm{ss}\top} H^+_{\mathrm{ss}} \dot{\boldsymbol{x}}^{\mathrm{ss}}, \tag{68}$$

where H^+_{ss} is the friction matrix evaluated at steady-state fluxes $\boldsymbol{j}(\boldsymbol{x}^{\mathrm{ss}}(t), t)$ and $\dot{\boldsymbol{x}}^{\mathrm{ss}}$ is the change of the steady state due to driving. Equation (68) is quadratic in the speed of driving $\dot{\boldsymbol{x}}^{\mathrm{ss}}$, so the total excess EP incurred over the course of a process, $\Sigma_{\mathrm{ex}}(T) = \int_0^T \sigma_{\mathrm{ex}}(t)\, dt$, vanishes in the limit of slow driving [$T \to \infty$ and $\dot{\boldsymbol{x}}^{\mathrm{ss}} = O(1/T)$]. This can be interpreted as a generalized Clausius equality [39] for our excess EP.

For closed systems, the linear-response analysis highlights the similarity to Onsager's variational principle (OVP) [98–100], also called the "least dissipation principle." In its usual form, OVP says that the evolution $\dot{\boldsymbol{x}}$ is determined by a balance between free energy input and dissipation due to friction. In our notation, this can be expressed as a variational principle for the state evolution, given a fixed generalized potential $\boldsymbol{\phi}^\star$:

$$\sigma_{\mathrm{ex}} \approx \max_{\dot{\boldsymbol{y}} \in \mathrm{im}\nabla^\top} [-2\dot{\boldsymbol{y}}^\top \boldsymbol{\phi}^\star - \dot{\boldsymbol{y}}^\top H^+_{\mathrm{ss}} \dot{\boldsymbol{y}}], \tag{69}$$

where the optimal $\dot{\boldsymbol{y}}$ recovers the system's actual evolution,

$$\dot{\boldsymbol{x}} \approx -H_{\mathrm{ss}} \boldsymbol{\phi}^\star. \tag{70}$$

In our setting, however, we are interested in the "inverse" problem of inferring the generalized potential $\boldsymbol{\phi}^\star$ from a given state evolution $\dot{\boldsymbol{x}}$, giving Eq. (62). For a closed system near steady state, this gives

$$\boldsymbol{\phi}^\star \approx -H^+_{\mathrm{ss}} \dot{\boldsymbol{x}}. \tag{71}$$

This inverse form of the OVP is sometimes called "Gyarmati's variational principle" [100–102].

VI. LARGE DEVIATIONS AND THERMODYNAMIC UNCERTAINTY RELATIONS

Above, we introduced the variational principle (38) and used it to define our generalized potential and excess/housekeeping decomposition. Here, we show that this variational principle has an intuitive physical interpretation in terms of dynamical large deviations. In particular, we show that our excess EPR governs the irreversibility of state dynamics and that the generalized potential can be understood as the "most irreversible" state observable. We also derive a thermodynamic uncertainty relation that allows for practical thermodynamic inference.

A. Dynamical large deviations: Excess EPR

Imagine that one makes n independent measurements of a Markovian stochastic system over a short time interval $[t, t + dt]$. These copies may represent different particles, unit volumes in a chemical reactor, or trial runs of the same experiment [103]. In these measurements, the number of times each reaction occurs will exhibit stochastic fluctuations. For each reaction ρ, let the random variable $J^{(n)}_\rho$ indicate the *empirical flux*, the mean number of reaction events per copy and per unit time. The reactions may correspond to transitions between microstates, as in a general MJP, or to macroscopic chemical reactions, as in a chemical master equation. We write $\boldsymbol{J}^{(n)} = (J^{(n)}_1, \ldots, J^{(n)}_m)$ to indicate the vector of empirical fluxes across all reactions.

The empirical fluxes are said to obey a *large-deviation principle* if, for large n, their probability distribution scales as

$$\Pr[\boldsymbol{J}^{(n)} \approx \boldsymbol{g}] \asymp e^{-n\,\Psi(\boldsymbol{g})\,dt},$$

where $\Psi(\boldsymbol{g})$ is the rate function. This may be termed a many-particle dynamical LDP, since it concerns fluctuations of a dynamical quantity (the fluxes) across independent copies.

In MJPs [104] and ideal CRNs [105], the rate function is known to be the generalized relative entropy $\mathcal{D}$ (13). In particular, the empirical flux LDP is

$$\Pr[\boldsymbol{J}^{(n)} \approx \boldsymbol{g}] \asymp e^{-n\,\mathcal{D}(\boldsymbol{g}\|\boldsymbol{j})\,dt}, \tag{72}$$

where $\boldsymbol{j} = \mathbb{E}[\boldsymbol{J}^{(n)}]$ are the mean fluxes. (See Lazarescu *et al.* [105] for the derivation for well-mixed CRNs with elementary reactions and mass-action kinetics.) In the rest of this section, we assume that the rate function has this form.

EPR has a simple interpretation in terms of statistical fluctuations. In the limit of many copies, the empirical fluxes converge to their expectation values, $\boldsymbol{j}$. At a finite n, however, there is a finite probability that the empirical fluxes move "backward in time" due to a statistical fluctuation, taking the value of the expected *reverse* fluxes $\tilde{\boldsymbol{j}}$. Given Eq. (72), the probability of this time reversal scales as the EPR,

$$\Pr[\boldsymbol{J}^{(n)} \approx \tilde{\boldsymbol{j}}] \asymp e^{-n\,\sigma(\boldsymbol{j})\,dt}, \tag{73}$$

which follows from $\mathcal{D}(\tilde{\boldsymbol{j}}\|\boldsymbol{j}) = \mathcal{D}(\boldsymbol{j}\|\tilde{\boldsymbol{j}}) = \sigma(\boldsymbol{j})$ (assuming no odd variables). Time reversal is exponentially rare, except in equilibrium ($\sigma = 0$) when the system has no direction of time. Equation (73) gives a statistical interpretation of EPR, independent of thermodynamic quantities like heat and work.

Instead of considering time reversal of all fluxes, as in Eq. (73), one may consider the time reversal of net production, i.e., the state change due to reactions. The empirical net production is given by the random variable $\nabla^\top \boldsymbol{J}^{(n)}$, and the probability of its time reversal is $\Pr[\nabla^\top \boldsymbol{J}^{(n)} \approx -\nabla^\top \boldsymbol{j}] \asymp$

$e^{-n\Psi' dt}$. The rate function Ψ' may be computed using the "contraction principle" [106, Sec. 2.3], which says that LDP governing a set of events is determined by the most likely event. In this case, the rate function is expressed via the optimization problem,

$$\Psi' = \min_{\boldsymbol{g}\in\mathbb{R}_+^m} \mathcal{D}(\boldsymbol{g}\|\boldsymbol{j}), \quad \text{where } \nabla^\top \boldsymbol{g} = -\nabla^\top \boldsymbol{j}.$$

Using a change of variables and antisymmetry (3), we see that this is equivalent to the variational expression (45) for excess EPR. Thus, for the empirical net production, the probability of time reversal scales as the excess EPR,

$$\Pr[\nabla^\top \boldsymbol{J}^{(n)} \approx -\nabla^\top \boldsymbol{j}] \asymp e^{-n\sigma_{\rm ex}(\boldsymbol{j})\,dt}. \tag{74}$$

This provides a physically meaningful and operational interpretation of excess EPR in terms of statistical irreversibility.

It is illuminating to compare Eq. (74), the large-deviation expression for excess EPR, and Eq. (28), the large-deviation expression for the steady-state quasipotential (described in Sec. III B). The quasipotential governs occupation fluctuations in stationarity. It is useful for studying steady-state fluctuations and stability, including classic problems such as escape from a stochastic attractor. Our excess EPR, on the other hand, governs the dynamical fluctuations of local-in-time fluxes. As we will see, it is useful for studying the dynamical properties, including thermodynamic uncertainty relations and thermodynamic speed limits. As we will see in the next section, it is also possible to interpret our generalized potential in terms of dynamical fluctuations.

B. Dynamical large deviations: generalized potential

In this section, we interpret our generalized potential $\boldsymbol{\phi}^\star$ in terms of large deviations.

Suppose that one is interested in a fluctuating state observable $\boldsymbol{\phi}$ (e.g., energy, position, etc.) across n independent measurements over time $[t, t+dt]$. The empirical mean change of the observable is captured by the random variable

$$\overline{\Delta\phi}_n = \boldsymbol{\phi}^\top \nabla^\top \boldsymbol{J}^{(n)}\, dt. \tag{75}$$

Specifically, this is the empirical change due to internal reactions. In a closed system without external fluxes, $\overline{\Delta\phi}_n$ is equal to the change over time.

In the limit of many copies, $\overline{\Delta\phi}_n$ converges to $\mathbb{E}[\Delta\phi] = \boldsymbol{\phi}^\top \nabla^\top \boldsymbol{j}\, dt$. As above, we may consider the irreversibility of $\boldsymbol{\phi}$ in terms of the probability that, due to a statistical fluctuation, $\overline{\Delta\phi}_n$ moves "backward in time" relative to the expectation. This probability can be expressed as

$$\Pr[\overline{\Delta\phi}_n \approx -\mathbb{E}[\Delta\phi]] \asymp e^{-n\mathcal{L}(\boldsymbol{\phi})\,dt}. \tag{76}$$

In Fig. 4, we illustrate the meaning of the rate function $\mathcal{L}(\boldsymbol{\phi})\,dt$, which is a fundamental measure of the irreversibility of $\boldsymbol{\phi}$. It is large for "fast" observables that change quickly relative to the scale of their dynamical fluctuations.

$\mathcal{L}(\boldsymbol{\phi})$ can be computed using the contraction principle, as $\min_{\boldsymbol{g}} \mathcal{D}(\boldsymbol{g}\|\boldsymbol{j})$ subject to $\boldsymbol{\phi}^\top \nabla^\top \boldsymbol{g} = -\boldsymbol{\phi}^\top \nabla^\top \boldsymbol{j}$. This optimization can be expressed in its dual form as

$$\mathcal{L}(\boldsymbol{\phi}) = \max_{\lambda\in\mathbb{R}} [\lambda(-\boldsymbol{j}^\top \nabla\boldsymbol{\phi}) - \boldsymbol{j}^\top (e^{\lambda\nabla\boldsymbol{\phi}} - \mathbf{1})]. \tag{77}$$

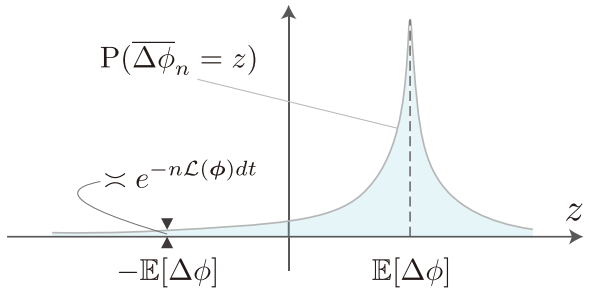

FIG. 4. Large-deviation interpretation of the irreversibility measure $\mathcal{L}(\phi)$. The change of state observable ϕ due to reactions over time $[t, t+dt]$ is measured in n independent copies, and the empirical mean change is captured by the random variable $\overline{\Delta\phi}_n$. Here, we show schematically the probability of different outcomes $\overline{\Delta\phi}_n = z$. For large n, the probability distribution is peaked at its expectation value $z = \mathbb{E}[\Delta\phi]$. The probability that the empirical mean moves in reverse, $z = -\mathbb{E}[\Delta\phi]$, decays exponentially in n as $\asymp e^{-n\mathcal{L}(\phi)\,dt}$. The generalized potential $\phi^\star$ is the "most irreversible" observable, having the largest $\mathcal{L}(\phi)$, and the excess EPR $\sigma_{\rm ex}$ is its degree of irreversibility [see Eq. (78)].

By comparing with our variational principle (38) and with a bit of rearranging, we may write the excess EPR as

$$\sigma_{\rm ex} = \max_{\boldsymbol{\phi}\in\mathbb{R}^d} \mathcal{L}(\boldsymbol{\phi}) = \mathcal{L}(\boldsymbol{\phi}^\star). \tag{78}$$

This shows that the generalized potential $\boldsymbol{\phi}^\star$ is the "most irreversible" observable, and that the excess EPR $\sigma_{\rm ex}$ quantifies its degree of irreversibility.

C. TUR and thermodynamic inference

Here, we discuss our variational principle as a thermodynamic uncertainty relation, which relates fluctuation statistics of observables to excess EPR. This TUR sets a fundamental bound on allowed statistics, limiting the "speed" of state observables relative to the scale of their fluctuations. The TUR also provides a way to perform thermodynamic inference of excess EPR and the generalized potential from empirical measurements.

Before proceeding, we remark that our large-deviation results, such as Eqs. (74), (76), and (78), in theory provide a relationship between excess EPR, the generalized potential, and measurement statistics. In practice, however, such results are often not practical because they involve the statistics of exponentially rare events. Our TUR, on the other hand, does not depend on exponentially rare events.

To introduce the TUR, consider again the change of observable $\boldsymbol{\phi}$ due to reactions over time interval $[t, t+dt]$. We use the random variable $\Delta\phi$ to indicate this change *in a single system* (it should be contrasted with the random variable $\overline{\Delta\phi}_n$ (75), the empirical mean change across n system copies).

The statistics of $\Delta\phi$ are encoded in the cumulant generating function, $\Lambda_{\boldsymbol{\phi}}(\lambda) = \ln \mathbb{E} e^{\lambda\Delta\phi}$. As an example, in an MJP where ϕ_i is the position of site i on a one-dimensional lattice, $\Delta\phi$ is the short-time displacement. The first derivative of the CGF gives the mean displacement, the second derivative gives the mean square displacement, etc.

A key result in large-deviation theory, called Cramér's theorem [106], states that the rate function of $\overline{\Delta\phi}_n$ is the Legendre transform of the CGF of $\Delta\phi$. Given Eq. (77), the rate function $\mathcal{L}(\boldsymbol{\phi})\,dt$ is the Legendre transform of $\Lambda_{\boldsymbol{\phi}}(\lambda) = \boldsymbol{j}^\top(e^{\lambda\nabla\boldsymbol{\phi}} - \mathbf{1})\,dt$, where we assume Poissonian fluctuations and work to first order in dt. The cumulants are explicitly given by $K^{(k)}_{\boldsymbol{\phi}} = \partial^{(k)}_\lambda \Lambda_{\boldsymbol{\phi}}(0) \approx \boldsymbol{j}^\top[\nabla\boldsymbol{\phi}]^k\,dt$. For example, the mean and variance are

$$K^{(1)}_{\boldsymbol{\phi}} = \mathbb{E}[\Delta\phi] = \boldsymbol{\phi}^\top\nabla^\top\boldsymbol{j}\,dt,$$
$$K^{(2)}_{\boldsymbol{\phi}} = \mathrm{Var}(\Delta\phi) = \boldsymbol{\phi}^\top\nabla^\top\mathrm{diag}(\boldsymbol{j})\nabla\boldsymbol{\phi}\,dt.$$

We note that the variance is equal to the mean-squared displacement (MSD) $\mathbb{E}[(\Delta\phi)^2]$ to first order in dt, and that it can be expressed as $K^{(2)}_{\boldsymbol{\phi}} = 2\boldsymbol{\phi}^\top H\boldsymbol{\phi}\,dt$ in terms of our mobility matrix H (63). In this sense, H is a diffusion matrix that governs the short-time MSD of all state observables.

We can also express the variational principle for excess EPR in terms of cumulants. Expanding the exponentials in Eq. (40) gives

$$\sigma_{\mathrm{ex}} = \max_{\boldsymbol{\phi}\in\mathbb{R}^d} \frac{1}{dt}\left[-2K^{(1)}_{\boldsymbol{\phi}} - \sum_{k=2}\frac{1}{k!}K^{(k)}_{\boldsymbol{\phi}}\right]. \tag{79}$$

This provides a set of bounds on the excess EPR, one for each state observable $\boldsymbol{\phi}$. Each bound involves two terms: The first is proportional to the observable's "speed" (the mean displacement $-2K^{(1)}_{\boldsymbol{\phi}}$), while the second quantifies the size of its fluctuations (via a sum of higher cumulants). This fluctuation term is non-negative, as discussed near Eq. (40). Note that the sum of higher cumulants converges, so in practice it may be truncated at some finite order.

Equation (79) can be interpreted as a short-time TUR [107] for excess EPR that specifies a trade-off between speed and fluctuations. It is also an example of a "higher-order" TUR [108], since it involves not only the mean and variance but also higher cumulants. Higher-order TURs can be made tight by appropriate choice of observable, unlike traditional TURs (involving only the mean and variance) that cannot always be made tight in far-from-equilibrium discrete systems [109]. In SM3 in the Supplemental Material [53], we derive a weaker but simpler bound that involves simpler statistics.

The TUR (79) permits thermodynamic inference of σ_{ex} and $\boldsymbol{\phi}^\star$, using similar techniques as for inference of EPR and thermodynamic forces from short-time trajectories [87,88,109,110]. Suppose that, in a closed system, one makes n two-point measurements of the change of a state observable during a short time interval, $\phi(t) \mapsto \phi(t+dt)$. These measurements can be used to compute the cumulants $K^{(k)}_{\boldsymbol{\phi}}$, which give a lower bound on the excess EPR when plugged into the TUR (79). The TUR can be further tightened by scaling the observable as $\lambda\boldsymbol{\phi}$ by $\lambda \in \mathbb{R}$ and optimizing over λ, in this way recovering the overall irreversibility $\mathcal{L}(\boldsymbol{\phi})$ (77).

If it is possible to measure more than one observable $(\boldsymbol{\phi}_{(1)}, \boldsymbol{\phi}_{(2)}, \ldots)$, tighter bounds can be found by numerical optimization of the TUR (79) over linear combinations of these observables. In the fine-grained regime where d linearly independent observables can be acquired (e.g., the identity of starting and ending microstates), then the measurements span the entire space of observables $\boldsymbol{\phi} \in \mathbb{R}^d$ and numerical optimization should recover the full excess EPR and the generalized potential $\boldsymbol{\phi}^\star$ at optimality. We note that the objective in TUR (79) is a concave function of $\boldsymbol{\phi}$; thus, optimization can be carried out using very efficient numerical algorithms. We also note that, unlike some EPR estimation techniques, our method never needs to distinguish which reaction or reservoir mediated a given transition.

VII. THERMODYNAMIC SPEED LIMITS

In this section, we derive thermodynamic speed limits. We first introduce a short-time TSL, which relates excess EPR to the instantaneous speed of evolution at a given time. We then introduce a finite-time TSL, which relates integrated excess EP to time and trajectory length. Because excess EPR is always smaller than EPR, our TSLs also imply bounds on the total entropy production.

A. Short-time speed limits

We first derive a "short-time" TSL that relates excess EPR to two simple quantities. The first quantity is the *activity*, the total number of reactions per unit time,

$$A := \sum_\rho j_\rho = \|\boldsymbol{j}\|_1. \tag{80}$$

The second quantity is the net production due to reactions, $\nabla^\top\boldsymbol{j}$. In a closed system without external fluxes, this is equal to the state evolution, $\nabla^\top\boldsymbol{j} = \dot{\boldsymbol{x}}$. For an open system with external fluxes, it includes both the state evolution and the external fluxes, $\nabla^\top\boldsymbol{j} = \dot{\boldsymbol{x}} + \boldsymbol{I}$. Finally, for an open system in steady state, it is equal to the external fluxes, $\nabla^\top\boldsymbol{j} = \boldsymbol{I}$.

Our TSL quantifies the minimal excess EPR compatible with a given net production $\nabla^\top\boldsymbol{j}$ and activity A,

$$\sigma_{\mathrm{TSL}} := \min_{\boldsymbol{g}\in\mathbb{R}^m_+} \sigma_{\mathrm{ex}}(\boldsymbol{g}), \quad \text{where } \nabla^\top\boldsymbol{g} = \nabla^\top\boldsymbol{j},\ \|\boldsymbol{g}\|_1 = A. \tag{81}$$

σ_{TSL} is a function of $\nabla^\top\boldsymbol{j}$ and A, though we usually leave this implicit in our notation. Also note that without the activity constraint, the minimum would be zero, since one can increase forward and reverse fluxes while keeping their difference fixed. In SM4 in the Supplemental Material [53], we derive the form of the optimal fluxes in Eq. (81) and show that they are conservative.

In that Supplemental Material [53], we also show that our TSL is closely related to the *(1-)Wasserstein speed* from optimal transport [21,111,112]. The Wasserstein speed quantifies the minimum activity required to achieve net production $\nabla^\top\boldsymbol{j}$ in a system with stoichiometric matrix ∇,

$$\dot{W} := \min_{\boldsymbol{a}\in\mathbb{R}^m_+} \|\boldsymbol{a}\|_1, \quad \text{where } \nabla^\top\boldsymbol{a} = \nabla^\top\boldsymbol{j}. \tag{82}$$

Importantly, $\dot{W}$ is sensitive to the system's network topology, as encoded in the matrix ∇. For example, the same net production requires a faster Wasserstein speed on a one-dimensional chain than on a fully connected network. In the Supplemental Material [53], we show that $\dot{W}$ may be efficiently computed using its dual form.

Our TSL can be expressed in terms of $\dot{W}$ as

$$\sigma_{\mathrm{TSL}} = 2\dot{W}\,\tanh^{-1}\frac{\dot{W}}{A}. \tag{83}$$

The right side depends on the ratio between the minimum required activity ($\dot{W}$) and the actual activity (A). The bound diverges as $\dot{W} \to A$, which corresponds to the limit where all activity is channeled into production. We note that this limit can only be achieved by making reactions absolutely irreversible ($j_\rho > 0$, $\tilde{j}_\rho = 0$), since reversibility increases activity without contributing to net production.

To summarize, we have the following hierarchy of bounds:

$$\sigma \geqslant \sigma_{\mathrm{ex}} \geqslant \sigma_{\mathrm{TSL}} \geqslant \frac{2\dot{W}^2}{A}. \tag{84}$$

The last quadratic bound follows from Eq. (83) and $x\tanh^{-1}x \geqslant x^2$. It is tight in the slow-production limit ($\dot{W} \ll A$), but it is much weaker than Eq. (83) for fast production, since it does not diverge when $\dot{W} \to A$. In principle, all of these bounds in Eq. (84) are achievable.

By inverting Eq. (83), we may write our TSL in a more familiar "speed limit" form, as an upper bound on the speed as a function of excess EPR and activity:

$$\dot{W} = A\,\Phi^{-1}(\sigma_{\mathrm{TSL}}/2A) \leqslant A\,\Phi^{-1}(\sigma_{\mathrm{ex}}/2A), \tag{85}$$

where Φ^{-1} is the inverse function of $\Phi(x) := x\tanh^{-1}x$.

Finally, we remark on the case of steady-state open systems, since these play an important role in many biological and chemical studies. For an open system in steady state, net production balances the external fluxes, $\nabla^\top \boldsymbol{j} = \boldsymbol{I}$, and our TSL (81) identifies the steady-state fluxes that minimize excess EPR under an activity constraint. This has some similarity to "flux-balance analysis" (FBA) [113], a biological modeling technique that identifies steady-state metabolic fluxes that optimize some (typically linear) objective under constraints. (See Refs. [114–116] for connections between FBA and nonequilibrium thermodynamics.) From a practical perspective, our TSL bounds the dissipation using three accessible pieces of information: the stoichiometry ∇, the external fluxes $\boldsymbol{I}$, and the overall activity A. We demonstrate the utility of this bound when we analyze metabolic networks in Sec. VIII C.

B. Relation to other TSLs

Previous work [19,21,112] has considered the analog of σ_{TSL} but for total EPR,

$$\sigma'_{\mathrm{TSL}} := \min_{\boldsymbol{g}\in\mathbb{R}^m_+} \sigma(\boldsymbol{g}), \quad \text{where } \nabla^\top \boldsymbol{g} = \nabla^\top \boldsymbol{j}, \quad \|\boldsymbol{g}\|_1 = A. \tag{86}$$

However, as mentioned above, the optimal fluxes in Eq. (81) are conservative, which implies that the two TSLs are equivalent: $\sigma'_{\mathrm{TSL}} = \sigma_{\mathrm{TSL}}$. This also shows that the activity TSL (86) only recovers the excess component of total EPR. We note that the Wasserstein expression (83) for σ'_{TSL} was derived for MJPs in Refs. [19,21] and CRNs in Ref. [112].

For MJPs, the Wasserstein speed can be bounded as $\dot{W} \geqslant \|\dot{\boldsymbol{p}}\|_1/2$. The quantity $\|\dot{\boldsymbol{p}}\|_1/2$, called the *total variation (TV) speed*, is a simple and optimization-free quantity that is not sensitive to system topology [111]. Our results imply the bounds

$$\sigma_{\mathrm{ex}} \geqslant \|\dot{\boldsymbol{p}}\|_1\,\tanh^{-1}\frac{\|\dot{\boldsymbol{p}}\|_1}{2A} \geqslant \frac{\|\dot{\boldsymbol{p}}\|_1^2}{2A}. \tag{87}$$

The last bound recalls a quadratic TSL for HS excess EPR [9],

$$\sigma_{\mathrm{ex}}^{\mathrm{HS}} \geqslant \frac{\|\dot{\boldsymbol{p}}\|_1^2}{2A}. \tag{88}$$

HS excess EPR does not obey the stronger TV bound, $\sigma_{\mathrm{ex}}^{\mathrm{HS}} \not\geqslant \|\dot{\boldsymbol{p}}\|_1\,\tanh^{-1}(\|\dot{\boldsymbol{p}}\|_1/2A)$, although some intermediate inequalities have been shown [12,117]. It also does not obey the analog of the Wasserstein TSL (83), as will be seen in the unicyclic system considered in Sec. VIII [118].

C. Finite-time speed limits

In this section, we derive "finite-time" TSLs for the excess EP incurred by a time-extended process.

Recall that the excess EP incurred during time $t \in [0, T]$ is given by $\Sigma_{\mathrm{ex}} = \int_0^T \sigma_{\mathrm{ex}}(t)\,dt$. We define our finite-time TSL as the minimal excess EP required by any trajectory with time-dependent net production $\nabla^\top \boldsymbol{j}(t)$ and activity $A(t) = \|\boldsymbol{j}(t)\|_1$. Formally, we optimize over all time-dependent fluxes compatible with these constraints:

$$\Sigma_{\mathrm{TSL}} := \min_{\boldsymbol{g}(t)} \int_0^T \sigma_{\mathrm{ex}}(\boldsymbol{g}(t))\,dt,$$
$$\text{where} \quad \nabla^\top \boldsymbol{g}(t) = \nabla^\top \boldsymbol{j}(t), \quad \|\boldsymbol{g}(t)\|_1 = A(t)\ \forall t, \tag{89}$$

where $\sigma_{\mathrm{ex}}(\boldsymbol{g}(t))$ refers to the excess EPR (38) evaluated at fluxes $\boldsymbol{j} \equiv \boldsymbol{g}(t)$. The minimization decouples over different time points, allowing us to express it as an integral over the short-time TSL:

$$\Sigma_{\mathrm{TSL}} = \int_0^T \sigma_{\mathrm{TSL}}(t)\,dt, \tag{90}$$

where $\sigma_{\mathrm{TSL}}(t)$ is the short-time TSL (81) corresponding to the net production $\nabla^\top \boldsymbol{j}(t)$ and activity $A(t)$ at time t. As mentioned, σ_{TSL} can be equivalently defined as a minimization of total EPR (rather than excess EPR); thus, Eq. (89) can also be equivalently defined as a minimization of total EP.

Using Eq. (83), we may relate our TSL to the Wasserstein distance from optimal transport:

$$\Sigma_{\mathrm{TSL}} = \int_0^T 2\dot{W}(t)\,\tanh^{-1}\frac{\dot{W}(t)}{A(t)}\,dt, \tag{91}$$

where $\dot{W}(t)$ is the Wasserstein speed (82) corresponding to net production $\nabla^\top \boldsymbol{j}(t)$. The bound $\Sigma_{\mathrm{ex}} \geqslant \Sigma_{\mathrm{TSL}}$ is achieved when the fluxes are chosen to be optimal at each time point $t \in [0, T]$, as described in SM4 in the Supplemental Material [53].

We can derive another TSL for the time-integrated activity and Wasserstein length:

$$\mathcal{A} := \int_0^T A(t)\,dt, \quad \mathscr{L} := \int_0^T \dot{W}(t)\,dt. \tag{92}$$

Using this definition, we may lower bound Σ_{TSL} as

$$\Sigma_{\mathrm{TSL}} \geqslant \Sigma_{\mathrm{TSL}}^{\mathscr{L}} := 2\mathscr{L}\,\tanh^{-1}\frac{\mathscr{L}}{\mathcal{A}}. \tag{93}$$

To derive this, we multiply and divide the term inside the integral in Eq. (91) by $A(t)/\mathcal{A}$. We then apply Jensen's inequality to the convex function $x \tanh^{-1} x$ with $x = \dot{W}/A$.

From a geometric perspective, the Wasserstein length $\mathscr{L}$ is the length of a given trajectory through concentration space. The Wasserstein length also has an operational interpretation: It is the minimal integrated activity (total number of reaction events) needed to implement the trajectory with a given stoichiometric matrix.

The above inequalities require detailed information about the system's trajectory. Our final result is another TSL that does not depend on such detailed information. To derive it, we introduce the integrated net production:

$$\boldsymbol{V} := \int_0^T \nabla^\top \boldsymbol{j}(t)\, dt. \tag{94}$$

For a closed system, the integrated net production is the change of the state, $\boldsymbol{V} = \boldsymbol{x}(T) - \boldsymbol{x}(0)$. For an open system, it accounts for the integrated outflow as well as the change of state, $\boldsymbol{V} = \boldsymbol{x}(T) - \boldsymbol{x}(0) + \int_0^T \boldsymbol{I}(t)\, dt$. Finally, for an open system in steady state, or (more generally) a cyclic process $\boldsymbol{x}(T) = \boldsymbol{x}(0)$, the integrated net production is the integrated outflow: $\boldsymbol{V} = \int_0^T \boldsymbol{I}(t)\, dt$. In all cases, $\boldsymbol{V}$ depends only on the initial and final states and the integrated flows.

The Wasserstein distance associated with integrated net production is the integrated-time analog of Eq. (82):

$$\mathscr{W} := \min_{\boldsymbol{a} \in \mathbb{R}_+^m} \|\boldsymbol{a}\|_1, \quad \text{where } \nabla^\top \boldsymbol{a} = \boldsymbol{V}. \tag{95}$$

For a closed system, $\mathscr{W}$ is the geodesic distance in Wasserstein space between $\boldsymbol{x}(0)$ and $\boldsymbol{x}(T)$. More generally, using the triangle inequality, we may show $\mathscr{W} \leqslant \mathscr{L}$, with equality for constant-speed trajectories. This leads to our final TSL:

$$\Sigma_{\mathrm{TSL}}^{\mathscr{L}} \geqslant \Sigma_{\mathrm{TSL}}^{\mathscr{W}} := 2\mathscr{W} \tanh^{-1} \frac{\mathscr{W}}{\mathcal{A}}. \tag{96}$$

Summarizing, we have the sequence of bounds

$$\Sigma \geqslant \Sigma_{\mathrm{ex}} \geqslant \Sigma_{\mathrm{TSL}} \geqslant \Sigma_{\mathrm{TSL}}^{\mathscr{L}} \geqslant \Sigma_{\mathrm{TSL}}^{\mathscr{W}}. \tag{97}$$

We may also derive a more familiar form of the TSL, as a bound on the minimal time required to traverse a given distance. Consider the time-averaged dynamical activity, defined as the number of reactions per time, $\langle A \rangle := \mathcal{A}/T$. Rearranging the results above gives

$$T \geqslant \frac{\mathscr{L}}{\langle A \rangle} \coth \frac{\Sigma_{\mathrm{ex}}}{2\mathscr{L}} \geqslant \frac{\mathscr{W}}{\langle A \rangle} \coth \frac{\Sigma_{\mathrm{ex}}}{2\mathscr{W}}. \tag{98}$$

Since $\coth(x) \geqslant 1$, there is a finite minimal time needed to undergo a trajectory of a given length,

$$T \geqslant \mathscr{L}/\langle A \rangle \geqslant \mathscr{W}/\langle A \rangle.$$

This latter bound is relevant in the highly irreversible limit where Σ_{ex} diverges. Conversely, these TSLs imply that Σ_{ex} diverges as $-\ln(T - \mathscr{L}/\langle A \rangle)$ as $T \to \mathscr{L}/\langle A \rangle$. This is stronger than the $1/T$ scaling reported in conventional TSLs [7,10,17,20,22,119,120], which is only tight in the limit of slow production [121–123].

D. Information-theoretic speed limit for relaxing MJPs

Our final result is a finite-time TSL that relates excess EP and the speed of relaxation. This result is different from the TSLs derived above: It only applies to MJPs, it only applies to time-symmetric processes, and it does not reference Wasserstein distance. Nonetheless, it shows the connection between our approach and previous literature [90], and it provides a useful bound based on an interpretable and empirically accessible information-theoretic notion of distance.

Consider an MJP that undergoes a process over $t \in [0, T]$, during which the system's probability distribution goes from $\boldsymbol{p}(0)$ to $\boldsymbol{p}(T)$. Suppose that the process involves an autonomous relaxation (time-independent control parameters), or more generally that the driving is time-symmetric, such that the control parameters at time t are the same as at time $T - t$. Given these assumptions, we show in SM2 in the Supplemental Material [53] that the variational principle for excess EPR (38) can be expressed as

$$\sigma_{\mathrm{ex}} = \max_{\boldsymbol{q}} \left[-\frac{d}{dt} D(\boldsymbol{p}(t) \| \boldsymbol{q}(-t)) \right], \tag{99}$$

where the maximization is over all probability distributions over the microstates. The notation $\boldsymbol{q}(-t)$ indicates that $\boldsymbol{q}$ evolves "backward in time,"

$$-\frac{d}{dt} q_i(-t) = \sum_{j(\neq i),\alpha} \left(q_j(-t) R_{ij}^\alpha - q_i(-t) R_{ji}^\alpha \right). \tag{100}$$

Thus, σ_{ex} is the fastest rate of contraction of relative entropy between the actual distribution $\boldsymbol{p}$ evolving forward in time and any other distribution evolving backward in time. The optimizer in Eq. (99) is the pseudo-canonical distribution $\boldsymbol{p}^\star \propto \boldsymbol{p} \circ e^{-\boldsymbol{\phi}^\star}$ defined by the generalized potential $\boldsymbol{\phi}^\star$.

Equation (99) gives the following information-theoretic bound:

$$\Sigma_{\mathrm{ex}}(T) \geqslant D(\boldsymbol{p}(0) \| \boldsymbol{p}(T)). \tag{101}$$

The relative entropy $D(\boldsymbol{p}(0) \| \boldsymbol{p}(T))$ is an information-theoretic measure of the distance traversed by the system's state over time $t \in [0, T]$. It is interpretable and practically accessible by measuring the state at two time points.

Equations (99) and (101) generalize the main results of Ref. [90], which derived the same expressions for EPR in conservative MJPs. In particular, the derivation of Eq. (101) proceeds in the same way as Eq. (3) in Ref. [90]. Specifically, we choose $\boldsymbol{q}(0) = \boldsymbol{p}(T)$ in Eq. (99) and then integrate over $t \in [0, T/2]$. Under the assumption of time-symmetric driving, $\boldsymbol{q}(t) = \boldsymbol{p}(T - t)$ is a solution to Eq. (100), so the time integral gives

$$\begin{aligned} \Sigma_{\mathrm{ex}}(T/2) &\geqslant -\int_0^{T/2} \frac{d}{dt} D[\boldsymbol{p}(t) \| \boldsymbol{p}(T - t)]\, dt \\ &= D[\boldsymbol{p}(0) \| \boldsymbol{p}(T)] - D[\boldsymbol{p}(T/2) \| \boldsymbol{p}(T/2)] \\ &= D[\boldsymbol{p}(0) \| \boldsymbol{p}(T)]. \end{aligned}$$

Since $\sigma_{\mathrm{ex}}(t) \geqslant 0$ for $t \in [T/2, T]$, we have the sequence of bounds:

$$\Sigma_{\mathrm{ex}}(T) \geqslant \Sigma_{\mathrm{ex}}(T/2) \geqslant D(\boldsymbol{p}(0) \| \boldsymbol{p}(T)].$$

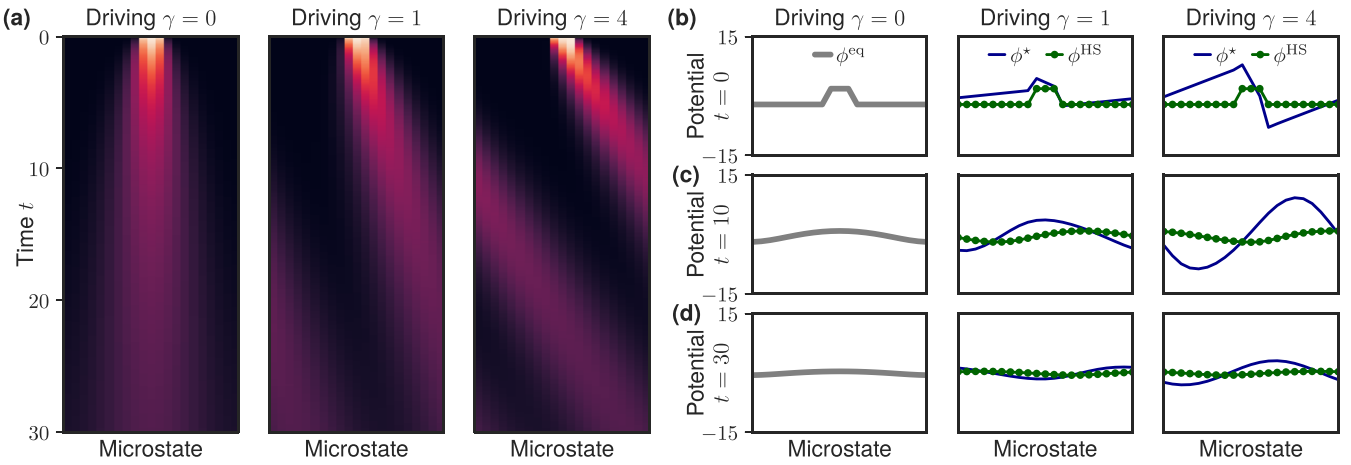

FIG. 5. Time evolution and generalized potentials for unicyclic MJP. (a) Time evolution of the probability distribution for three driving strengths ($\gamma = 0, 1, 4$). The system has 21 microstates, and the initial distribution is concentrated on three microstates $i \in \{10, 11, 12\}$. (b)–(d) Our generalized potential $\phi^{\star}$ (blue) and steady-state potential ϕ^{ss} (dotted green) for different driving strengths (columns) evaluated at three timepoints: $t = 0$ (b), $t = 10$ (c), and $t = 30$ (d). Both potentials vanish as the system approaches the uniform steady state.

A bound like Eq. (101) was conjectured for HS excess EP in Ref. [124]. However, that conjecture does not hold, as we will see in the unicyclic system considered in the next section.

VIII. EXAMPLES

We now illustrate our approach on three examples: a unicyclic MJP, a nonlinear CRN (Brusselator), and real-world metabolic networks.

A. Unicyclic MJP

In our first example, we consider a uniform cyclic MJP. This example will be useful to illustrate the difference between our decomposition and the HS one.

Our MJP has $d = 21$ microstates, and its transition rates are parametrized as

$$R_{i+1,i} = \frac{1}{1+e^{-\gamma}}, \quad R_{i-1,i} = \frac{e^{-\gamma}}{1+e^{-\gamma}}, \tag{102}$$

where the indexing of microstates i is taken mod d. The parameter γ determines the strength of nonconservative driving around the cycle. The overall timescale is normalized, so that the escape rates do not depend on γ ($-R_{ii} = 1$ for all i). Therefore, the dynamical activity is $A = 1$, regardless of γ or the distribution $\boldsymbol{p}$. As the initial probability distribution $\boldsymbol{p}(0)$, we assign elevated probability to three microstates, $p_{10}(0) = p_{11}(0) = p_{12}(0) = 0.3$, with the remaining probability split among the other 18 microstates, $p_i(0) = 0.1/18$.

Figure 5(a) shows the time evolution of the system's state for three values of the driving strength: $\gamma = 0$, $\gamma = 1$, and $\gamma = 4$. For $\gamma = 0$, the system is conservative and it relaxes to equilibrium by diffusing symmetrically. For $\gamma = 1$ and $\gamma = 4$, clockwise/counterclockwise symmetry is broken and the system exhibits a decaying oscillation around the cycle in the direction $i \to i + 1$.

Figure 5(b) shows the steady-state potential $\boldsymbol{\phi}^{\mathrm{ss}} = \ln(\boldsymbol{p}/\boldsymbol{p}^{\mathrm{ss}})$ and our generalized potential $\boldsymbol{\phi}^{\star}$ (found by numerical optimization) at the initial state $\boldsymbol{p}(0)$. For the conservative system with $\gamma = 0$, the potentials are equal to each other and to the standard nonequilibrium potential, $\boldsymbol{\phi}^{\star} = \boldsymbol{\phi}^{\mathrm{ss}} = \boldsymbol{\phi}^{\mathrm{eq}} = \ln(\boldsymbol{p}/\boldsymbol{p}^{\mathrm{eq}})$. For stronger driving, $\boldsymbol{\phi}^{\star}$ becomes larger in magnitude, more asymmetric, and increasingly "clifflike," reflecting the system's increasingly fast and asymmetric relaxation dynamics. On the other hand, because the initial state is symmetric and the steady state is uniform for all γ, the steady-state potential $\boldsymbol{\phi}^{\mathrm{ss}}$ is always symmetric and does not vary with γ.

Figures 5(c) and 5(d) show the steady-state potential $\boldsymbol{\phi}^{\mathrm{ss}}$ and our generalized potential $\boldsymbol{\phi}^{\star}$ for the same three driving strengths, but now evaluated on the time-evolved states $\boldsymbol{p}(t)$ at $t = 10$ and $t = 30$. Both potentials become smoother over time, gradually approaching the zero vector as the system approaches the uniform steady-state distribution. Our generalized potential maintains the asymmetry that reflects the asymmetry of relaxation dynamics, and its magnitude reflects the strength of driving. On the other hand, the steady-state potential remains essentially symmetric about its maximum, similar to the time-evolved distribution $\boldsymbol{p}(t)$. It has some dependence on γ due to the fact that different γ lead to different time-evolved states $\boldsymbol{p}(t)$, but this dependence is weak.

Figure 6 illustrates our thermodynamic speed limit. Figure 6(a) shows the Wasserstein speed $\dot{W}$ (82) as a function of the driving strength, evaluated at the initial state $\boldsymbol{p}(0)$ and the time-evolved state $\boldsymbol{p}(t)$ at $t = 10$ and $t = 30$. At all times, the Wasserstein speed increases with driving strength. Thus, the state evolves faster when it undergoes a directed force around the cycle, compared to when it undergoes an undriven symmetric diffusion.

Figure 6(b) illustrates our short-time Wasserstein TSL on the initial state $\boldsymbol{p}(0)$. This shows the EPR, our excess EPR, the HS excess EPR, and the TSL lower bound $\sigma_{\mathrm{TSL}} = 2\dot{W}\tanh^{-1}(\dot{W}/A)$ (83) as a function of the driving strength γ. The EPR, our excess EPR, and the Wasserstein TSL increase with stronger driving. However, the HS excess EPR remains constant, being insensitive to the driving strength. We verify the inequality $\sigma_{\mathrm{ex}} \geqslant \sigma_{\mathrm{ex}}^{\mathrm{HS}}$ from Eq. (116). We also verify that the Wasserstein TSL holds for our excess EPR, but not the

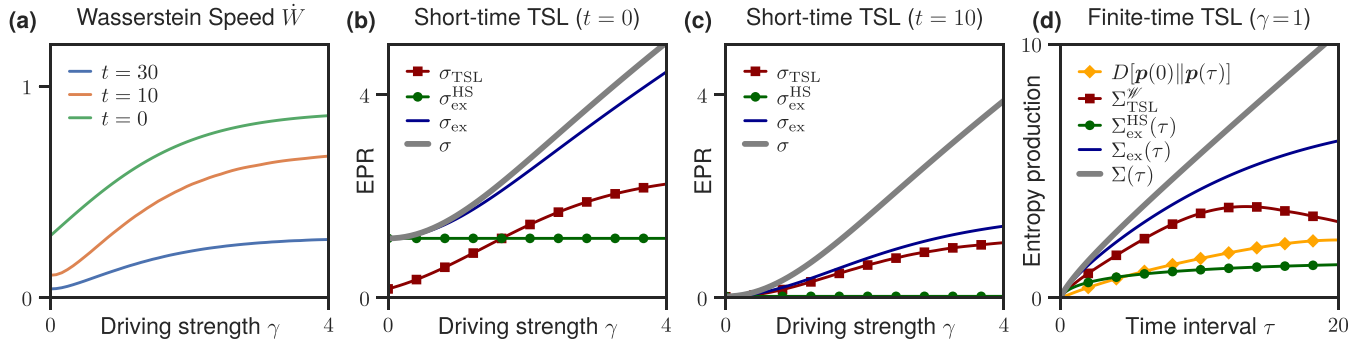

FIG. 6. Wasserstein TSL on unicyclic MJP. (a) Wasserstein speed increases with driving strength when evaluated on state $\boldsymbol{p}(t)$ at $t \in \{0, 10, 30\}$. (b) EPR, our excess EPR and HS excess EPR for the initial state $\boldsymbol{p}(0)$, as a function of the driving strength. The Wasserstein short-time TSL $\sigma_{\rm TSL}$ (83) provides a lower bound on our excess EPR. The HS excess EPR is not sensitive to the driving strength γ. (c) Same as in panel (b), but now shown for the time-evolved state $\boldsymbol{p}(t)$ at $t = 10$. (d) Time-integrated EP, excess EP, and HS excess EP as a function of the time interval, considering the spatiotemporal trajectory shown in Fig. 5(a) for $\gamma = 1$. We also show the finite-time Wasserstein TSL $\Sigma_{\rm TSL}^{\mathscr{W}}$ (96) and information-theoretic bound (101). Observe that the Wasserstein and information-theoretic bounds only hold for our excess EP, not the HS excess EP.

HS excess EPR. Figure 6(c) shows the same quantities, but now evaluated on the time-evolved state $\boldsymbol{p}(t)$ at $t = 10$. Observe that by time $t = 10$, HS excess EPR has almost entirely vanished. Our excess EPR and TSL bound also decrease over time, but not as quickly. The short-time Wasserstein TSL is nearly tight at $t = 10$ across a range of driving strengths.

To summarize, Figs. 5(b) and 6(b) show that our generalized potential and excess EPR are sensitive to the direction and speed of evolution of the system's state. On the other hand, in this system, the steady-state potential and HS excess EPR are not sensitive to the strength of driving or speed of evolution.

Finally, Fig. 6(d) illustrates our finite-time TSLs. Given the spatiotemporal trajectory from Fig. 5(a) for $\gamma = 1$, we plot the time-integrated EP, $\Sigma(\tau) = \int_0^\tau \sigma(t)\,dt$, as a function of the time interval τ, and similarly for $\Sigma_{\rm ex}(\tau)$ and $\Sigma_{\rm ex}^{\rm HS}(\tau)$. We verify the bounds $\Sigma(\tau) \geqslant \Sigma_{\rm ex}(\tau) \geqslant \Sigma_{\rm ex}^{\rm HS}(\tau)$, the information-theoretic bound $\Sigma_{\rm ex}(\tau) \geqslant D[\boldsymbol{p}(0)\|\boldsymbol{p}(\tau)]$, and the Wasserstein TSL $\Sigma_{\rm TSL}^{\mathscr{W}}$ (96). The Wasserstein TSL captures most of the excess EPR until $\tau \approx 10$, at which point it begins to decrease. We observe that the Wasserstein and information-theoretic bounds do not hold for HS excess EPR.

B. Brusselator CRN

In our second example, we consider the Brusselator [58], a well-known CRN model of a chemical oscillator, shown above in Fig. 2.

To analyze this system, we use the fact that two of the reactions, $X_1 \rightleftharpoons X_2$ and $3X_1 \rightleftharpoons 2X_1 + X_2$, have the same stoichiometry. We then introduce the reaction-level coarse graining described in Sec. V D, giving the following coarse-grained stoichiometric matrix and forward/reverse fluxes:

$$\overline{\nabla} = \begin{bmatrix} 1 & 0 \\ -1 & 0 \\ -1 & 1 \\ 1 & -1 \end{bmatrix}, \quad \bar{\boldsymbol{j}} = \begin{bmatrix} k_1^+ \\ c_1 k_1^- \\ c_1\left(k_2^+ + c_1^2 k_3^-\right) \\ c_2\left(k_2^- + c_1^2 k_3^+\right) \end{bmatrix},$$

$$\bar{\bar{\boldsymbol{j}}} = \begin{bmatrix} c_1 k_1^- \\ k_1^+ \\ c_2\left(k_2^- + c_1^2 k_3^+\right) \\ c_1\left(k_2^+ + c_1^2 k_3^-\right) \end{bmatrix}.$$

The resulting coarse-grained forces are conservative,

$$\bar{\boldsymbol{f}} = \left(\ln \frac{\bar{\bar{j}}_1}{\bar{\bar{j}}_2}, \ \ln \frac{\bar{\bar{j}}_2}{\bar{\bar{j}}_1}, \ \ln \frac{\bar{\bar{j}}_3}{\bar{\bar{j}}_4}, \ \ln \frac{\bar{\bar{j}}_4}{\bar{\bar{j}}_3}\right) = -\overline{\nabla}\boldsymbol{\phi}^\star,$$

where the generalized potential is

$$\boldsymbol{\phi}^\star = \left(\ln \frac{\bar{\bar{j}}_2}{\bar{\bar{j}}_1}, \ \ln \frac{\bar{\bar{j}}_2}{\bar{\bar{j}}_1} + \ln \frac{\bar{\bar{j}}_4}{\bar{\bar{j}}_3}\right), \tag{103}$$

$$= \left(\ln \frac{c_1 k_1^-}{k_1^+}, \ \ln \frac{c_1 k_1^-}{k_1^+} + \ln \frac{c_2\left(k_2^- + c_1^2 k_3^+\right)}{c_1\left(k_2^+ + c_1^2 k_3^-\right)}\right). \tag{104}$$

Since the coarse-grained forces are conservative, the excess EPR is the EPR of the coarse-grained fluxes (see Sec. V D):

$$\sigma_{\rm ex} = \sigma(\bar{\boldsymbol{j}}). \tag{105}$$

The housekeeping EPR is the remainder, $\sigma_{\rm hk} = \sigma(\boldsymbol{j}) - \sigma(\bar{\boldsymbol{j}})$.

The excess/housekeeping decomposition can also be performed at the level of individual reactions, as described above in Sec. V E. For the Brusselator, the reactions $\varnothing \rightleftharpoons X_1$ are not coarse grained, and one can verify that their forces obey $f_\rho = -[\nabla\boldsymbol{\phi}^\star]_\rho$. Thus, according to Eqs. (57) and (58), their fluxes and forces are purely excess, with no housekeeping contribution. On the other hand, the reactions $X_1 \rightleftharpoons X_2$ and $3X_1 \rightleftharpoons 2X_1 + X_2$ are coarse grained, so they have $f_\rho \neq -[\nabla\boldsymbol{\phi}^\star]_\rho$ and exhibit nonvanishing housekeeping forces and fluxes. Their housekeeping fluxes quantify the flow around the futile cycle $X_1 \rightharpoonup X_2$, followed by $2X_1 + X_2 \rightharpoonup 3X_1$, which dissipates EP but has no net effect on species counts.

We now provide a concrete numerical example. We consider rate constants $k_1^+ = k_1^- = k_2^- = k_3^- = 1, k_2^+ = 17$, while varying k_3^+ in the range $k_3^+ \in [6, 11]$. For these parameter values, the system exhibits limit-cycle oscillations. Time-dependent concentrations $c_1(t), c_2(t)$ for three different choices of k_3^+ are shown in Figs. 7(b)–7(d). In addition,

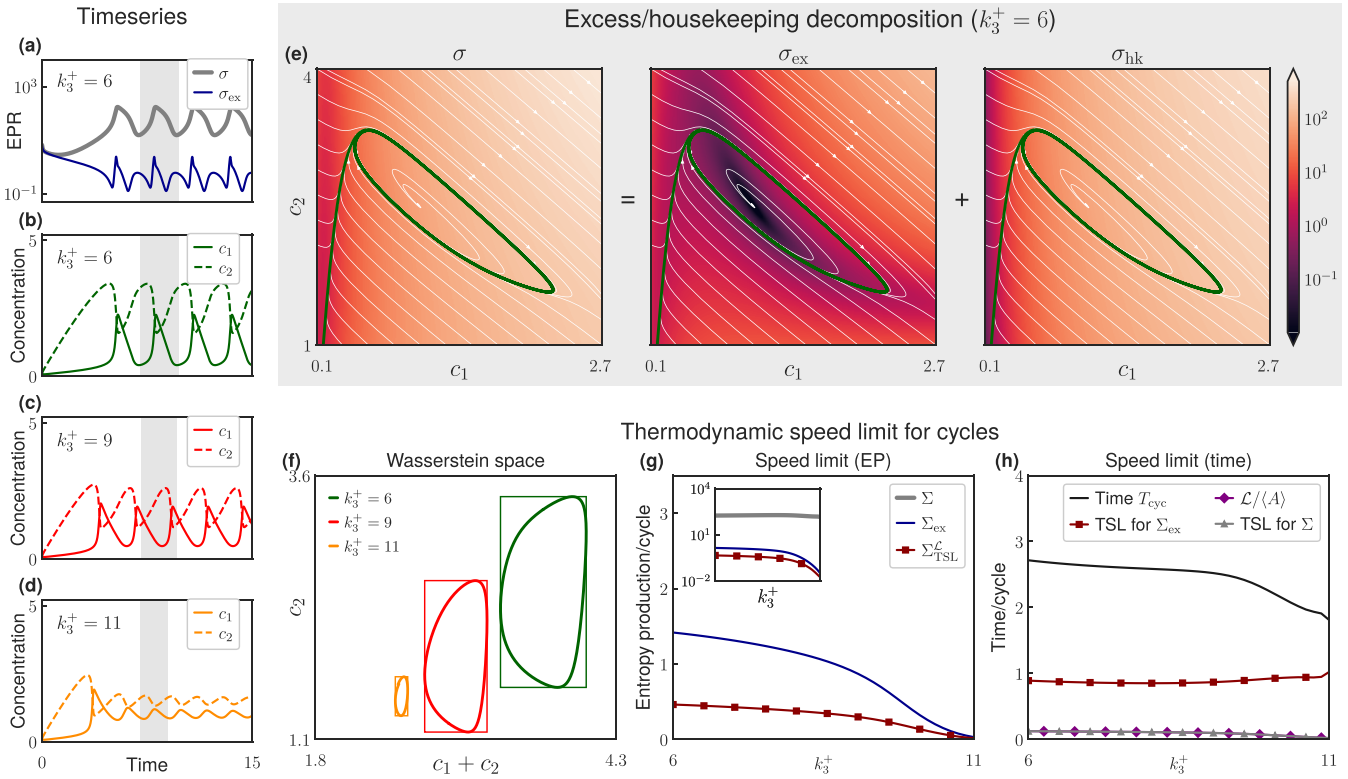

FIG. 7. Brusselator example (see also Fig. 2). (a) Time series of σ and $\sigma_{\rm ex}$ during a trajectory ($k_1^+ = k_1^- = k_2^- = k_3^- = 1$, $k_2^+ = 17$, $k_3^+ = 6$). (b)–(d) Evolution of concentrations $c_1(t), c_2(t)$ for $k_3^+ \in \{6, 9, 11\}$ (other parameters as above). The length of a single oscillator cycle is marked in gray. Excess EPR tends to be larger during faster evolution. (e) σ decomposed into $\sigma_{\rm ex}$ and $\sigma_{\rm hk}$ as a function of concentrations, with parameters as in panel (a). Streamlines (white) and actual trajectory (green) are overlaid. The excess EPR is sensitive to dynamical features, such as the limit cycles and fixed points. (f) The cycles from panels (b)–(d) are shown in coordinates $(c_1 + c_2, c_2)$ in Wasserstein space. The cycle Wasserstein length $\mathscr{L}$ is the perimeter length of the bounding boxes (see text for details). (g) The Wasserstein TSL $\Sigma_{\rm TSL}^{\mathscr{L}}$ (93) bounds EP and excess EP. Inset shows semilogarithmic scale. (h) Cycle time is bounded by the time TSLs (107).

Fig. 7(a) shows the EPR $\sigma(t)$ and excess EPR $\sigma_{\rm ex}(t)$ for $k_3^+ = 6$. Excess EPR tends to be large when the concentrations are changing rapidly.

Figure 7(e) shows the decomposition of EPR into excess and housekeeping components, as a function of the two concentrations $\boldsymbol{c} = (c_1, c_2)$. We focus on the $k_3^+ = 6$ system, the same one shown in Figs. 7(a) and 7(b). Streamlines show the dynamical evolution $\dot{\boldsymbol{x}}$ at each point in concentration space, and the actual trajectory from Fig. 7(b), shown in green. Excess EPR highlights dynamical structure: Bright regions correspond to fast evolution, and dark regions to slow evolution (the limit cycle and unstable fixed point). This is in contrast to the EPR and the housekeeping EPR, which do not appear to have any clear relationship to state dynamics.

We do not compare our analysis to the HS decomposition because, for the parameters considered here, the Brusselator does not have a steady state (no stable fixed point). Formally, the HS decomposition could be defined by taking $\boldsymbol{c}^{\rm ss}$ as the unstable fixed point; however, this is not physically meaningful and it leads to negative values of $\sigma_{\rm ex}^{\rm HS}$ (see Fig. 7 of Ref. [22]).

Next, we illustrate our finite-time TSL by deriving a bound on the dissipation and time needed to complete one oscillator cycle. For a trajectory that enters a limit cycle, we measure the cycle time $T_{\rm cyc}$ and the integrated activity $\mathcal{A}$, excess EP $\Sigma_{\rm ex}$, and EP Σ incurred during one cycle. In addition, we measure the Wasserstein length of the cycle, $\mathscr{L} = \int_0^{T_{\rm cyc}} \dot{W}(t)\, dt$. We remind the reader that, in a system without external fluxes, $\mathscr{L}$ is the *minimal activity* required to traverse the cycle in concentration space. Our finite-time TSL for $\mathscr{L}$ (93) gives the lower bound $\Sigma_{\rm TSL}^{\mathscr{L}} = 2\mathscr{L} \tanh^{-1}(\mathscr{L}/\mathcal{A})$ on EP Σ and excess EP $\Sigma_{\rm ex}$.

For Brusselator dynamics $\dot{\boldsymbol{c}} = (\dot{c}_1, \dot{c}_2)$, the Wasserstein speed can be written in closed form,

$$\dot{W} = |\dot{c}_1 + \dot{c}_2| + |\dot{c}_2|. \tag{106}$$

To see why, note that only the (coarse-grained) reactions $X_1 \rightleftharpoons X_2$ affect the concentration of X_2, so either the forward or reverse direction must have a flux of at least $|\dot{c}_2|$ (which one depends on the sign of $\dot{c}_2$). One of the two remaining reactions must have a minimal flux $|\dot{c}_1 + \dot{c}_2|$ to induce the correct evolution of X_1. Equation (106) implies that the Brusselator's Wasserstein length $\mathscr{L}$ is given by the "taxicab geometry" in coordinates $(c_1 + c_2, c_2)$.

To illustrate this, Fig. 7(f) plots the three cycles from Figs. 7(b)–7(d) in this coordinate system. In this space, the cycle's Wasserstein length $\mathscr{L}$ is equal to the perimeter of the minimal bounding box. Comparing Figs. 7(b)–7(d) and 7(f), we see that larger k_3^+ leads to shorter cycle times and smaller Wasserstein lengths.

We illustrate our TSL $\Sigma_{\rm TSL}^{\mathscr{L}}$ in Fig. 7(g) for $k_3^+ \in [6, 11]$. The bound is relatively tight, capturing about a third of the excess EP at $k_3^+ \approx 6$ and about half around $k_3^+ \approx 11$. EP

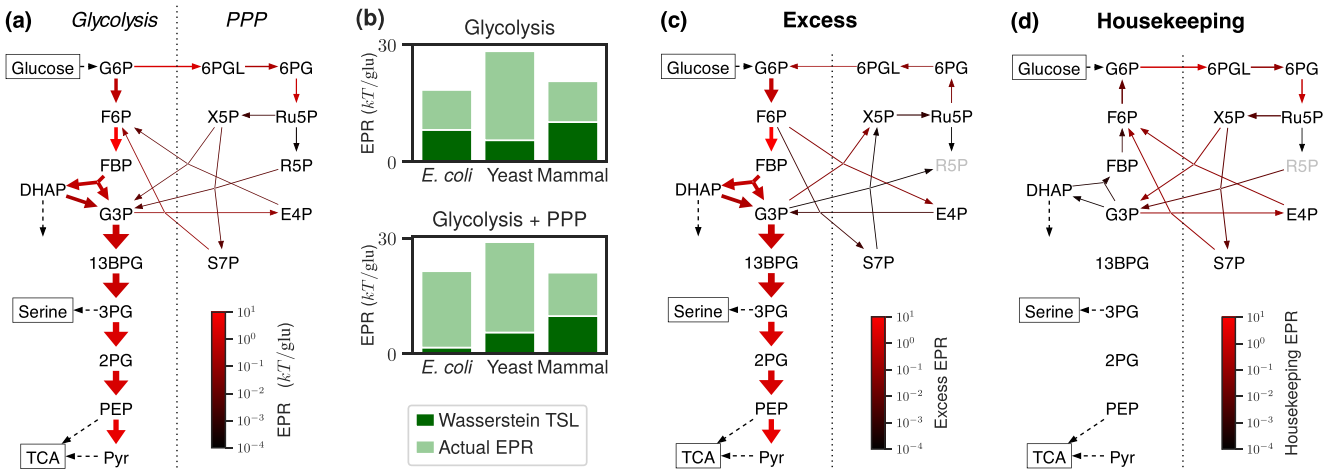

FIG. 8. Metabolic networks for *E. coli*, yeast, and mammalian cells, based on data from Ref. [125]. (a) Net fluxes and dissipation for individual reactions in *E. coli* for glycolysis and the PPP pathways. For each reversible reaction, arrow width indicates net flux ($\mathcal{J}_\rho = j_\rho - j_{\tilde{\rho}}$) and color indicates reaction-level EPR $\sigma_{\rm rev}^{(\rho)}$, see below Eq. (16). EPR is in units of kT per incoming glucose molecule. (b) Our TSL $\sigma \geqslant 2\dot{W}\tanh^{-1}(\dot{W}/A)$ bounds EPR needed to sustain external fluxes in steady state. (c), (d) The original network does not have housekeeping EPR. To illustrate our excess/housekeeping decomposition, we suppose that R5P (gray) is chemostatted by other cellular processes. This creates an emergent housekeeping cycle that runs through PPP and back to upper glycolysis. Here, we show the excess/housekeeping contributions to net fluxes (58) and EPR (59) at the level of individual reversible reactions.

is many orders of magnitude larger than excess EP (see inset in semilogarithmic scale) and it does not lead to a useful TSL.

Finally, we use Eq. (98) to bound the minimal time needed to complete a cycle as

$$T_{\rm cyc} \geqslant \frac{\mathscr{L}}{\langle A\rangle}\coth\frac{\Sigma_{\rm ex}}{2\mathscr{L}} \geqslant \frac{\mathscr{L}}{\langle A\rangle}\coth\frac{\Sigma}{2\mathscr{L}} \geqslant \frac{\mathscr{L}}{\langle A\rangle}. \tag{107}$$

The last bound $T_{\rm cyc} \geqslant \mathscr{L}/\langle A\rangle$ holds even in the absolutely irreversible regime where $\sigma_{\rm ex}$ diverges. These bounds are illustrated in Fig. 7(h) for $k_3^+ \in [6, 11]$. The excess EP bound is relatively tight across the range of k_3^+ values. The EP-based bound and absolute irreversibility bound are essentially equivalent, since EP is very large for this system.

C. Metabolic networks

In our last example, we illustrate our approach on real-world metabolic networks from three species: *Escherichia coli* bacteria, yeast, and a mammalian cell line. Using our TSL, we demonstrate that central metabolism operates with high efficiency, given constraints imposed by its overall activity and stoichiometry. We also show that our excess/housekeeping decomposition is able to identify fluxes and dissipation associated with futile metabolic cycles.

As standard in metabolic modeling and flux-balance analysis [113], we model the metabolic networks as open deterministic CRNs in steady state. We focus on two core pathways of central metabolism: glycolysis, responsible for breaking down glucose into pyruvate while producing ATP, and the pentose phosphate pathway (PPP), responsible for production of nicotinamide adenine dinucleotide phosphate (NADPH) and ribose-5-phosphate (R5P) (used downstream for nucleotide synthesis). Central metabolism is highly conserved across different branches of life. Accordingly, the networks of all three biological species involve the same $d = 18$ metabolites and $m = 34$ one-way reactions (17 reversible reactions). For illustration, we show the *E. coli* network in Fig. 8(a), with glycolysis reactions in the left column and PPP reactions in the right column.

To determine flux and thermodynamic parameters, we use published data from isotopic labeling experiments [125]. In these experiments, each of the three cell types was grown on a nutrient-rich medium with high glucose. Following Ref. [125], we treat the first reaction in glycolysis (phosphorylation of glucose to G6P) as external. Also, we do not explicitly represent internal cofactors and small molecules (ATP, ADP, Pi, CO2, etc.), assuming that they are chemostatted at constant concentrations by cellular homeostasis. For details about the dataset, see SM6 in the Supplemental Material [53].

In addition to metabolic reactions, the network is also subject to external fluxes of nutrients and metabolites from the environment, most importantly the inflow of phosphorylated glucose (G6P), outflow of phosphoenolpyruvic acid (PEP) and pyruvate to the citric acid cycle, outflow of 3-Phosphoglyceric acid (3PG) (used for serine production), and outflow of dihydroxyacetone phosphate (DHAP). External fluxes are set to balance net production, $\boldsymbol{I} = \nabla^\top \boldsymbol{j}$, ensuring steady-state conditions.

Some of the flux and thermodynamic data are visualized in the *E. coli* network in Fig. 8(a). The direction and width of arrows indicate the direction and net flux of different reactions, and color indicates the reaction EPR (summed for both forward and reverse reactions). Fluxes are normalized relative to the rate of glucose uptake; thus, EPR values should be understood in units of kT/glucose. The yeast and mammalian networks (not shown) are qualitatively similar to the *E. coli* network, though they exhibit less flux through the PPP pathway.

The stoichiometric matrix has rank $17 = m/2$; hence, according to algebraic condition (56), the forces are conservative and housekeeping EPR vanishes. This is reasonable from a biological perspective, since glycolytic metabolism appears to avoid futile cycles [126]. Although the forces are conservative, the steady state is nonequilibrium due to external fluxes.

Next, we use our TSL from Eq. (83) to bound the EPR:

$$\sigma \geqslant \sigma_{\mathrm{TSL}} = 2\dot{W} \tanh^{-1} \frac{\dot{W}}{A}. \tag{108}$$

The Wasserstein speed $\dot{W}$ (82) quantifies the minimal activity required to balance external fluxes $\boldsymbol{I} = \nabla^{\top} \boldsymbol{j}$, given the system's fixed stoichiometry. It can be understood as a topology-dependent measure of the intensity of the external fluxes.

Results for the three biological species (*E. coli*, yeast, and mammalian) are shown in Fig. 8(b). In the top row, we compute the TSL only for the glycolysis pathway. The dark green bars show the minimal bound (108), while the lighter green bars show the actual EPR. We observe that glycolysis is remarkably efficient: Its dissipation is close to the fundamental minimum permitted by its stoichiometry and overall activity (reactions per unit time). *E. coli* and mammalian cells achieve nearly $\sigma_{\mathrm{TSL}}/\sigma \approx 50\%$ efficiency, while yeast achieves $\sigma_{\mathrm{TSL}}/\sigma \approx 20\%$ efficiency. This accords with biochemical studies, which have shown that the glycolysis pathway appears highly optimized for thermodynamic performance [127,128].

In the bottom row of Fig. 8(b), we compute the TSL when including both the glycolysis and PPP pathways. For yeast and mammalian cells, the results are essentially the same as without PPP (top row), because these networks have small flux through the PPP pathway in these data. The *E. coli* network has significant flux through the PPP pathway. For this reason, when accounting for PPP, *E. coli* is further away from the fundamental bound set by the TSL, indicating that the PPP pathway is less thermodynamically efficient. Arguably, this is not surprising, since the main function of PPP is synthesis of precursor metabolites rather than energy conversion.

In our final analysis, we demonstrate our excess/housekeeping decomposition. As mentioned, the metabolic networks considered above have maximum rank and thus no housekeeping EPR. In biological terms, their stoichiometric matrix does not permit "futile cycles" (sequences of reactions that contribute to EPR but have no net effect on metabolites). However, the definition of futile cycle depends on which metabolites are tracked internally, and "emergent cycles" [129] can appear when metabolites are removed from the stoichiometric matrix and treated as chemostatted.

To illustrate this phenomenon, we removed ribose-5-phosphate from the stoichiometric matrix of the *E. coli* network. Biologically, this may represent chemostatting of R5P by homeostatic mechanisms that buffer nucleotide demand. This reduces the rank of the stoichiometric matrix to 16, leading to an emergent cycle and a nontrivial excess/housekeeping decomposition. Figures 8(c) and 8(d) show excess/housekeeping fluxes and dissipation at the level of individual reversible reactions [see Eqs. (58) and (59)]. Thus, by treating R5P as chemostatted, we reveal a maintenance cycle that circulates carbon through the PPP back to upper glycolysis. This cycle is captured by the housekeeping fluxes $\boldsymbol{j}_{\mathrm{hk}}$ shown in Fig. 8(d).

To our knowledge, our excess/housekeeping decomposition provides the first method for quantifying the dissipation associated with futile cycles in general open CRNs with external fluxes. Notably, our decomposition does not depend on a decomposition of the null-space of the stoichiometric matrix, which in general is nonunique [129].

TABLE II. Comparison between our approach and several other generalized potentials and excess/housekeeping decompositions. See Sec. IX for details.

	Hatano-Sasa	Euclidean-Onsager [22]	Hessian geometry [51]	Ours
Variational definition		✓	✓	✓
Large deviations	✓		✓	✓
Thermodynamic inference				✓
Minimum EP principle		✓		
Additive invariance				✓
Coarse graining	MJPs			✓

IX. COMPARISON TO PREVIOUS APPROACHES

Several other generalizations of the free energy potential and excess/housekeeping decompositions have been proposed, including in some recent publications. Here, we compare our approach to several existing proposals. In the first subsection, we compare to other "variational" proposals, which define the generalized potential and EPR decomposition via an optimization problem. In the second subsection, we compare to the steady-state approach, as previously reviewed in Sec. III B.

For purposes of comparison, we will draw attention to several important properties (summarized in Table II):

(1) *Variational definition*: Are the quantities defined in a variational way, via the optimization of some objective function? Such definitions typically guarantee existence of a well-defined potential and decomposition in a general class of systems, and they allow derivations of TURs and related thermodynamic bounds.

(2) *Large deviations*: Do the quantities have an interpretation in terms of a large-deviation principle? This provides physical meaning in terms of fluctuation statistics.

(3) *Thermodynamic inference*: Do the quantities permit thermodynamic inference using local-in-time statistics? This allows estimation and validation from real-world experimental data.

(4) *Minimum EP principle*: Does the variational principle identify minimal EP required to achieve some desired state evolution? This is relevant for thermodynamic optimal control, i.e., implementing a given evolution while minimizing entropy production.

(5) *Additive invariance*: Are the definitions invariant under joining of independent systems? This enforces the fundamental physical principle of extensivity.

(6) *Coarse graining*: Are the definitions invariant under merging of reactions with the same stoichiometry? This guarantees that the potential and excess EPR do not depend on detailed information about which reactions mediate state changes, and has practical benefits (see Sec. V D).

A. Variational approaches

Our definition of the generalized potential and excess/housekeeping decomposition is based on the variational principle (38). Several previous proposals have also defined generalized potentials and EPR decompositions using different (though related) variational principles. Here, we summarize and compare these approaches.

Given a system with stoichiometry ∇ and fluxes $\boldsymbol{j}$, we say that a generalized potential $\hat{\boldsymbol{\phi}}$ is defined variationally when it is expressed as

$$\hat{\boldsymbol{\phi}} = \operatorname*{argmax}_{\boldsymbol{\phi}\in\mathbb{R}^d} \mathcal{L}(\boldsymbol{\phi}, \boldsymbol{j}, \nabla), \tag{109}$$

where $\mathcal{L}$ is some objective function that depends on the "test potential" $\boldsymbol{\phi}$, the fluxes $\boldsymbol{j}$, and the stoichiometric matrix ∇ (it could also depend on the state $\boldsymbol{x}$, though we leave this dependence implicit in our notation). The optimal potential exists as long as $\mathcal{L}$ satisfies some mild regularity conditions, though it may not be unique (e.g., due to conserved quantities). As we will see below, in existing proposals, the objective involves the Legendre transform of some convex function.

One appealing feature of the variational approach is that a generalized free energy can be defined for any system, whether stochastic or deterministic, as long as one can define an appropriate objective $\mathcal{L}$. Of course, different objectives produce different results, and the choice of the objective must be justified by physical, operational, or formal considerations.

In our approach, the generalized potential $\boldsymbol{\phi}^\star$ is defined using the following objective:

$$\boldsymbol{\phi}^\star = \operatorname*{argmax}_{\boldsymbol{\phi}\in\mathbb{R}^d}[-\boldsymbol{j}^\top\nabla\boldsymbol{\phi} - \Phi(\boldsymbol{j}, -\nabla\boldsymbol{\phi})], \tag{110}$$

where for convenience we introduced the function $\Phi(\boldsymbol{j}, \boldsymbol{\theta}) := \boldsymbol{j}^\top(e^{-\boldsymbol{\theta}} - \mathbf{1})$. The excess EPR is the maximum value of the objective, $\sigma_{\rm ex} = -\boldsymbol{j}^\top\nabla\boldsymbol{\phi}^\star - \Phi(\boldsymbol{j}, -\nabla\boldsymbol{\phi}^\star)$. The optimality condition, $\nabla^\top\boldsymbol{j} = \nabla^\top(\tilde{\boldsymbol{j}} \circ e^{-\nabla\boldsymbol{\phi}^\star})$, implies that $\boldsymbol{\phi}^\star$ exponentially tilts the reverse fluxes so as to recover the actual net production.

As discussed above, Φ is the cumulant generating function of short-time dynamical fluctuations, and the objective (110) is defined in terms of its Legendre transform. For this reason, our definitions have a natural interpretation in terms of large deviations: $\sigma_{\rm ex}$ is the rate function that controls irreversibility of the state evolution, and the potential $\boldsymbol{\phi}^\star$ is the "most irreversible" state observable. Our definitions are directly accessible to thermodynamic inference. In particular, local-in-time statistics provide a tight lower bound on the excess EPR via the TUR (79).

Several other variational definitions have been advanced in the literature. For instance, the present authors previously proposed the Euclidean-Onsager potential [22], defined as

$$\boldsymbol{\phi}^{\rm ons} = \operatorname*{argmax}_{\boldsymbol{\phi}\in\mathbb{R}^d}[-2\boldsymbol{j}^\top\nabla\boldsymbol{\phi} - \boldsymbol{\phi}^\top\nabla^\top L(\boldsymbol{j})\nabla\boldsymbol{\phi}]. \tag{111}$$

Here, we introduced the diagonal matrix $L_{\rho\rho} := \frac{1}{2}(j_\rho - \tilde{j}_\rho)/\ln(j_\rho/\tilde{j}_\rho)$ that encodes reaction-level Onsager coefficients that map forces to net fluxes. The associated excess EPR $\sigma_{\rm ex}^{\rm ons}$ is defined as the maximum of the objective (111). Equation (111) involves the Legendre transform of the quadratic function $\boldsymbol{\phi}^\top\nabla^\top L(\boldsymbol{j})\nabla\boldsymbol{\phi}$. In SM5.1 in the Supplemental Material [53], we discuss this definition in more depth, and we derive the inequality $\sigma_{\rm ex} \leqslant \sigma_{\rm ex}^{\rm ons}$.

The Euclidean-Onsager approach does not have a direct relation to large deviations and, outside the linear-response regime, it is not easily accessible to thermodynamic inference. Nonetheless, it has a natural interpretation in terms of a minimum EP principle. In particular, it was shown that $\sigma_{\rm ex}^{\rm ons}$ is the minimum EPR incurred by any fluxes $\boldsymbol{j}'$ that have the same net production ($\nabla^\top\boldsymbol{j}' = \nabla^\top\boldsymbol{j}$) and the same Onsager coefficients [$L(\boldsymbol{j}') = L(\boldsymbol{j})$] as $\boldsymbol{j}$ [22]. Moreover, these optimal fluxes have conservative forces $-\nabla\boldsymbol{\phi}^{\rm ons}$.

Another generalized potential was proposed in Ref. [51]. It is helpful to introduce the "symmetrized" fluxes $j_\rho^s := \sqrt{j_\rho\tilde{j}_\rho}$, which are equilibrium fluxes that have the same "frenetic activity" as the actual fluxes, $\sqrt{j_\rho^s\tilde{j}_\rho^s} = \sqrt{j_\rho\tilde{j}_\rho}$. The authors define two different potentials, which in our notation can be written as

$$\boldsymbol{\phi}^{\rm hess} = \operatorname*{argmax}_{\boldsymbol{\phi}\in\mathbb{R}^d}[-\boldsymbol{j}^\top\nabla\boldsymbol{\phi} - \Phi(\boldsymbol{j}^s, -\nabla\boldsymbol{\phi})], \tag{112}$$

$$\boldsymbol{\psi}^{\rm hess} = \operatorname*{argmax}_{\boldsymbol{\phi}\in\mathbb{R}^d}[-\Phi(\boldsymbol{j}^s, \ln(\boldsymbol{j}/\tilde{\boldsymbol{j}}) + \nabla\boldsymbol{\phi})], \tag{113}$$

where Φ is the short-time CGF mentioned above. The corresponding excess EPR $\sigma_{\rm ex}^{\rm hess}$ has a somewhat complicated expression; for details and a numerical comparison, see SM5.2 in the Supplemental Material [53]. The authors refer to their proposal as one based on "Hessian geometry," to contrast with Riemannian geometry as might be induced by a quadratic metric [e.g., $L(\boldsymbol{j})$ from Eq. (111)].

The potentials $\boldsymbol{\phi}^{\rm hess}$ and $\boldsymbol{\psi}^{\rm hess}$ can be interpreted in terms of the family of fluxes $\boldsymbol{\theta} \mapsto \boldsymbol{j}^s \circ e^{\boldsymbol{\theta}/2}$, characterized by varying thermodynamic forces $\boldsymbol{\theta}$ and fixed frenetic activity $\boldsymbol{j}^s$. Within this family, the forces $-\nabla\boldsymbol{\phi}^{\rm hess}$ induce the correct state dynamics, $\nabla^\top(\boldsymbol{j}^s \circ e^{-\nabla\boldsymbol{\phi}^{\rm hess}/2}) = \nabla^\top\boldsymbol{j}$, while the forces $\ln(\boldsymbol{j}/\tilde{\boldsymbol{j}}) + \nabla\boldsymbol{\psi}^{\rm hess}$ make the fluxes stationary, $\nabla^\top(\boldsymbol{j}^s \circ e^{(\ln(\boldsymbol{j}/\tilde{\boldsymbol{j}})+\nabla\boldsymbol{\psi}^{\rm hess})/2}) = \mathbf{0}$. We note that in the Euclidean-Onsager approach, with its simpler quadratic structure, the potential $\boldsymbol{\phi}^{\rm ons}$ plays both of these dynamical roles.

The Hessian-geometry approach is closely related to research on large deviations of CRNs [65,69,130–133], and the two objectives [Eqs. (112) and (113)] involve Legendre transforms of the CGF. However, the CGFs are evaluated under the "reference" equilibrium fluxes $\boldsymbol{j}^s$, rather than the actual fluxes $\boldsymbol{j}$. For this reason, this approach is not directly amenable to thermodynamic inference using local-in-time statistics. Also, to our knowledge, the excess EPR $\sigma_{\rm ex}^{\rm hess}$ does not have an interpretation as a minimum EPR expression.

To summarize, we have shown that our approach is closely related to local-in-time statistics and thermodynamic inference. The Euclidean-Onsager approach is related to the minimum EP principle with prescribed Onsager coefficients. The Hessian approach is related to dynamical large deviations of equilibrium systems with prescribed frenetic activity.

Interestingly, there is another aspect that distinguishes our approach from the others: Among the objectives listed above, ours (110) is the only one that is linear in the fluxes $\boldsymbol{j}$. This formal property has important physical consequences. Consider two systems characterized by fluxes $\boldsymbol{j}$ and $\boldsymbol{j}'$ that have the same stoichiometry ∇ and the same generalized potential $\boldsymbol{\phi}^\star$. For example, these might represent reactions occurring within two different regions of a reactor volume. Then, when the systems are coarse grained into a single system with fluxes $\boldsymbol{j}+\boldsymbol{j}'$, our definition is unique in guaranteeing that the generalized potential $\boldsymbol{\phi}^\star$ remains unchanged and the excess EPR adds extensively, $\sigma_{\rm ex}(\boldsymbol{j}+\boldsymbol{j}') = \sigma_{\rm ex}(\boldsymbol{j}) + \sigma_{\rm ex}(\boldsymbol{j}')$. This invariance mirrors that of conservative systems: When systems with the same conservative forces are combined, the forces do not change and the EPR adds extensively. This additive invariance also underlies the coarse-graining property, discussed in Sec. V D, which states that our definitions are invariant when merging reactions with the same stoichiometry.

All of the above proposals become equivalent in the linear-response regime near equilibrium. Specifically, when $\boldsymbol{j} \approx \boldsymbol{j}^{\rm eq}$ (where $\boldsymbol{j}^{\rm eq} = \tilde{\boldsymbol{j}}^{\rm eq}$ are some equilibrium fluxes), all of the objectives reduce to the quadratic form of Eq. (62). Thus, all objectives share the same properties and interpretations (large deviations, thermodynamic inference, minimum EPR, additive invariance) near equilibrium, but become distinct in the far-from-equilibrium regime.

Finally, we may consider Fokker-Planck dynamics as the continuum limit of MJPs in the linear-response regime [22, Appendix B]. In this limit, the dynamics converge to the continuity equation $\partial_t \rho_t = -\mathrm{div}\mathbf{j}$, and the generalized potentials converge to the potential proposed by Maes and Netočný (MN) for Fokker-Planck systems [41]. As described in SM5.3 in the Supplemental Material [53], it has a variational characterization [85],

$$\varphi^{\rm MN} = \operatorname*{argmax}_{\varphi:\mathbb{R}^k\to\mathbb{R}} \int (2\,\mathbf{j}\cdot \mathrm{grad}\varphi - \rho_t \|\mathrm{grad}\varphi\|^2) d\boldsymbol{r}, \qquad (114)$$

analogous to Eq. (62). The variational definitions in Eqs. (110)–(113) provide different generalizations of the Maes-Netočný decomposition to far-from-equilibrium discrete systems.

B. Steady-state approach

We now briefly compare our proposal to the steady-state approach, as previously discussed in Sec. III B. For simplicity, we focus on the case of MJPs and complex-balanced CRNs, where the steady-state potential has the simple form $\boldsymbol{\phi}^{\rm ss} = \mathrm{grad}_{\boldsymbol{x}} D(\boldsymbol{x}\|\boldsymbol{x}^{\rm ss}) = \ln(\boldsymbol{x}/\boldsymbol{x}^{\rm ss})$ and the excess/housekeeping decomposition is the HS decomposition.

As we show in SM5.4 in the Supplemental Material [53], the HS housekeeping EPR can be expressed in the notation of Eq. (53) as

$$\sigma_{\rm hk}^{\rm HS} = \mathscr{D}(\boldsymbol{f}\|-\nabla\boldsymbol{\phi}^{\rm ss}). \qquad (115)$$

This has a simple connection to our information-geometric interpretation, as visualized in Fig. 3(b): The HS housekeeping EPR is represented by a line from the actual forces $\boldsymbol{f}$ to the point $-\nabla\boldsymbol{\phi}^{\rm ss}$ on the conservative manifold. Since our housekeeping EPR satisfies $\sigma_{\rm hk} = \min_{\boldsymbol{\phi}} \mathscr{D}(\boldsymbol{f}\|-\nabla\boldsymbol{\phi})$, it is always smaller than the HS housekeeping EPR:

$$\sigma_{\rm hk} \leqslant \sigma_{\rm hk}^{\rm HS}, \quad \sigma_{\rm ex} \geqslant \sigma_{\rm ex}^{\rm HS}. \qquad (116)$$

The difference $\sigma_{\rm ex} - \sigma_{\rm ex}^{\rm HS} \geqslant 0$ has been previously termed "coupling EPR" in Fokker-Planck systems [86], where it quantifies the gap between the HS and Maes-Netočný [41] definitions of housekeeping EPR.

In terms of the properties listed in Table II, the steady-state potential is closely related to the large deviations of steady-state fluctuations (see Sec. III B). However, for nonstationary systems, it is not directly accessible to thermodynamic inference using local-in-time statistics. In addition, the HS excess EPR does not satisfy a minimum EPR principle.

The steady-state potential does not satisfy "additive invariance," since combining systems with the same steady-state potential $\boldsymbol{\phi}^{\rm ss}$ does not necessarily preserve $\boldsymbol{\phi}^{\rm ss}$. However, in MJPs, it obeys the coarse-graining condition discussed in Sec. V D: $\boldsymbol{\phi}^{\rm ss}$ and $\sigma_{\rm ex}^{\rm HS}$ are invariant under merging of transitions with same stoichiometry (i.e., same $i \to j$) mediated by different reservoirs. For MJPs where the coarse-grained fluxes have a conservative form, as in Eq. (49) and the two-level MJP discussed in Sec. V B, our potential $\boldsymbol{\phi}^\star$ and EPR decomposition is the same one as the HS one.

Finally, we consider the linear-response regime, for simplicity focusing on the case of MJPs. Near steady state $\boldsymbol{p} \approx \boldsymbol{p}^{\rm ss}$, we may approximate $\phi_i^{\rm ss} = \ln(p_i/p_i^{\rm ss}) \approx (p_i - p_i^{\rm ss})/p_i^{\rm ss}$. Multiplying both sides by $R_{ji}p_i^{\rm ss}$ and summing gives

$$\dot{\boldsymbol{p}} \approx -G(\boldsymbol{p}^{\rm ss})\boldsymbol{\phi}^{\rm ss}, \qquad (117)$$

where $G_{ji}(\boldsymbol{p}^{\rm ss}) := -R_{ji}p_i^{\rm ss}$ acts as the HS mobility matrix. This may be compared to our linear-response equation [Eq. (70)]. However, in nonconservative systems, the matrix G is generally not symmetric; thus, Eq. (117) does not satisfy Onsager's reciprocal relations. Interestingly, our steady-state mobility matrix is the additive symmetrization of the HS one, $H_{\rm ss}(\boldsymbol{p}^{\rm ss}) = (G(\boldsymbol{p}^{\rm ss}) + G(\boldsymbol{p}^{\rm ss})^\top)/2$ [see Eq. (64)]. As expected, the two agree for conservative systems.

Finally, Mandal and Jarzynski [134] analyzed the linear-response regime of the HS decomposition for MJPs. They showed that the HS excess EPR defines a Riemannian geometry over thermodynamic states. However, the HS friction tensor is not local in time, but rather determined by fluctuations under the stationary process [134, Eq. (24)].

X. DISCUSSION

In this paper, we proposed a generalized free energy and excess/housekeeping decomposition for nonconservative systems. We demonstrated that our approach is applicable to a broad class of stochastic and deterministic systems. We also showed that our definitions are amenable to thermodynamic inference, and that they lead to useful bounds such as thermodynamic uncertainty relations and thermodynamic speed limits.

In conservative systems, the nonequilibrium free energy can be understood from two different perspectives. The first perspective is state based (static): The nonequilibrium free energy is related to the deviation between the actual state and the equilibrium state, and it controls properties like equilibrium fluctuations and extractable work in nonequilibrium

states. The second perspective is dynamical (flux based): The nonequilibrium free energy gives rise to conservative thermodynamic forces, and it controls properties like dynamical fluctuations and gradient flow dynamics. This dynamical perspective can be traced back to the visionary work of Onsager [99,135,136].

These two perspectives lead to different generalizations to nonconservative systems. Until now, most research has considered the first (static) perspective by defining the generalized potential in terms of the deviation between the actual state and (nonequilibrium) steady state.

We proposed an alternative approach based on the dynamical perspective. Here, the generalized potential is characterized by offering the "conservative approximation" to the forces. From the perspective of large deviations, the generalized potential is the most irreversible state observable, and the excess EPR is defined as its degree of irreversibility.

We mention several possible directions for future work.

First, to our knowledge, our analysis of metabolic networks represents the first application of a TSL to a complex biological reaction network. A promising direction for further research is the application of theoretical results from nonequilibrium thermodynamics to large-scale metabolic modeling and flux-balance analysis [137–139].

Second, as we discuss near Eq. (65), in the linear-response regime, our excess EPR defines a Riemannian geometry over the set of thermodynamic states. It would be interesting to explore thermodynamic length and optimal protocols using this geometry in nonconservative systems.

Third, in this paper, we explored thermodynamic bounds, such as TURs and TSLs, for the excess EPR. However, we also showed that the housekeeping EPR, the nonexcess contribution to dissipation, quantifies the nonconservative nature of the forces. Future work may investigate the housekeeping contribution in more depth, e.g., from the perspective of large deviations, TURs, and TSLs, as well as decompositions into elementary cycles [22].

Fourth, here we considered the excess/housekeeping decomposition applied to the EPR at a given instant in time. However, the variational principle that defines excess EPR may also be considered for time-extended stochastic processes, where entropy production is quantified as the relative entropy between trajectory distributions. This may lead to fluctuation theorems for excess/housekeeping EPR and generalizations to non-Markovian stochastic processes. There are also interesting connections to other trajectory-level variational principles, such as Schrödinger-bridge problems [140,141] and maximum caliber inference [142].

Finally, it is interesting to consider our generalized potential and information-geometric decomposition in other types of systems, including reaction-diffusion [23], hydrodynamic [143], and quantum [144] systems, which have recently been explored using Euclidean geometry.

ACKNOWLEDGMENTS

The authors thank Masafumi Oizumi and Praful Gagrani for fruitful discussions. A.D. was supported by JSPS KAKENHI Grants No. 19H05795 and No. 22K13974. K.Y. was supported by Grant-in-Aid for JSPS Fellows (Grant No. 22J21619). S.I. was supported by JSPS KAKENHI Grants No. 21H01560, No. 22H01141, No. 23H00467, and No. 24H00834, JST ERATO Grant No. JPMJER2302, and UTEC-UTokyo FSI Research Grant Program. A.K. received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Sklodowska-Curie Grant Agreement No. 101068029.

DATA AVAILABILITY

No new data were created in this study. All data analyzed in this work are publicly available from the original sources cited in the references.

APPENDIX A: NONLINEAR MJPS

In this Appendix, we discuss applications of our approach to nonlinear MJPs. Nonlinear MJPs are stochastic master equations in which the transition rates may depend on the probability distribution. In terms of the MJP notation of Eq. (7), this means that the rates R^{α}_{ij} may depend on $\boldsymbol{p}$. Nonlinear MJPs are often used in physics [145–147] and biology [148–150] to model many-body systems (many particles, regions, species, etc.) with mean-field interactions [150–152].

We illustrate our formalism on a classic model by Malek-Mansour and Nicolis of local fluctuations in a reaction-diffusion system [145]. Let us consider a system that contains a single chemical species, and let $i \in \{0, 1, 2, \ldots\}$ indicate the number of particles of that species within a given small region. It is assumed that the region is locally well mixed, that it undergoes the same fluctuations as its environment, and that it is statistically independent of its environment. Then, the probability distribution $\boldsymbol{p}$ of i evolves according to a birth-death process [145],

$$\dot{p}_i = R_i(\boldsymbol{p}) + \mathscr{D}\langle i\rangle_{\boldsymbol{p}}(p_{i-1} - p_i) + \mathscr{D}((i+1)p_{i+1} - i p_i).$$

(The term $\mathscr{D}\langle i\rangle_{\boldsymbol{p}} p_{i-1}$ is omitted at the boundary $i = 0$.) The first term $R_i(\boldsymbol{p})$ is the contribution due to local reactions, the second term due to particle exchange from the environment, and the third term due to particle exchange to the environment. $\mathscr{D}$ is an effective diffusion frequency and $\langle i\rangle_{\boldsymbol{p}}$ is the expected particle count under distribution $\boldsymbol{p}$ (i.e., particle density in the environment). Because the transition rate $\mathscr{D}\langle i\rangle_{\boldsymbol{p}}$ depends on $\boldsymbol{p}$, this process is a nonlinear MJP.

This system can be put in the form of the discrete continuity equation (1) by introducing an appropriate incidence matrix ∇ and probability fluxes. In particular, the fluxes due to exchanges with the environment are defined as

$$\begin{aligned} j_{i\to i+1}(\boldsymbol{p}) &= \mathscr{D}\langle i\rangle_{\boldsymbol{p}} p_i, \\ j_{i+1\to i}(\boldsymbol{p}) &= \mathscr{D}(i+1)p_{i+1}. \end{aligned} \tag{A1}$$

The force associated with gaining one particle from the reservoir can be written as a sum of three terms,

$$f_{i\to i+1} = \ln \frac{j_{i\to i+1}}{j_{i+1\to i}} = \ln \frac{p_i}{p_{i+1}} + \ln \frac{1}{i+1} - \ln \frac{1}{\langle i\rangle_{\boldsymbol{p}}},$$

representing the change of the region's statistical entropy, its internal entropy, and the environment's entropy. When the reaction term $R_i(\boldsymbol{p})$ arises from a chemical master equation,

it can also be expressed in terms of an incidence matrix and reaction fluxes.

Using these fluxes, EPR is defined as in the main text [Eq. (4)]. The generalized free energy and the excess/housekeeping decomposition are also defined as in the main text [Eqs. (38) and (39)]. As usual, the excess term quantifies the nonstationary contribution, while the housekeeping term quantifies the nonconservative contribution.

To understand the meaning of conservative forces, it is useful to note that any Poisson distribution ($p_i^* = e^{-\lambda}\lambda^i/i!$) satisfies detailed balance for the exchange fluxes,

$$j_{i\to i+1}(\boldsymbol{p}^*) = j_{i+1\to i}(\boldsymbol{p}^*),$$

irrespective of the mean λ. Similarly, the forces for the exchange transitions are conservative,

$$f_{i\to i+1} = \phi_i^* - \phi_{i+1}^*, \tag{A2}$$

for the potential $\phi_i^* = \ln(p_i/p_i^*)$. Therefore, the overall system is conservative, and the housekeeping EPR vanishes, as long as the reaction fluxes that specify $R_i(\boldsymbol{p})$ obey detailed balance for some Poisson equilibrium distribution.

APPENDIX B: SYSTEMS WITH ODD VARIABLES

1. Entropy production rate

We show how our approach applies to MJPs with odd variables, such as velocity or momentum, whose sign changes under time reversal.

We first write the expression of EPR. Consider an MJP coupled to a single heat bath that evolves over a small time interval $[t, t+dt]$. The conditional probability that the system is in state j at time $t+dt$, given state i at time t, is

$$T_{j|i} = dt\, R_{ji} + O(dt^2).$$

The EP is the relative entropy between the forward and backward joint distributions,

$$\Sigma = D\big(p_i T_{j|i} \| p_j T_{\epsilon i|\epsilon j}\big), \tag{B1}$$

where ϵi indicates the conjugation of microstate i (odd-parity variables flipped in sign). The conjugation of the reverse conditional probability $T_{\epsilon i|\epsilon j}$ follows from the principle of local detailed balance for systems with odd variables (see Refs. [153–155], also Sec. 5.3.4 in Ref. [156]). We may write Eq. (B1) more explicitly as

$$\Sigma = \underbrace{\sum_{i\neq j} p_i T_{j|i} \ln \frac{p_i T_{j|i}}{p_j T_{\epsilon i|\epsilon j}}}_{\text{Transitions}} + \underbrace{\sum_i p_i T_{i|i} \ln \frac{p_i T_{i|i}}{p_i T_{\epsilon i|\epsilon i}}}_{\text{Diagonals}}.$$

The second term is the contribution to EP due to different escape rates under the forward and reverse dynamics. This contribution only appears in systems with odd variables, since it vanishes when $i = \epsilon i$.

The EPR is the time derivative of EP, $\sigma = d_t \Sigma$. With a bit of algebra, this derivative can be found as

$$\sigma = \sum_{j\neq i} \left(p_i R_{ji} \ln \frac{p_i R_{ji}}{p_j R_{\epsilon i\epsilon j}} - p_i R_{ji} + p_j R_{\epsilon i\epsilon j} \right). \tag{B2}$$

For derivations, see Eq. (4.10) of Ref. [157] or Eq. (23) of Ref. [158]. For a system coupled to multiple reservoirs indexed by α, this may be generalized as

$$\sigma = \sum_{j\neq i,\alpha} \left(p_i R^\alpha_{ji} \ln \frac{p_i R^\alpha_{ji}}{p_j R^\alpha_{\epsilon i\epsilon j}} - p_i R^\alpha_{ji} + p_j R^\alpha_{\epsilon i\epsilon j} \right). \tag{B3}$$

The EPR can be expressed as a relative entropy between flux vectors. We define a reaction ρ for each one-way transition $(i \to j, \alpha)$ with flux $j_\rho = p_i R^\alpha_{ji}$ and reverse flux $\tilde{j}_\rho = p_j R^\alpha_{\epsilon i\epsilon j}$. Importantly, unlike in systems without odd variables, the reverse flux does not have to correspond to the forward flux of any reaction. The force across reaction ρ is

$$f_\rho = \ln \frac{j_\rho}{\tilde{j}_\rho} = \ln \frac{p_i R^\alpha_{ji}}{p_j R^\alpha_{\epsilon i\epsilon j}}, \tag{B4}$$

as in Eq. (2). Finally, the EPR (B3) can be written as the relative entropy between forward and reverse fluxes, as in Eq. (14):

$$\sigma = \mathcal{D}(\boldsymbol{j} \| \tilde{\boldsymbol{j}}).$$

Note that Eq. (B3) does not have the simple "flux-force" form $\sigma = \sum_\rho j_\rho f_\rho$, as it does in systems without odd variables. This reflects the irreversibility due to different forward and reverse escape rates.

2. Generalized free energy and excess/housekeeping decomposition

Many of our results continue to hold for MJPs with odd variables, including the definition of the generalized potential and the excess/housekeeping decomposition. One important caveat is that we do not simplify our expressions by using Eq. (3), the relationship between forward and reverse fluxes that holds without odd variables. Without using this equation, excess EPR (38) is defined as

$$\sigma_{\text{ex}} = \max_{\boldsymbol{\phi}\in\mathbb{R}^d} [-\boldsymbol{\phi}^\top \nabla^\top \boldsymbol{j} - \tilde{\boldsymbol{j}}^\top (e^{-\nabla\boldsymbol{\phi}} - \mathbf{1})], \tag{B5}$$

which is the analog of Eq. (33). The optimality condition that defines $\boldsymbol{\phi}^\star$ is given by

$$\nabla^\top \boldsymbol{j} = \nabla^\top (\tilde{\boldsymbol{j}} \circ e^{-\nabla\boldsymbol{\phi}^\star}), \tag{B6}$$

rather than Eq. (39). The expression of excess EPR as information-theoretic optimal transport, as in Eq. (45), remains unchanged. In the special case where Eq. (3) holds, Eq. (B5) reduces to Eq. (38), and Eq. (B6) reduces to Eq. (39).

Some results must be qualified in the presence of odd variables. For instance, in MJPs with odd variables, excess EPR does not necessarily vanish in stationarity, unless the steady state is symmetric under conjugation: $p_i^{\text{ss}} = p_{\epsilon i}^{\text{ss}}$. Similarly, we do not have the inequality between our excess EPR and HS excess EPR [Eq. (116)], unless the steady-state distribution is symmetric under conjugation. Finally, our thermodynamic speed limits do not hold in general for systems with odd variables. These differences arise because, for systems with odd variables and asymmetric escape rates, the steady state may be nonequilibrium even when the thermodynamic forces are conservative. The steady state is equilibrium if the forces are

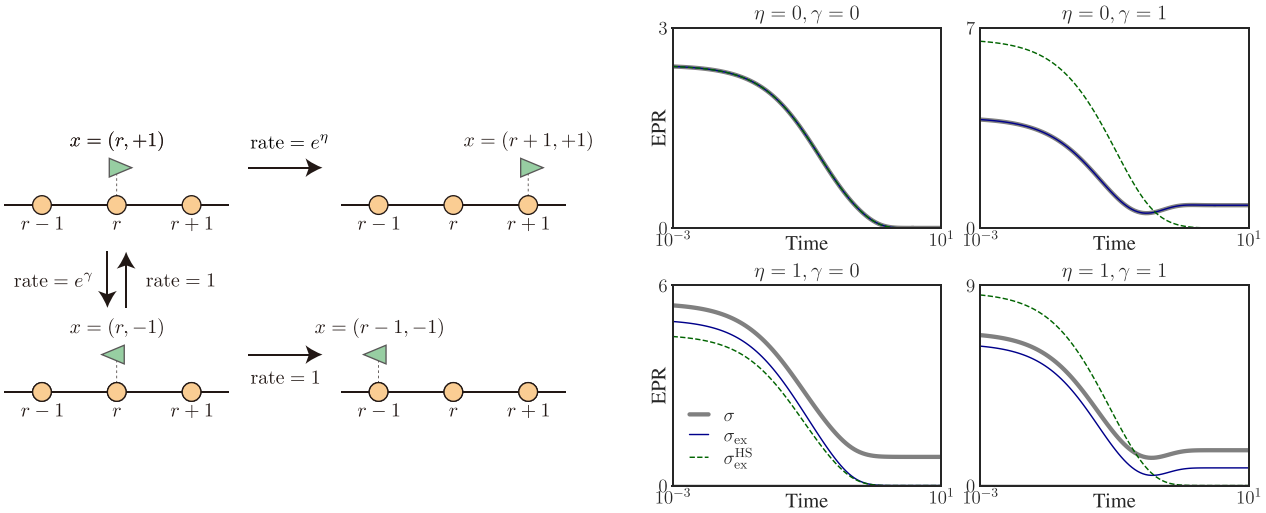

FIG. 9. Left: a discrete system with odd variables, consisting of a particle on a ring with position $r \in \{1, \ldots, k\}$ and odd velocity $v \in \{-1, +1\}$ [154,155,159]. Right: four panels showing time courses of EPR σ, our excess EPR σ_{ex}, and HS excess EPR $\sigma_{\mathrm{ex}}^{\mathrm{HS}}$ for different parameters. $\eta \neq 0$ means forces are nonconservative, and $\gamma \neq 0$ means steady-state distribution is not symmetric under conjugation of odd variables. HS decomposition can give unphysical values ($\sigma_{\mathrm{ex}}^{\mathrm{HS}} > \sigma$, $\sigma_{\mathrm{hk}}^{\mathrm{HS}} < 0$) when $\gamma \neq 0$.

conservative and the steady state is symmetric under conjugation of odd variables. (See Ref. [155] for further discussion.)

Odd variables are problematic for the HS decomposition, e.g., they can lead to negative values of HS housekeeping EPR [154,155,159]. To our knowledge, no universally applicable housekeeping/excess decomposition has been previously proposed for systems with odd variables.

3. Example: Particle on a ring

We provide an example to illustrate our excess/housekeeping decomposition in the presence of odd variables. We also compare to the HS decomposition, where the housekeeping EPR can take unphysical negative values [154,155,159].

We use an existing model from the literature [154,155,159], illustrated in Fig. 9 (left). There is a particle on a ring with k locations; the particle also has an odd "velocity" degree of freedom, indicating whether it is moving clockwise or counterclockwise. The system's microstate is specified by $i = (r, v)$, where $r \in \{1, \ldots, k\}$ is the position of the particle on the ring and $v \in \{-1, +1\}$ is the velocity. Conjugation involves flipping the sign of the velocity variables,

$$\epsilon(r, v) = (r, -v).$$

The rate matrix contains two types of jumps: around the ring $(r, v) \to (r + v, v)$ and velocity flips $(r, v) \to (r, -v)$.

The transition rates are parametrized as

$$R_{(r+v,v)\leftarrow(r,v)} = e^{\delta_{v,1}\eta},$$
$$R_{(r,-v)\leftarrow(r,v)} = e^{\delta_{v,1}\gamma}.$$

In words, the particle moves in the direction of its velocity with rate e^{η} when $v = +1$ and rate 1 when $v = -1$; the velocity flips with rate e^{γ} when $v = +1$ and rate 1 when $v = -1$.

The steady state is given by

$$p^{\mathrm{ss}}_{r,v} = \frac{\delta_{v,1} + \delta_{v,-1}e^{\gamma}}{k(e^{\gamma} + 1)}. \tag{B7}$$

The parameter η controls the strength of driving around the ring, leading to nonconservative forces when $\eta \neq 0$. The parameter γ controls the breaking of symmetry of velocity flips, leading to a steady-state distribution that is asymmetric under conjugation ($p^{\mathrm{ss}}_{r,v} \neq p^{\mathrm{ss}}_{r,-v}$) when $\gamma \neq 0$. The steady state is equilibrium only when $\eta = 0$ and $\gamma = 0$. Using Eq. (B4), the forces for the two types of transitions are

$$\begin{aligned} 1f_{(r+v,v)\leftarrow(r,v)} &= \ln \frac{p_{r,v}R_{(r+v,v)\leftarrow(r,v)}}{p_{r+v,v}R_{(r,-v)\leftarrow(r+v,-v)}} \\ &= \ln \frac{p_{r,v}}{p_{r+v,v}} + \delta_{v,1}\eta - \delta_{-v,1}\eta \\ &= \ln \frac{p_{r,v}}{p_{r+v,v}} + v\eta, \end{aligned} \tag{B8}$$

$$1f_{(r,-v)\leftarrow(r,v)} = \ln \frac{p_{r,v}R_{(r,-v)\leftarrow(r,v)}}{p_{r,-v}R_{(r,v)\leftarrow(r,-v)}} = \ln \frac{p_{r,v}}{p_{r,-v}}. \tag{B9}$$

In Fig. 9 (right), we visualize the time-dependent values of EPR σ, our excess EPR σ_{ex}, and the HS excess EPR $\sigma_{\mathrm{ex}}^{\mathrm{HS}}$. We consider a system with $k = 4$ positions and the non-steady-state initial distribution $p_{r,v} \propto 10\delta_{r,0}\delta_{v,1} + 1$. We consider four parameter values.

In the first condition (top left), $\eta = 0$ and $\gamma = 0$, so all transitions are symmetric. Here, the forces are conservative, $f = -\nabla\boldsymbol{\phi}$ for $\boldsymbol{\phi} = \ln \boldsymbol{p}$ (up to a constant) and the steady-state distribution is symmetric under conjugation. The steady state is in equilibrium and $\sigma = \sigma_{\mathrm{ex}} = \sigma_{\mathrm{ex}}^{\mathrm{HS}}$ at all times.

In the second condition (top right), $\eta = 0$ and $\gamma = 1$, so velocity flips $(r, -1) \to (r, +1)$ occur more frequently than $(r, +1) \to (r, -1)$. The forces are conservative, $f = -\nabla\boldsymbol{\phi}$ for $\boldsymbol{\phi} = \ln \boldsymbol{p}$, but the steady state is not symmetric under conjugation. For this reason, the steady state is nonequilibrium

($\sigma > 0$ in steady state). Since the forces are conservative, our housekeeping EPR vanishes and $\sigma = \sigma_{\rm ex}$ at all times. The HS decomposition gives different numerical values, and it can produce unphysical negative values ($\sigma_{\rm ex}^{\rm HS} > \sigma, \sigma_{\rm hk}^{\rm HS} < 0$).

In the third condition (bottom left), $\eta = 1$ and $\gamma = 0$, so movements along the ring with positive velocity are faster than those with negative velocity. The steady state is symmetric under time reversal but the forces $f_{(r+v,v)\leftarrow(r,v)}$ are not conservative, so the steady state is nonequilibrium. Our decomposition and HS decomposition both obey $0 \leqslant \sigma_{\rm ex} \leqslant \sigma$ and $0 \leqslant \sigma_{\rm ex}^{\rm HS} \leqslant \sigma$. We verify that, in systems with time-reversal-symmetric steady states, $\sigma_{\rm ex} \geqslant \sigma_{\rm ex}^{\rm HS}$ and $\sigma_{\rm ex} = \sigma_{\rm ex}^{\rm HS} = 0$ in steady state.

In the fourth condition (bottom right), $\eta = 1$ and $\gamma = 1$, so the forces are not conservative and the steady state is not symmetric under conjugation. The HS decomposition again gives unphysical values $\sigma_{\rm ex}^{\rm HS} > \sigma, \sigma_{\rm hk}^{\rm HS} < 0$. Under our decomposition, neither $\sigma_{\rm ex}$ nor $\sigma_{\rm hk}$ vanishes in steady state.

[1] J. Maas, Gradient flows of the entropy for finite Markov chains, J. Funct. Anal. **261**, 2250 (2011).

[2] A. Mielke, A gradient structure for reaction–diffusion systems and for energy-drift-diffusion systems, Nonlinearity **24**, 1329 (2011).

[3] F. Schlögl, Fluctuations in thermodynamic non equilibrium states, Z. Phys. A **244**, 199 (1971).

[4] H. Qian, Relative entropy: Free energy associated with equilibrium fluctuations and nonequilibrium deviations, Phys. Rev. E **63**, 042103 (2001).

[5] I. Procaccia and R. Levine, Potential work: A statistical-mechanical approach for systems in disequilibrium, J. Chem. Phys. **65**, 3357 (1976).

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

[7] N. Shiraishi, K. Funo, and K. Saito, Speed limit for classical stochastic processes, Phys. Rev. Lett. **121**, 070601 (2018).

[8] S. Ito, Stochastic thermodynamic interpretation of information geometry, Phys. Rev. Lett. **121**, 030605 (2018).

[9] V. T. Vo, T. Van Vu, and Y. Hasegawa, Unified approach to classical speed limit and thermodynamic uncertainty relation, Phys. Rev. E **102**, 062132 (2020).

[10] K. Yoshimura and S. Ito, Thermodynamic uncertainty relation and thermodynamic speed limit in deterministic chemical reaction networks, Phys. Rev. Lett. **127**, 160601 (2021).

[11] N. Shiraishi, Wasserstein distance in speed limit inequalities for Markov jump processes, J. Stat. Mech. (2024) 074003.

[12] J. S. Lee, S. Lee, H. Kwon, and H. Park, Speed limit for a highly irreversible process and tight finite-time Landauer's bound, Phys. Rev. Lett. **129**, 120603 (2022).

[13] P. Salamon, J. D. Nulton, and R. S. Berry, Length in statistical thermodynamics, J. Chem. Phys. **82**, 2433 (1985).

[14] P. Salamon and R. S. Berry, Thermodynamic length and dissipated availability, Phys. Rev. Lett. **51**, 1127 (1983).

[15] G. E. Crooks, Measuring thermodynamic length, Phys. Rev. Lett. **99**, 100602 (2007).

[16] D. A. Sivak and G. E. Crooks, Thermodynamic metrics and optimal paths, Phys. Rev. Lett. **108**, 190602 (2012).

[17] E. Aurell, C. Mejía-Monasterio, and P. Muratore-Ginanneschi, Optimal protocols and optimal transport in stochastic thermodynamics, Phys. Rev. Lett. **106**, 250601 (2011).

[18] 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).

[19] A. Dechant and Y. Sakurai, Thermodynamic interpretation of Wasserstein distance, arXiv:1912.08405.

[20] M. Nakazato and S. Ito, Geometrical aspects of entropy production in stochastic thermodynamics based on Wasserstein distance, Phys. Rev. Res. **3**, 043093 (2021).

[21] T. Van Vu and K. Saito, Thermodynamic unification of optimal transport: Thermodynamic uncertainty relation, minimum dissipation, and thermodynamic speed limits, Phys. Rev. X **13**, 011013 (2023).

[22] K. Yoshimura, A. Kolchinsky, A. Dechant, and S. Ito, Housekeeping and excess entropy production for general nonlinear dynamics, Phys. Rev. Res. **5**, 013017 (2023).

[23] R. Nagayama, K. Yoshimura, A. Kolchinsky, and S. Ito, Geometric thermodynamics of reaction-diffusion systems: Thermodynamic trade-off relations and optimal transport for pattern formation, Phys. Rev. Res. **7**, 033011 (2025).

[24] S. Ito, Geometric thermodynamics for the Fokker–Planck equation: Stochastic thermodynamic links between information geometry and optimal transport, Inf. Geom. **7**, 441 (2024).

[25] É. Fodor, R. L. Jack, and M. E. Cates, Irreversibility and biased ensembles in active matter: Insights from stochastic thermodynamics, Annu. Rev. Condens. Matter Phys. **13**, 215 (2022).

[26] R. Graham and T. Tél, Nonequilibrium potential for coexisting attractors, Phys. Rev. A **33**, 1322 (1986).

[27] M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems* (Springer Science & Business Media, 1998).

[28] B. Derrida, J. L. Lebowitz, and E. R. Speer, Free energy functional for nonequilibrium systems: An exactly solvable case, Phys. Rev. Lett. **87**, 150601 (2001).

[29] P. Ao, Potential in stochastic differential equations: Novel construction, J. Phys. A: Math. Gen. **37**, L25 (2004).

[30] J. Wang, L. Xu, and E. Wang, Potential landscape and flux framework of nonequilibrium networks: Robustness, dissipation, and coherence of biochemical oscillations, Proc. Natl. Acad. Sci. USA **105**, 12271 (2008).

[31] H. Ge and H. Qian, Dissipation, generalized free energy, and a self-consistent nonequilibrium thermodynamics of chemically driven open subsystems, Phys. Rev. E **87**, 062125 (2013).

[32] Q. Li and W. E, The free action of nonequilibrium dynamics, J. Stat. Phys. **161**, 300 (2015).

[33] X. Fang, K. Kruse, T. Lu, and J. Wang, Nonequilibrium physics in biology, Rev. Mod. Phys. **91**, 045004 (2019).

[34] P. Glansdorff and I. Prigogine, Non-equilibrium stability theory, Physica **46**, 344 (1970).

[35] Y. Oono and M. Paniconi, Steady state thermodynamics, Prog. Theor. Phys. Suppl. **130**, 29 (1998).

[36] T. Hatano and S.-I. Sasa, Steady-state thermodynamics of Langevin systems, Phys. Rev. Lett. **86**, 3463 (2001).

[37] T. Speck and U. Seifert, Integral fluctuation theorem for the housekeeping heat, J. Phys. A: Math. Gen. **38**, L581 (2005).

[38] M. Esposito, U. Harbola, and S. Mukamel, Entropy fluctuation theorems in driven open systems: Application to electron counting statistics, Phys. Rev. E **76**, 031132 (2007).

[39] T. S. Komatsu, N. Nakagawa, S.-I. Sasa, and H. Tasaki, Steady-state thermodynamics for heat conduction: Microscopic derivation, Phys. Rev. Lett. **100**, 230602 (2008).

[40] T. Sagawa and H. Hayakawa, Geometrical expression of excess entropy production, Phys. Rev. E **84**, 051110 (2011).

[41] C. Maes and K. Netočný, A nonequilibrium extension of the Clausius heat theorem, J. Stat. Phys. **154**, 188 (2014).

[42] E. Smith, Intrinsic and extrinsic thermodynamics for stochastic population processes with multi-level large-deviation structure, Entropy **22**, 1137 (2020).

[43] J. Keizer, Heat, work, and the thermodynamic temperature at nonequilibrium steady states, J. Chem. Phys. **82**, 2751 (1985).

[44] G. Falasco and M. Esposito, Macroscopic stochastic thermodynamics, Rev. Mod. Phys. **97**, 015002 (2025).

[45] M. Esposito and C. Van den Broeck, Three detailed fluctuation theorems, Phys. Rev. Lett. **104**, 090601 (2010).

[46] R. Rao and M. Esposito, Nonequilibrium thermodynamics of chemical reaction networks: Wisdom from stochastic thermodynamics, Phys. Rev. X **6**, 041064 (2016).

[47] H. Ge and H. Qian, Nonequilibrium thermodynamic formalism of nonlinear chemical reaction systems with Waage–Guldberg's law of mass action, Chem. Phys. **472**, 241 (2016).

[48] U. Seifert, From stochastic thermodynamics to thermodynamic inference, Annu. Rev. Condens. Matter Phys. **10**, 171 (2019).

[49] T. Van Vu and K. Saito, Topological speed limit, Phys. Rev. Lett. **130**, 010402 (2023).

[50] S.-I. Amari, *Information Geometry and its Applications* (Springer, 2016), Vol. 194.

[51] T. J. Kobayashi, D. Loutchko, A. Kamimura, and Y. Sughiyama, Hessian geometry of nonequilibrium chemical reaction networks and entropy production decompositions, Phys. Rev. Res. **4**, 033208 (2022).

[52] T. J. Kobayashi, D. Loutchko, A. Kamimura, S. A. Horiguchi, and Y. Sughiyama, Information geometry of dynamics on graphs and hypergraphs, Inf. Geom. **7**, 97 (2024).

[53] See Supplemental Material at http://link.aps.org/supplemental/10.1103/r48t-dghl for supporting derivations and additional numerical comparisons, which includes Refs. [143,160–163].

[54] M. Feinberg, *Foundations of Chemical Reaction Network Theory*, Applied Mathematical Sciences (Springer International Publishing, Cham, 2019), Vol. 202.

[55] D. Kondepudi and I. Prigogine, *Modern Thermodynamics: From Heat Engines to Dissipative Structures* (John Wiley & Sons, 2014).

[56] C. Maes, Local detailed balance, SciPost Phys. Lect. Notes **32** (2021).

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

[58] I. Prigogine and R. Lefever, Symmetry breaking instabilities in dissipative systems. II, J. Chem. Phys. **48**, 1695 (1968).

[59] F. Avanzini, E. Penocchio, G. Falasco, and M. Esposito, Nonequilibrium thermodynamics of non-ideal chemical reaction networks, J. Chem. Phys. **154**, 094114 (2021).

[60] D. A. Beard and H. Qian, Relationship between thermodynamic driving force and one-way fluxes in reversible processes, PLoS ONE **2**, e144 (2007).

[61] A. Wachtel, R. Rao, and M. Esposito, Thermodynamically consistent coarse graining of biocatalysts beyond Michaelis–Menten, New J. Phys. **20**, 042002 (2018).

[62] Y. Peng, H. Qian, D. A. Beard, and H. Ge, Universal relation between thermodynamic driving force and one-way fluxes in a nonequilibrium chemical reaction with complex mechanism, Phys. Rev. Res. **2**, 033089 (2020).

[63] F. Avanzini, G. Falasco, and M. Esposito, Thermodynamics of non-elementary chemical reaction networks, New J. Phys. **22**, 093040 (2020).

[64] S. Adams, N. Dirr, M. Peletier, and J. Zimmer, Large deviations and gradient flows, Philos. Trans. R. Soc. A **371**, 20120341 (2013).

[65] A. Mielke, D. R. M. Renger, and M. A. Peletier, On the relation between gradient flows and the large-deviation principle, with applications to Markov chains and diffusion, Potential Anal. **41**, 1293 (2014).

[66] R. C. Kraaij, A. Lazarescu, C. Maes, and M. Peletier, Fluctuation symmetry leads to GENERIC equations with non-quadratic dissipation, Stoch. Proc. Appl. **130**, 139 (2020).

[67] R. Rao and M. Esposito, Conservation laws shape dissipation, New J. Phys. **20**, 023007 (2018).

[68] H. Spohn, *Large Scale Dynamics of Interacting Particles* (Springer, Berlin, Heidelberg, 1991).

[69] A. Mielke, R. I. A. Patterson, M. A. Peletier, and D. R. Michiel Renger, Non-equilibrium thermodynamical principles for chemical reactions with mass-action kinetics, SIAM J. Appl. Math. **77**, 1562 (2017).

[70] To derive this result, observe that $\sigma = \boldsymbol{j}^\top \boldsymbol{f} = -\dot{\boldsymbol{x}}^\top \boldsymbol{\phi}^{\mathrm{eq}}$, since $\boldsymbol{f} = -\nabla \boldsymbol{\phi}^{\mathrm{eq}}$ and $\dot{\boldsymbol{x}} = \nabla^\top \boldsymbol{j}$.

[71] H. Ge and H. Qian, Mesoscopic kinetic basis of macroscopic chemical thermodynamics: A mathematical theory, Phys. Rev. E **94**, 052150 (2016).

[72] H. Touchette, The large deviation approach to statistical mechanics, Phys. Rep. **478**, 1 (2009).

[73] D. F. Anderson, G. Craciun, M. Gopalkrishnan, and C. Wiuf, Lyapunov functions, stationary distributions, and non-equilibrium potential for reaction networks, Bull. Math. Biol. **77**, 1744 (2015).

[74] H. Ge and H. Qian, Mathematical formalism of nonequilibrium thermodynamics for nonlinear chemical reaction systems with general rate law, J. Stat. Phys. **166**, 190 (2017).

[75] H. Gang, Stationary solution of master equations in the large-system-size limit, Phys. Rev. A **36**, 5782 (1987).

[76] J. Keizer, *Statistical Thermodynamics of Nonequilibrium Processes* (Springer, New York, NY, 1987).

[77] H. Ge and H. Qian, Non-equilibrium phase transition in mesoscopic biochemical systems: From stochastic to nonlinear dynamics and beyond, J. R. Soc. Interface **8**, 107 (2011).

[78] E. Weinan, W. Ren, and E. Vanden-Eijnden, Minimum action method for the study of rare events, Commun. Pure Appl. Math. **57**, 637 (2004).

[79] D. Dahiya and M. Cameron, Ordered line integral methods for computing the quasi-potential, J. Sci. Comput. **75**, 1351 (2018).

[80] Y. Li, S. Xu, J. Duan, X. Liu, and Y. Chu, A machine learning method for computing quasi-potential of stochastic dynamical systems, Nonlin. Dyn. **109**, 1877 (2022).

[81] R. Zakine and E. Vanden-Eijnden, Minimum-action method for nonequilibrium phase transitions, Phys. Rev. X **13**, 041044 (2023).

[82] P. Gagrani and E. Smith, Action functional gradient descent algorithm for estimating escape paths in stochastic chemical reaction networks, Phys. Rev. E **107**, 034305 (2023).

[83] B. C. Nolting and K. C. Abbott, Balls, cups, and quasi-potentials: Quantifying stability in stochastic systems, Ecology **97**, 850 (2016).

[84] H. Ge and H. Qian, Thermodynamic limit of a nonequilibrium steady state: Maxwell-type construction for a bistable biochemical system, Phys. Rev. Lett. **103**, 148103 (2009).

[85] A. Dechant, S.-I. Sasa, and S. Ito, Geometric decomposition of entropy production in out-of-equilibrium systems, Phys. Rev. Res. **4**, L012034 (2022).

[86] A. Dechant, S.-I. Sasa, and S. Ito, Geometric decomposition of entropy production into excess, housekeeping, and coupling parts, Phys. Rev. E **106**, 024125 (2022).

[87] D.-K. Kim, Y. Bae, S. Lee, and H. Jeong, Learning entropy production via neural networks, Phys. Rev. Lett. **125**, 140604 (2020).

[88] S. Otsubo, S. K. Manikandan, T. Sagawa, and S. Krishnamurthy, Estimating time-dependent entropy production from non-equilibrium trajectories, Commun. Phys. **5**, 11 (2022).

[89] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, 2004).

[90] N. Shiraishi and K. Saito, Information-theoretical bound of the irreversibility in thermal relaxation processes, Phys. Rev. Lett. **123**, 110603 (2019).

[91] I. Gentil, C. Léonard, and L. Ripani, About the analogy between optimal transport and minimal entropy, Ann. Fac. Sci. Toulouse: Math. **26**, 569 (2017).

[92] C. Villani, *Optimal Transport: Old and New* (Springer, 2009), Vol. 338.

[93] To derive Eq. (56) formally, recall that $f_\rho = -f_{\tilde{\rho}}$ by antisymmetry; hence, $\boldsymbol{f}$ belongs to the antisymmetric subspace $V = \{\boldsymbol{\theta} \in \mathbb{R}^m : \theta_\rho = -\theta_{\tilde{\rho}}\}$ of dimension $m/2$. The stoichiometric matrix also obeys the antisymmetry $\nabla_{\rho\cdot} = -\nabla_{\tilde{\rho}\cdot}$, therefore $\mathrm{Im}\,\nabla \subseteq V$. When $\mathrm{rank}\,\nabla = m/2$, we must have $\mathrm{Im}\,\nabla = V \ni \boldsymbol{f}$.

[94] M. Collins, R. E. Schapire, and Y. Singer, Logistic regression, AdaBoost and Bregman distances, Mach. Learn. **48**, 253 (2002).

[95] I. Csiszár, I-Divergence geometry of probability distributions and minimization problems, Ann. Probab. **3**, 146 (1975).

[96] S. Ito, M. Oizumi, and S.-I. Amari, Unified framework for the entropy production and the stochastic interaction based on information geometry, Phys. Rev. Res. **2**, 033048 (2020).

[97] The equation $\nabla^\top \boldsymbol{j} = -H\boldsymbol{\phi}^\star$ only determines $\boldsymbol{\phi}^\star$ up to the null-space of ∇. In the linear-response regime, the choice $\boldsymbol{\phi}^\star = -H^+(\boldsymbol{j})\nabla^\top \boldsymbol{j}$ satisfies the correct conservation laws.

[98] M. Doi, Onsager principle in polymer dynamics, Prog. Polym. Sci. **112**, 101339 (2021).

[99] L. Onsager, Reciprocal relations in irreversible processes. I., Phys. Rev. **37**, 405 (1931).

[100] I. Gyarmati, *Non-equilibrium Thermodynamics: Field Theory and Variational Principles* (Springer, Berlin Heidelberg, 1970).

[101] L. M. Martyushev and V. D. Seleznev, Maximum entropy production principle in physics, chemistry and biology, Phys. Rep. **426**, 1 (2006).

[102] J. Verhás, Gyarmati's variational principle of dissipative processes, Entropy **16**, 2362 (2014).

[103] For physical relevance, we should assume that short time intervals are longer than the hydrodynamic timescale, so that assumptions of Markovian and overdamped dynamics are justified.

[104] C. Maes, K. Netočný, and B. Wynants, On and beyond entropy production: The case of Markov jump processes, Markov Proc. Relat. Fields **14**, 445 (2008).

[105] A. Lazarescu, T. Cossetto, G. Falasco, and M. Esposito, Large deviations and dynamical phase transitions in stochastic chemical networks, J. Chem. Phys. **151**, 064117 (2019).

[106] S. R. S. Varadhan, *Large Deviations*, *Courant Lecture Notes in Mathematics*, Courant Institute of Mathematical Sciences (American Mathematical Society, New York, 2016), Vol. 27.

[107] S. K. Manikandan, D. Gupta, and S. Krishnamurthy, Inferring entropy production from short experiments, Phys. Rev. Lett. **124**, 120603 (2020).

[108] A. Dechant and S.-I. Sasa, Fluctuation–response inequality out of equilibrium, Proc. Natl. Acad. Sci. USA **117**, 6430 (2020).

[109] S. Otsubo, S. Ito, A. Dechant, and T. Sagawa, Estimating entropy production by machine learning of short-time fluctuating currents, Phys. Rev. E **101**, 062106 (2020).

[110] M. Aguilera, S. Ito, and A. Kolchinsky, Inferring entropy production in many-body systems using nonequilibrium maximum entropy, Phys. Rev. Lett. **136**, 077101 (2026).

[111] A. Dechant, Minimum entropy production, detailed balance and Wasserstein distance for continuous-time Markov processes, J. Phys. A: Math. Theor. **55**, 094001 (2022).

[112] R. Nagayama, K. Yoshimura, and S. Ito, Infinite variety of thermodynamic speed limits with general activities, Phys. Rev. Res. **7**, 013307 (2025).

[113] J. D. Orth, I. Thiele, and B. Ø. Palsson, What is flux balance analysis?, Nat. Biotechnol. **28**, 245 (2010).

[114] R. M. Fleming, C. M. Maes, M. A. Saunders, Y. Ye, and B. Ø. Palsson, A variational principle for computing nonequilibrium fluxes and potentials in genome-scale biochemical networks, J. Theor. Biol. **292**, 71 (2012).

[115] M. Kschischo, A gentle introduction to the thermodynamics of biochemical stoichiometric networks in steady state, Eur. Phys. J.: Spec. Top. **187**, 255 (2010).

[116] J. Piñero, R. Solé, and A. Kolchinsky, Optimization of nonequilibrium free energy harvesting illustrated on bacteriorhodopsin, Phys. Rev. Res. **6**, 013275 (2024).

[117] J.-C. Delvenne and G. Falasco, Thermokinetic relations, Phys. Rev. E **109**, 014109 (2024).

[118] A slightly weaker version of the Wasserstein TSL (83) was recently proposed for HS excess EPR in Ref. [117]. Unfortunately, the bound proposed in that work is not always valid; a counter-example may be found using the unicyclic system from Sec. VIII A.

[119] T. Van Vu and Y. Hasegawa, Geometrical bounds of the irreversibility in Markovian systems, Phys. Rev. Lett. **126**, 010601 (2021).

[120] R. Hamazaki, Speed limits for macroscopic transitions, PRX Quantum **3**, 020319 (2022).

[121] 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).

[122] Y.-Z. Zhen, D. Egloff, K. Modi, and O. Dahlsten, Universal bound on energy cost of bit reset in finite time, Phys. Rev. Lett. **127**, 190602 (2021).

[123] Y.-Z. Zhen, D. Egloff, K. Modi, and O. Dahlsten, Inverse linear versus exponential scaling of work penalty in finite-time bit reset, Phys. Rev. E **105**, 044147 (2022).

[124] J. Gu, Speed limit, dissipation bound, and dissipation-time trade-off in thermal relaxation processes, Phys. Rev. E **108**, L052103 (2023).

[125] J. O. Park, S. A. Rubin, Y.-F. Xu, D. Amador-Noguez, J. Fan, T. Shlomi, and J. D. Rabinowitz, Metabolite concentrations, fluxes and free energies imply efficient enzyme usage, Nat. Chem. Biol. **12**, 482 (2016).

[126] J. B. Russell and G. M. Cook, Energetics of bacterial growth: Balance of anabolic and catabolic reactions, Microbiol. Rev. **59**, 48 (1995).

[127] R. Heinrich, F. Montero, E. Klipp, T. G. Waddell, and E. Meléndez-Hevia, Theoretical approaches to the evolutionary optimization of glycolysis: Thermodynamic and kinetic constraints, Eur. J. Biochem. **243**, 191 (1997).

[128] E. Meléndez-Hevia, T. G. Waddell, R. Heinrich, and F. Montero, Theoretical approaches to the evolutionary optimization of glycolysis: Chemical analysis, Eur. J. Biochem. **244**, 527 (1997).

[129] M. Polettini and M. Esposito, Irreversible thermodynamics of open chemical networks. I. Emergent cycles and broken conservation laws, J. Chem. Phys. **141**, 024117 (2014).

[130] M. Kaiser, R. L. Jack, and J. Zimmer, Canonical structure and orthogonality of forces and currents in irreversible Markov chains, J. Stat. Phys. **170**, 1019 (2018).

[131] D. M. Renger, Gradient and GENERIC systems in the space of fluxes, applied to reacting particle systems, Entropy **20**, 596 (2018).

[132] D. R. M. Renger and J. Zimmer, Orthogonality of fluxes in general nonlinear reaction networks, Discrete Contin. Dyn. Syst. Ser. S **14**, 205 (2021).

[133] R. I. A. Patterson, D. R. M. Renger, and U. Sharma, Variational structures beyond gradient flows: A macroscopic fluctuation-theory perspective, J. Stat. Phys. **191**, 18 (2024).

[134] D. Mandal and C. Jarzynski, Analysis of slow transitions between nonequilibrium steady states, J. Stat. Mech. (2016) 063204.

[135] L. Onsager, Reciprocal relations in irreversible processes. II., Phys. Rev. **38**, 2265 (1931).

[136] L. Onsager and S. Machlup, Fluctuations and irreversible processes, Phys. Rev. **91**, 1505 (1953).

[137] A. Wachtel, R. Rao, and M. Esposito, Free-energy transduction in chemical reaction networks: From enzymes to metabolism, J. Chem. Phys. **157**, 024109 (2022).

[138] P. Gagrani, N. Lauber, E. Smith, and C. Flamm, Thermodynamic ranking of pathways in reaction networks, J. Stat. Mech. (2025) 123208.

[139] K. Sugie, D. Loutchko, and T. J. Kobayashi, Transitions and thermodynamics on species graphs of chemical reaction networks, Phys. Rev. E **112**, 044112 (2025).

[140] A. Baradat and C. Léonard, Minimizing relative entropy of path measures under marginal constraints, arXiv:2001.10920.

[141] O. Movilla Miangolarra, A. Eldesoukey, and T. T. Georgiou, Inferring potential landscapes: A Schrödinger bridge approach to maximum caliber, Phys. Rev. Res. **6**, 033070 (2024).

[142] P. D. Dixit, J. Wagoner, C. Weistuch, S. Pressé, K. Ghosh, and K. A. Dill, Perspective: Maximum caliber is a general variational principle for dynamical systems, J. Chem. Phys. **148**, 010901 (2018).

[143] K. Yoshimura and S. Ito, Two applications of stochastic thermodynamics to hydrodynamics, Phys. Rev. Res. **6**, L022057 (2024).

[144] K. Yoshimura, Y. Maekawa, R. Nagayama, and S. Ito, Force-current structure in Markovian open quantum systems and its applications: Geometric housekeeping-excess decomposition and thermodynamic trade-off relations, Phys. Rev. Res. **7**, 013244 (2025).

[145] M. Malek-Mansour and G. Nicolis, A master equation description of local fluctuations, J. Stat. Phys. **13**, 197 (1975).

[146] G. Nicolis and I. Prigogine, *Self-Organization in Nonequilibrium Systems: From Dissipative Structures to Order Through Fluctuations*, 1st ed. (Wiley, New York, 1977).

[147] J. Korbel and D. H. Wolpert, Stochastic thermodynamics and fluctuation theorems for non-linear systems, New J. Phys. **23**, 033049 (2021).

[148] T. D. Frank, Strongly nonlinear stochastic processes in physics and the life sciences, Int. Scholarly Res. Not. **2013**, 149169 (2013).

[149] A. E. Allahverdyan and A. Galstyan, Le Chatelier's principle in replicator dynamics, Phys. Rev. E **84**, 041117 (2011).

[150] D. D. Patterson, S. A. Levin, C. Staver, and J. D. Touboul, Probabilistic foundations of spatial mean-field models in ecology and applications, SIAM J. Appl. Dyn. Syst. **19**, 2682 (2020).

[151] S. Feng and X. Zheng, Solutions of a class of nonlinear master equations, Stoch. Proc. Appl. **43**, 65 (1992).

[152] V. N. Kolokoltsov, J. Li, and W. Yang, Mean field games and nonlinear Markov processes, arXiv:1112.3744.

[153] R. E. Spinney and I. J. Ford, Entropy production in full phase space for continuous stochastic dynamics, Phys. Rev. E **85**, 051113 (2012).

[154] R. E. Spinney and I. J. Ford, Nonequilibrium thermodynamics of stochastic systems with odd and even variables, Phys. Rev. Lett. **108**, 170603 (2012).

[155] H. K. Lee, C. Kwon, and H. Park, Fluctuation theorems and entropy production with odd-parity variables, Phys. Rev. Lett. **110**, 050602 (2013).

[156] C. Gardiner, *Handbook of Stochastic Methods: For Physics, Chemistry and the Natural Sciences*, 3rd ed. (Springer, Berlin, 2004).

[157] R. E. Spinney, The use of stochastic methods to explore the thermal equilibrium distribution and define entropy production out of equilibrium, Ph.D. thesis, UCL (University College London, 2012).

[158] F. Liu and H. Lei, Splitting of the rate matrix as a definition of time reversal in master equation systems, J. Phys. A: Math. Theor. **45**, 125004 (2012).

[159] I. J. Ford and R. E. Spinney, Entropy production from stochastic dynamics in discrete full phase space, Phys. Rev. E **86**, 021127 (2012).

[160] X. Li, F. Wu, F. Qi, and D. A. Beard, A database of thermodynamic properties of the reactions of glycolysis, the tricarboxylic acid cycle, and the pentose phosphate pathway, Database **2011**, bar005 (2011).

[161] H. L. Messiha, E. Kent, N. Malys, K. M. Carroll, N. Swainston, P. Mendes, and K. Smallbone, Enzyme characterisation and kinetic modelling of the pentose phosphate pathway in yeast, PeerJ PrePrints 2:e146v4.

[162] P. Raux, C. Goupil, and G. Verley, Thermodynamic circuits: Modeling chemical reaction networks with nonequilibrium conductance matrices, Phys. Rev. E **112**, 034112 (2025).

[163] H. Vroylandt, D. Lacoste, and G. Verley, Degree of coupling and efficiency of energy converters far-from-equilibrium, J. Stat. Mech. (2018) 023205.