Inferring Entropy Production in Many-Body Systems Using Nonequilibrium Maximum Entropy

Miguel Aguilera,[1,2,*,†] Sosuke Ito,[3,4] and Artemy Kolchinsky[5,3,*,‡]
[1]BCAM—Basque Center for Applied Mathematics, 48009 Bilbao, Spain
[2]IKERBASQUE, Basque Foundation for Science, 48009 Bilbao, Spain
[3]Universal Biology Institute, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
[4]Department of Physics, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
[5]ICREA-Complex Systems Lab, Universitat Pompeu Fabra, 08003 Barcelona, Spain

(Received 6 June 2025; revised 10 September 2025; accepted 18 December 2025; published 18 February 2026)

We propose a method for inferring entropy production (EP) in high-dimensional stochastic systems, including many-body systems and non-Markovian systems with long memory. Standard techniques for estimating EP become intractable in such systems due to computational and statistical limitations. We infer trajectory-level EP and lower bounds on average EP by exploiting a nonequilibrium analogue of the maximum entropy principle, along with convex duality. Our approach uses only samples of trajectory observables, such as spatiotemporal correlations. It does not require reconstruction of high-dimensional probability distributions or rate matrices, nor impose any special assumptions such as discrete states or multipartite dynamics. In addition, it may be used to compute a hierarchical decomposition of EP, reflecting contributions from different interaction orders, and it has an intuitive physical interpretation as a "thermodynamic uncertainty relation." We demonstrate its numerical performance on a disordered nonequilibrium spin model with 1000 spins and a large neural spike-train dataset.

DOI: 10.1103/xgkj-dxzh

The central quantity of interest in nonequilibrium thermodynamics is entropy production (EP) [1]. In microscopic physical systems, EP characterizes departure from thermodynamic equilibrium and quantifies the dissipation of thermodynamic free energy. More generally, EP provides an information-theoretic measure of temporal irreversibility, including in meso- and macroscopic systems [2–5].

Recently, there has been growing interest in thermodynamic inference [6], that is, the problem of estimating EP from empirical measurements of a stochastic system. Various methods have been developed for inferring EP from partial [7–14] information. For instance, the celebrated thermodynamic uncertainty relation (TUR) [15–18] bounds EP in terms of the mean and variance of a single current. Other techniques relate EP to the statistics of waiting times [19–22], observed transitions [11,12,23,24], and counting observables [25]. Many of these methods are designed to infer underlying dissipation from a small number of coarse-grained observables.

In this Letter, we consider thermodynamic inference in the high-dimensional setting, where a large number of observables are available. This is particularly relevant for multivariate measurements of complex systems—such as nonequilibrium disordered networks [26], biological active matter [27,28], and neural systems [29–33]—that have many degrees of freedom and/or long non-Markovian memory. A naïve approach requires the estimation of trajectory probability distributions, but in high dimensions, this is usually statistically and numerically infeasible.

To address this challenge, we propose to infer EP using an information-theoretic variational principle, which can be understood as the nonequilibrium analogue of the maximum entropy (MaxEnt) principle in statistical physics. Our variational principle has a simple physical interpretation: it quantifies the irreversibility captured by the expectations of some (possibly large) number of trajectory observables. Importantly, this variational principle has a dual form that leads to a tractable convex optimization problem. This optimization problem gives a lower bound on EP and an estimate of the fluctuating trajectory-level EP—directly from trajectory samples and without explicit use of trajectory probabilities.

As we show, our bound can also be understood as a higher-order TUR [34,35], which we use to derive further simple bounds on EP. Finally, we show that in multipartite systems, our variational principle can be split into smaller subproblems, greatly improving the performance scaling of our method with system size.

*These authors contributed equally to this work.
†Contact author: maguilera@bcamath.org
‡Contact author: artemyk@gmail.com

Entropy production—We consider a nonequilibrium stochastic system either in continuous or discrete time. For time $t \in [0, T]$, the system follows trajectories $\boldsymbol{x}$ according to the "forward" probability distribution $p(\boldsymbol{x})$. We write $\boldsymbol{x}_t$ for the system's state at time t, and $x_{i,t}$ for the state of degree of freedom i at t. As standard in stochastic thermodynamics, the system is also associated with a "reverse" trajectory distribution $\tilde{p}(\boldsymbol{x})$ produced by sampling from the final distribution of the forward process $p(\boldsymbol{x}_T)$, applying time-dependent driving in reverse, and finally time-reversing trajectories [1]. For simplicity, here we mostly focus on the common case of stationary systems without odd variables (such as velocity). In such a system, the reverse distribution is given simply by time-reversing forward trajectories, $\tilde{p}(\boldsymbol{x}) = p(\tilde{\boldsymbol{x}})$ where $\tilde{\boldsymbol{x}}_t \coloneqq \boldsymbol{x}_{T-t}$.

The trajectory EP of trajectory $\boldsymbol{x}$ and the average EP across all trajectories are defined in terms of symmetry breaking between forward and reverse distributions [1],

$$\sigma(\boldsymbol{x}) \coloneqq \ln \frac{p(\boldsymbol{x})}{\tilde{p}(\boldsymbol{x})}, \qquad \Sigma \coloneqq D(p\|\tilde{p}) = \left\langle \ln \frac{p(\boldsymbol{x})}{\tilde{p}(\boldsymbol{x})} \right\rangle_p, \tag{1}$$

where $\langle \ldots \rangle_p = \sum_{\boldsymbol{x}} \ldots p(\boldsymbol{x})$ is the expectation over p and $D(\cdot\|\cdot)$ is the Kullback-Leibler (KL) divergence. EP vanishes in equilibrium, when trajectory statistics are indistinguishable under the forward and reverse distributions.

EP bound—We suppose that one measures trajectory observables, represented by a vector-valued function $\boldsymbol{g}(\boldsymbol{x}) \in \mathbb{R}^d$, from the forward and reverse process. Natural choices of observables include correlation functions, e.g., two-point $(x_{i,t}x_{j,t'})$ spatiotemporal correlations. We do not assume antisymmetric observables such as $\boldsymbol{g}(\boldsymbol{x}) = -\boldsymbol{g}(\tilde{\boldsymbol{x}})$.

Our goal is to estimate $\sigma(\boldsymbol{x})$ and Σ from forward and reverse samples of observables $\boldsymbol{g}$. The overall irreversibility encoded in these observables is quantified by the KL divergence between the distributions of $\boldsymbol{g}$ under the forward and reverse processes, $\Sigma_{\boldsymbol{g}}^{\mathrm{DPI}} \coloneqq D(p_G\|\tilde{p}_G)$, where $p_G(\boldsymbol{g}') = \langle \delta_{\boldsymbol{g}',\boldsymbol{g}(\boldsymbol{x})} \rangle_p$ and $\tilde{p}_G(\boldsymbol{g}') = \langle \delta_{\boldsymbol{g}',\boldsymbol{g}(\boldsymbol{x})} \rangle_{\tilde{p}}$. The data processing inequality (DPI) implies that $\Sigma_{\boldsymbol{g}}^{\mathrm{DPI}} \leq \Sigma$, with equality when $\boldsymbol{g}(\boldsymbol{x})$ is an invertible function [36].

Estimating $\Sigma_{\boldsymbol{g}}^{\mathrm{DPI}}$ is very difficult for high-dimensional observables $\boldsymbol{g}$. Instead, we bound EP using a nonequilibrium generalization of the MaxEnt variational principle. Specifically, we choose the distribution that minimizes KL divergence to the reverse process while matching the forward-process expectations of the observables:

$$\Sigma_{\boldsymbol{g}} \coloneqq \min_q D(q\|\tilde{p}) \text{ subject to } \langle \boldsymbol{g} \rangle_q = \langle \boldsymbol{g} \rangle_p. \tag{2}$$

This quantity obeys the bounds $0 \leq \Sigma_{\boldsymbol{g}} \leq \Sigma_{\boldsymbol{g}}^{\mathrm{DPI}}$, as shown in Supplemental Material (SM) [37]. $\Sigma_{\boldsymbol{g}}$ quantifies the irreversibility captured by observable expectations, and it vanishes if and only if the expectation of $\boldsymbol{g}$ is the same under the forward and reverse processes. Large-deviations theory gives $\Sigma_{\boldsymbol{g}}$ a physical interpretation in terms of fluctuations: given n sample trajectories from the reverse process, the probability that the empirical average of observables $\boldsymbol{g}$ is equal to $\langle \boldsymbol{g} \rangle_p$ scales as $\asymp e^{-n\Sigma_{\boldsymbol{g}}}$ [41].

Importantly, by exploiting a remarkable information-theoretic duality, we may calculate $\Sigma_{\boldsymbol{g}}$ without optimizing or inferring any probability distributions. As shown in SM [37], the optimization problem (2) has the dual formulation

$$\Sigma_{\boldsymbol{g}} = \max_{\boldsymbol{\theta} \in \mathbb{R}^d} \left(\boldsymbol{\theta}^\top \langle \boldsymbol{g} \rangle_p - \ln \langle e^{\boldsymbol{\theta}^\top \boldsymbol{g}} \rangle_{\tilde{p}} \right). \tag{3}$$

Equation (3) is an unconstrained convex optimization problem over Lagrangian multipliers $\boldsymbol{\theta}$ that enforce the d expectation constraints on $\boldsymbol{g}$. The objective depends only on the expectations of $\boldsymbol{g}$ and $e^{\boldsymbol{\theta}^\top \boldsymbol{g}}$ under the forward and backward processes, which can be easily computed from empirical samples of $\boldsymbol{g}$. In the special case of antisymmetric observables and a steady-state system without odd variables, Eq. (3) can be estimated using only forward samples since $\ln \langle e^{\boldsymbol{\theta}^\top \boldsymbol{g}} \rangle_{\tilde{p}} = \ln \langle e^{-\boldsymbol{\theta}^\top \boldsymbol{g}} \rangle_p$. As we discuss in more detail below, a related (but different) variational bound on EP was proposed in Refs. [42,43].

As an example, consider a system with 1000 binary spins measured at two time points, with $\boldsymbol{g}$ encoding two-point correlations $x_{i,t}x_{j,t'}$. Equation (3) is an unconstrained convex optimization problem over 1000^2 variables, which can be solved using standard numerical techniques. Conversely, a naïve estimate of Σ or $\Sigma_{\boldsymbol{g}}^{\mathrm{DPI}}$ requires inference of $> 2^{1000}$ probabilities.

So far, we have made no assumptions regarding multipartite structure. Here, we say that the observables are "multipartite" if they can be partitioned into blocks, such that only a single block can be active in any given trajectory. In practice, this situation often arises when observables depend on local subsystems and only one subsystem can change state at a given time. We provide a formal definition of multipartite observables in the *End Matter*. There, we show that multipartiteness allows Eq. (3) to be split into smaller optimization problems, each involving a proportional fraction of variables and samples, dramatically reducing the computation and memory required for optimization.

Maximum likelihood EP decomposition—As we show in SM [37], the distribution q^* that optimizes Eq. (2) belongs to an exponential family,

$$q_{\boldsymbol{\theta}}(\boldsymbol{x}) = \tilde{p}(\boldsymbol{x}) e^{\boldsymbol{\theta}^\top \boldsymbol{g}(\boldsymbol{x}) - \ln \langle e^{\boldsymbol{\theta}^\top \boldsymbol{g}} \rangle_{\tilde{p}}}. \tag{4}$$

Specifically, $q^* = q_{\boldsymbol{\theta}^*}$ for optimal parameters $\boldsymbol{\theta}^*$ that assign maximum likelihood (ML) to forward trajectories:

$$\boldsymbol{\theta}^* = \operatorname*{argmax}_{\boldsymbol{\theta} \in \mathbb{R}^d} \langle \ln q_{\boldsymbol{\theta}}(\boldsymbol{x}) \rangle_p = \operatorname*{argmin}_{\boldsymbol{\theta} \in \mathbb{R}^d} D(p\|q_{\boldsymbol{\theta}}). \tag{5}$$

This also leads to the information-geometric Pythagorean theorem [44], which decomposes EP into two nonnegative terms:

$$\underbrace{D(p\|\tilde{p})}_{\Sigma} = \underbrace{D(q^*\|\tilde{p})}_{\Sigma_{\boldsymbol{g}}} + \underbrace{D(p\|q^*)}_{\Sigma_{\boldsymbol{g}}^{\perp}}. \tag{6}$$

$\Sigma_{\boldsymbol{g}}^{\perp} := \Sigma - \Sigma_{\boldsymbol{g}} = D(p\|q^*)$ quantifies the EP not captured by $\Sigma_{\boldsymbol{g}}$, expressed via the ML problem $\Sigma_{\boldsymbol{g}}^{\perp} = \min_{\boldsymbol{\theta}} D(p\|q_{\boldsymbol{\theta}})$.

We may further decompose $\Sigma_{\boldsymbol{g}}^{\perp}$ by using that Eq. (4) implies that q^* obeys $q^*_{X|G}(\boldsymbol{x}|\boldsymbol{g}') = \tilde{p}_{X|G}(\boldsymbol{x}|\boldsymbol{g}')$, where conditional probabilities are computed using Bayes' rule, $\tilde{p}_{X|G}(\boldsymbol{x}|\boldsymbol{g}') = \tilde{p}(\boldsymbol{x})\delta_{\boldsymbol{g}',\boldsymbol{g}(\boldsymbol{x})}/\tilde{p}_G(\boldsymbol{g}')$. Using the chain rule of KL divergence, we may write $\Sigma_{\boldsymbol{g}}^{\perp}$ as a sum of two terms:

$$\Sigma_{\boldsymbol{g}}^{\perp} = \underbrace{D(p_G\|q^*_G)}_{\Sigma_{\boldsymbol{g}}^{\mathrm{DPI}} - \Sigma_{\boldsymbol{g}}} + \underbrace{D(p_{X|G}\|\tilde{p}_{X|G})}_{\Sigma - \Sigma_{\boldsymbol{g}}^{\mathrm{DPI}}}. \tag{7}$$

The first term is the additional EP that could be inferred by constraining all statistics (not just the mean) of $\boldsymbol{g}$ in Eq. (2). The second term is the EP that cannot be inferred from *any* statistics of $\boldsymbol{g}$.

Our method also supports a hierarchical decomposition of EP, where contributions from interactions of increasing order (e.g., singletons, pairs, triplets) are successively added. This gives a sequence of bounds $0 \le \Sigma_1 \le \Sigma_2 \le \ldots \le \Sigma$, with $\Sigma_k = D(q_k^*\|\tilde{p})$ capturing contributions from functions $\boldsymbol{g}(\boldsymbol{x})$ containing interactions up to order k. This structure mirrors information-geometric decompositions in equilibrium MaxEnt [44,45] and allows breaking down EP into interpretable contributions, each satisfying $\Sigma_k - \Sigma_{k-1} = D(q_k^*\|q_{k-1}^*)$.

Finally, our method gives an estimate of trajectory EP,

$$\sigma_{\boldsymbol{\theta}^*}(\boldsymbol{x}) := \ln\frac{q_{\boldsymbol{\theta}^*}(\boldsymbol{x})}{\tilde{p}(\boldsymbol{x})} = \boldsymbol{\theta}^{*\top}\boldsymbol{g}(\boldsymbol{x}) - \ln\langle e^{\boldsymbol{\theta}^{*\top}\boldsymbol{g}}\rangle_{\tilde{p}}. \tag{8}$$

This estimator satisfies $\langle\sigma_{\boldsymbol{\theta}^*}\rangle_p = \Sigma_{\boldsymbol{g}}$, and it may be shown to provide the best approximation of trajectory EP $\sigma(\boldsymbol{x})$ within the exponential family $q_{\boldsymbol{\theta}}$. Specifically, if we use $\sigma_{\boldsymbol{\theta}}(\boldsymbol{x}) = \ln[q_{\boldsymbol{\theta}}(\boldsymbol{x})/\tilde{p}(\boldsymbol{x})]$ to indicate the approximation given by parameters $\boldsymbol{\theta}$ and $\delta\sigma_{\boldsymbol{\theta}}(\boldsymbol{x}) = \sigma(\boldsymbol{x}) - \sigma_{\boldsymbol{\theta}}(\boldsymbol{x})$ the corresponding residual, then $\sigma_{\boldsymbol{\theta}^*}$ minimizes the expected error [46]:

$$D(p\|q_{\boldsymbol{\theta}}) = \langle\delta\sigma_{\boldsymbol{\theta}} + e^{-\delta\sigma_{\boldsymbol{\theta}}} - 1\rangle_p. \tag{9}$$

Here, $\delta\sigma_{\boldsymbol{\theta}} + e^{-\delta\sigma_{\boldsymbol{\theta}}} - 1 \ge 0$ is an information-theoretic loss function that converges to squared error $(\delta\sigma_{\boldsymbol{\theta}})^2/2$ as $\delta\sigma_{\boldsymbol{\theta}} \to 0$.

The trajectory EP estimator is exact [$\sigma(\boldsymbol{x}) = \sigma_{\boldsymbol{\theta}^*}(\boldsymbol{x})$ for all $\boldsymbol{x}$] if and only if $\Sigma_{\boldsymbol{g}} = \Sigma$, since then $p = q^*$. In this case, if the system is in steady state and without odd variables, we may exploit antisymmetry $\sigma(\boldsymbol{x}) = -\sigma(\tilde{\boldsymbol{x}})$ to simplify Eq. (8) as

$$\sigma(\boldsymbol{x}) = \sigma_{\boldsymbol{\theta}^*}(\boldsymbol{x}) = \boldsymbol{\theta}^{*\top}[\boldsymbol{g}(\boldsymbol{x}) - \boldsymbol{g}(\tilde{\boldsymbol{x}})]/2, \tag{10}$$

which simplifies further to $\boldsymbol{\theta}^{*\top}\boldsymbol{g}(\boldsymbol{x})$ when $\boldsymbol{g}$ is antisymmetric.

Thermodynamic uncertainty relations (TURs)—Equation (3) can be interpreted as a TUR that relates EP and fluctuations of trajectory observables [47]. Consider any trajectory observable that can be expressed as a linear combination $o(\boldsymbol{x}) = \boldsymbol{\theta}^{\top}\boldsymbol{g}(\boldsymbol{x})$ for some $\boldsymbol{\theta} \in \mathbb{R}^d$. The first term in Eq. (3) is the expectation of $o(\boldsymbol{x})$, while the second term is the cumulant generating function (CGF) of $o(\boldsymbol{x})$ under $\tilde{p}$. Thus, $\Sigma_{\boldsymbol{g}}$ bounds the fluctuation-discounted expectation of all such $o(\boldsymbol{x})$, and it can be understood as a higher-order TUR [34] that constrains all cumulants of trajectory observables. This contrasts with quadratic TURs, which constrain only the mean and variance [18,48,49].

To relate our approach to quadratic TURs, we expand the CGF as $\ln\langle e^{\boldsymbol{\theta}^{\top}\boldsymbol{g}}\rangle_{\tilde{p}} \approx \boldsymbol{\theta}^{\top}\langle\boldsymbol{g}\rangle_{\tilde{p}} + \boldsymbol{\theta}^{\top}\mathsf{K}_{\tilde{p}}\boldsymbol{\theta}/2$, where $\mathsf{K}_{\tilde{p}}$ is the covariance matrix of $\boldsymbol{g}$ under $\tilde{p}$. Plugging into Eq. (3) gives the approximate optimizer $\hat{\boldsymbol{\theta}}$ as the solution to the linear system $\mathsf{K}_{\tilde{p}}\hat{\boldsymbol{\theta}} = \langle\boldsymbol{g}\rangle_p - \langle\boldsymbol{g}\rangle_{\tilde{p}} \equiv \langle\boldsymbol{g}\rangle_{p-\tilde{p}}$, giving the weaker bound

$$\Sigma_{\boldsymbol{g}} \ge \hat{\Sigma}_{\boldsymbol{g}} := \langle\boldsymbol{g}\rangle_{p-\tilde{p}}^{\top}\mathsf{K}_{\tilde{p}}^{-1}\langle\boldsymbol{g}\rangle_p - \ln\left\langle e^{\langle\boldsymbol{g}\rangle_{p-\tilde{p}}^{\top}\mathsf{K}_{\tilde{p}}^{-1}\boldsymbol{g}}\right\rangle_{\tilde{p}}. \tag{11}$$

$\hat{\Sigma}_{\boldsymbol{g}}$ can be shown to correspond to a single Newton-Raphson step for maximizing Eq. (3) starting from $\boldsymbol{\theta} = \boldsymbol{0}$ [37]. This bound has the advantage of not requiring optimization, only the solution of a linear system. The inequality $\Sigma_{\boldsymbol{g}} \ge \hat{\Sigma}_{\boldsymbol{g}}$ is tight when $\boldsymbol{g}$ has Gaussian statistics.

We may also compare our results to existing TURs. We consider the special case of a steady-state system without odd variables and antisymmetric $\boldsymbol{g}$. In this case, EP may be bounded as (see End Matter)

$$\Sigma \ge \Sigma_{\boldsymbol{g}}^{\mathrm{TUR}} := \ln(1 + 2\langle\boldsymbol{g}\rangle_p^{\top}\mathsf{K}_p^{-1}\langle\boldsymbol{g}\rangle_p). \tag{12}$$

This bound approaches the usual quadratic form, $\Sigma_{\boldsymbol{g}}^{\mathrm{TUR}} \simeq 2\langle\boldsymbol{g}\rangle_p^{\top}\mathsf{K}_p^{-1}\langle\boldsymbol{g}\rangle_p$ for sufficiently small $\langle\boldsymbol{g}\rangle_p^{\top}\mathsf{K}_p^{-1}\langle\boldsymbol{g}\rangle_p$. This regime holds in the short-time limit, since $\langle\boldsymbol{g}\rangle_p$ scales with the observation period.

The bounds (11) and (12) may require solving linear systems involving very large covariance matrices. For multipartite observables, it is often possible to split these bounds into subproblems that involve smaller matrices (see End Matter).

Related work—In the machine learning literature, Belghazi *et al.* [50] proposed an estimator of KL divergence inspired by the "Donsker-Varadhan" variational representation [51]. When applied to $\Sigma = D(p\|\tilde{p})$, it leads to the lower bound

$$\Sigma_{\boldsymbol{g}}^{\mathrm{B}} := \max_{\boldsymbol{\phi}\in\mathbb{R}^k}\left(\langle v_{\boldsymbol{\phi}}\rangle_p - \ln\langle e^{v_{\boldsymbol{\phi}}}\rangle_{\tilde{p}}\right), \tag{13}$$

where $v_{\boldsymbol{\phi}}(\boldsymbol{x})$ is a nonlinear function parameterized by $\boldsymbol{\phi} \in \mathbb{R}^k$ (e.g., output of a neural network) [50]. $\Sigma_{\boldsymbol{g}}$ in Eq. (3) is a special case of $\Sigma_{\boldsymbol{g}}^{\mathrm{B}}$ where $v_{\boldsymbol{\phi}}(\boldsymbol{x})$ contains all linear combinations of

$\boldsymbol{g}(\boldsymbol{x})$. The neural-network-based bound (13) has both practical and conceptual differences with respect to Eq. (3). For instance, it generally involves a difficult nonconvex optimization problem, it does not provide a Pythagorean decomposition of EP, and it cannot be interpreted in terms of large deviations. Exploration of $\Sigma_{\boldsymbol{g}}^{\mathrm{B}}$ for thermodynamic inference is left as an interesting future direction.

A related variational expression for thermodynamic inference was developed in Kim *et al.* [42] and Otsubo *et al.* [43]. In our notation, it is written as

$$\Sigma_{\boldsymbol{g}}^{\mathrm{KO}} := \max_{\boldsymbol{\theta}} \left(\boldsymbol{\theta}^{\top} \langle \boldsymbol{g} \rangle_p - \langle e^{\boldsymbol{\theta}^{\top} \boldsymbol{g}} \rangle_{\tilde{p}} + 1 \right). \tag{14}$$

It is possible to consider a neural-network-based version of $\Sigma_{\boldsymbol{g}}^{\mathrm{KO}}$ by replacing $\boldsymbol{\theta}^{\top} \boldsymbol{g}(\boldsymbol{x})$ with a parameterized function $v_{\boldsymbol{\phi}}(\boldsymbol{x})$, as in Eq. (13), as explored in Refs. [42,52].

Note that $\Sigma_{\boldsymbol{g}} \geq \Sigma_{\boldsymbol{g}}^{\mathrm{KO}}$, since $-\ln x \geq -x + 1$ [50,53], therefore $\Sigma_{\boldsymbol{g}}$ always provides a tighter bound on EP than $\Sigma_{\boldsymbol{g}}^{\mathrm{KO}}$. The two bounds become equivalent in the limit $\langle e^{\boldsymbol{\theta}^{\top} \boldsymbol{g}} \rangle_{\tilde{p}} \to 1$, such as the short-time limit $T \to 0$ with antisymmetric observables when $\langle e^{\boldsymbol{\theta}^{\top} \boldsymbol{g}} \rangle_{\tilde{p}} = 1 + O(T)$ [43]. In general, however, the two bounds give different results, and in the End Matter, we provide a simple example where $\Sigma_{\boldsymbol{g}}$ is arbitrarily better than $\Sigma_{\boldsymbol{g}}^{\mathrm{KO}}$. We note that outside the short-time limit, $\Sigma_{\boldsymbol{g}}^{\mathrm{KO}}$ does not have a straightforward interpretation in terms of maximum likelihood inference or large-deviations statistics.

Lynn *et al.* [32,54,55] proposed a way to decompose and bound steady-state EP using a different information-theoretic optimization. Although originally focused on local EP in systems with multipartite dynamics, it can be generalized to arbitrary observables $\boldsymbol{g}$ and nonmultipartite systems as

$$\Sigma_{\boldsymbol{g}}^{\mathrm{L}} = \min_{q, \tilde{q}} D(q \| \tilde{q}) \text{ where } \langle \boldsymbol{g} \rangle_q = \langle \boldsymbol{g} \rangle_p, \tilde{q}(\boldsymbol{x}) = q(\tilde{\boldsymbol{x}}) \tag{15}$$

where the constraint $\tilde{q}(\boldsymbol{x}) = q(\tilde{\boldsymbol{x}})$ imposes that the trajectory distribution is stationary and without odd variables. Like our bound $\Sigma_{\boldsymbol{g}}$, $\Sigma_{\boldsymbol{g}}^{\mathrm{L}}$ can be used to generate hierarchical decompositions of EP.

Although Eq. (15) resembles Eq. (2), it differs in that it simultaneously optimizes both arguments of the KL divergence, q and $\tilde{q}$. As discussed in the End Matter, the optimal distribution in Eq. (15) is not in an exponential family, and in general it lacks full support. For this reason, Eq. (15) does not have a tractable dual expression analogous to our dual Eq. (3), and there is no straightforward way to scale $\Sigma_{\boldsymbol{g}}^{\mathrm{L}}$ to large systems.

Other work has explored objective (15) in combination with other types of constraints, e.g., on waiting times [14,19] and hidden Markovian structure [13,23]. However, the resulting optimization typically does not scale to high-dimensional or even moderately-sized systems, in part because the considered constraints lead to nonconvex optimization problems.

Our Pythagorean theorem (6) is related to previous decompositions of EP in interacting systems [56] and systems with nonconservative forces [35,57]. Moreover, the general form of Eq. (2) is a maximum entropy (MaxEnt) problem over trajectory distributions [58–61], sometimes called "maximum caliber" (MaxCal) [62]. Related techniques have been used to infer models from which the entropy flow can be estimated, though without ensuring a lower bound on EP [63].

Our approach differs in several other ways from earlier work on MaxEnt and MaxCal. First, we minimize the KL divergence relative to an unknown prior distribution $\tilde{p}$, from which we typically only have samples. This makes standard approximation methods for large-scale MaxEnt problems, such as mean-field and Bethe approximations [26,64,65], not directly applicable. Second, we care not only about the parameters $\boldsymbol{\theta}$, as is typical of MaxEnt "inverse problems," but also about bounding EP by the quantity $\Sigma_{\boldsymbol{g}}$. Both of these issues are resolved by the dual formulation (3).

Other notable approaches to EP inference have employed compression algorithms [66–68] and deep learning of probability flows [69].

Examples: Nonequilibrium spin model and Neuropixels dataset—We illustrate our method on two examples: a nonequilibrium kinetic Ising model [26,70,71] and *in vivo* spike data from the Neuropixels Visual Behavior repository [72]. For both examples, we consider a single discrete time step and binary variables, $\boldsymbol{x} = (\boldsymbol{x}_0, \boldsymbol{x}_1)$ where $x_{i,t} \in \{-1, +1\}$ for $i \in \{1, \ldots, N\}$. We emphasize that our approach also works for continuous systems. In SM [37], we illustrate it on a linear Langevin system.

To calculate $\Sigma_{\boldsymbol{g}}$, we sample from steady state using the Monte Carlo method. We then optimize Eq. (3) using gradient ascent with Barzilai-Borwein step sizes [73]. We perform early stopping using held-out validation data, which avoids overfitting even in the far-from-equilibrium regime where many reverse transitions are not sampled. Reported EP estimates are computed on held-out test data. Details of the data generation, optimization, and analysis can be found in SM [37] and our code repository [74].

The nonequilibrium spin model is specified by transition probabilities

$$T(\boldsymbol{x}_1 | \boldsymbol{x}_0) = \frac{1}{N} \sum_i \left[W_i(\boldsymbol{x}_0) \delta_{\boldsymbol{x}_0, \boldsymbol{x}_1^{[i]}} + (1 - W_i(\boldsymbol{x}_0)) \delta_{\boldsymbol{x}_0, \boldsymbol{x}_1} \right]$$

where we introduced the spin-flip operator: $(\boldsymbol{x}^{[i]})_i = -x_i$ and $(\boldsymbol{x}^{[i]})_j = x_j$ for $j \neq i$. The flip probability for spin i is

$$W_i(\boldsymbol{x}) = \frac{\exp(-\beta x_i \sum_{j: j \neq i} w_{ij} x_j)}{2 \cosh(\beta \sum_{j: j \neq i} w_{ij} x_j)}, \tag{16}$$

where β is an inverse temperature and w_{ij} are (typically asymmetric) coupling parameters. We consider the diluted version of the model [75–77] with k average neighbors.

Here $w_{ij} = c_{ij} z_{ij}/\sqrt{k}$, with binary connections $c_{ij} \sim$ Bernoulli$[k/(N-1)]$ and real-valued weights $z_{ij} \sim \mathcal{N}(0,1)$.

The steady-state EP in this model has a simple closed-form expression, $\Sigma = \beta \sum_{i \neq j} w_{ij} \langle (x_{i,1} - x_{i,0}) x_{j,0} \rangle_{\mathrm{st}}$, where $\langle \cdots \rangle_{\mathrm{st}}$ indicates expectations under the stationary process p. This provides ground truth to evaluate our estimators.

Our observables of interest are time-lagged correlations,

$$g_{ij}(\boldsymbol{x}) = (x_{i,1} - x_{i,0}) x_{j,0} \quad \text{for all } i, j, \tag{17}$$

with conjugate parameters θ_{ij}. Because the system has multipartite observables, Eq. (3) can be decomposed into N independent problems, each involving only those transitions where spin i changes, improving computational performance for large systems (see SM).

Figure 1(a) shows the actual and inferred EP at different β for $N = 1000$, $k = 6$. EP increases with β, and all estimators ($\Sigma_{\boldsymbol{g}}^{\mathrm{TUR}}$, $\hat{\Sigma}_{\boldsymbol{g}}$, $\Sigma_{\boldsymbol{g}}$) agree in the near-equilibrium regime of small β. Importantly, $\Sigma_{\boldsymbol{g}}$ provides a tight bound on EP even in the far-from-equilibrium regime of large β. The gap between $\Sigma_{\boldsymbol{g}}$ and $\hat{\Sigma}_{\boldsymbol{g}}$ indicates the onset of highly non-Gaussian statistics from $\beta \approx 2$.

As discussed above, $\Sigma_{\boldsymbol{g}} \approx \Sigma$ implies that the trajectory EP $\sigma(\boldsymbol{x})$ is closely approximated by $\sigma_{\boldsymbol{\theta}^*}$ from Eq. (8). As shown in SM, this allows us to use the optimal parameters $\boldsymbol{\theta}^*$ to infer the asymmetry of the coupling constants, $\theta^*_{ij} - \theta^*_{ji} \approx \beta(w_{ij} - w_{ji})$. This relation is verified in Fig. 1(b).

In SM, we also consider other coupling matrices for which the model exhibits a nonequilibrium phase transition [26]. We show that our method accurately infers EP in this regime, a difficult task for existing approximations [78].

Next, as an example application to biological data, we estimate EP in the Visual Behavior Neuropixels dataset [72], which contains spike-train recordings from 81 mice in 103 sessions. In this example, EP is understood as a statistical measure of temporal irreversibility, rather than energetic dissipation [33]. The dataset includes spiking activity from multiple brain regions, including visual cortical areas (VISp, VISl, VISal, VISrl, VISam, and VISpm) and subcortical structures.

We analyze data from visual areas during three conditions: active behavior (visual change detection task), passive replay (same stimuli but without task engagement), and Gabor (receptive field characterization with Gabor stimuli and full-field flashes). We discretize spike trains into temporal bins of length 10 ms and verify that most bins contain no more than one spike. In this dataset, neurons can update in a parallel (nonmultipartite) manner.

Our observables are defined as

$$g_{ij}(\boldsymbol{x}) = x_{i,1} x_{j,0} - x_{i,0} x_{j,1} \quad \text{for all } i < j, \tag{18}$$

which naturally represent time-lagged correlations in nonmultipartite dynamics. They capture the antisymmetric part of Eq. (17) and give more robust numerical results on this dataset.

We estimate EP across different conditions, recording sessions, and system sizes N. We randomly select ten subsets of neurons and estimate $\Sigma_{\boldsymbol{g}}$ for each subset. For improved comparison, we normalize EP in each condition by the expected number of spikes per bin, $R = \sum_i (1 + \langle x_{i,0} \rangle_{\mathrm{st}})/2$.

Figure 2(a) shows that EP grows superlinearly with size (note that $R \propto N$). It also shows that the active condition is associated with the largest normalized EP. Figure 2(b)

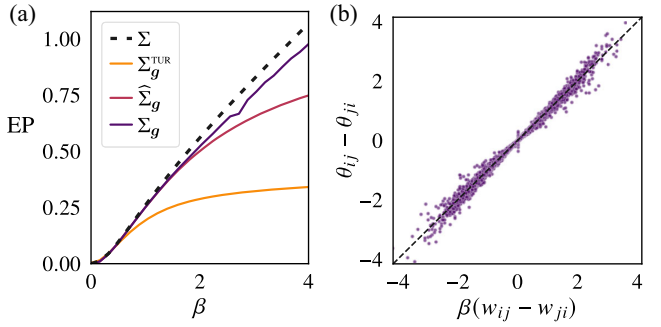

FIG. 1. Disordered nonequilibrium spin model with 1000 spins. (a) Steady-state EP estimates for different inverse temperatures β. (b) Asymmetry of inferred parameters, shown against the true coupling asymmetries in the model for $\beta = 2.5$ ($R^2 = 0.9831$). Estimates are based on 10^9 state transitions sampled by Monte Carlo.

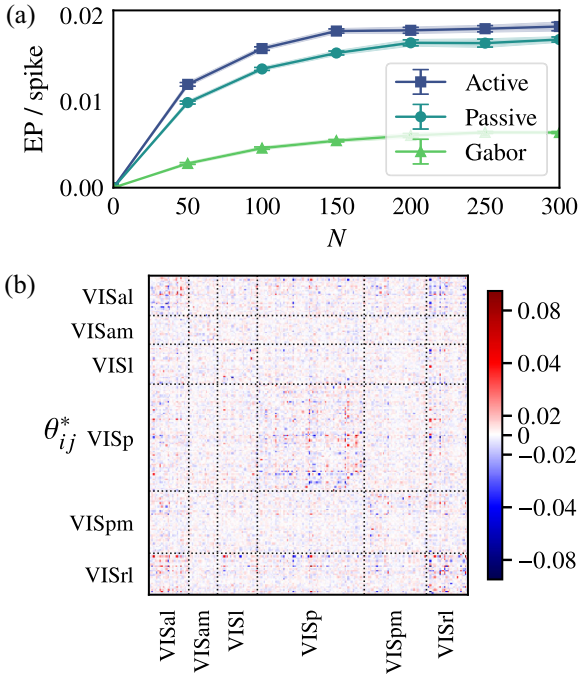

FIG. 2. (a) Estimated EP per expected number of spikes R, in the Neuropixels Visual Behavior dataset for three conditions. Error bars indicate the standard error of the mean. (b) A sample of inferred coupling coefficients θ^*_{ij} grouped by visual area. Here we select 200 neurons with the highest firing rate from an active trial. To improve visualization, the lower triangle shows $\theta^*_{ji} \equiv -\theta^*_{ij}$ for $i < j$.

illustrates the matrix of inferred parameters θ^*_{ij} for 200 neurons from an active trial, grouped by visual brain area. This matrix specifies a network of functional connectivity [79], showing how interactions between individual neurons contribute to temporal irreversibility of *in vivo* brain dynamics. This network reveals a clustered organization aligned with anatomical brain regions.

Acknowledgments—M. A. is partly supported by John Templeton Foundation (Grant No. 62828) and Grant No. PID2023-146869NA-I00 funded by MICIU/AEI/10.13039/501100011033 and cofunded by the European Union, and supported by the Basque Government through the No. BERC 2022-2025 program and by the Spanish State Research Agency through BCAM Severo Ochoa excellence accreditation No. CEX2021-01142-S funded by MICIU/AEI/10.13039/501100011033. S. I. is supported by JSPS KAKENHI Grants No. 22H01141, No. 23H00467, and No. 24H00834, JST ERATO Grant No. JPMJER2302, and UTEC-UTokyo FSI Research Grant Program. A. K. is partly supported by John Templeton Foundation (Grant No. 62828) and by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 101068029.

Data availability—The code reproducing the findings of this article is openly available at "Inferring Entropy Production in Many-Body Systems Using Nonequilibrium Maximum Entropy: Code repository (2026)" [74].

[1] U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines, Rep. Prog. Phys. **75**, 126001 (2012).
[2] G. Gallavotti and E. G. D. Cohen, Dynamical ensembles in nonequilibrium statistical mechanics, Phys. Rev. Lett. **74**, 2694 (1995).
[3] J. Kurchan, Fluctuation theorem for stochastic dynamics, J. Phys. A **31**, 3719 (1998).
[4] C. Maes, The fluctuation theorem as a Gibbs property, J. Stat. Phys. **95**, 367 (1999).
[5] C. Jarzynski, Hamiltonian derivation of a detailed fluctuation theorem, J. Stat. Phys. **98**, 77 (2000).
[6] U. Seifert, From stochastic thermodynamics to thermodynamic inference, Annu. Rev. Condens. Matter Phys. **10**, 171 (2019).
[7] S. Rahav and C. Jarzynski, Fluctuation relations and coarse-graining, J. Stat. Mech. (2007) P09012.
[8] M. Esposito, Stochastic thermodynamics under coarse graining, Phys. Rev. E **85**, 041125 (2012).
[9] S. Bo and A. Celani, Entropy production in stochastic systems with fast and slow time-scales, J. Stat. Phys. **154**, 1325 (2014).
[10] S.-W. Wang, K. Kawaguchi, S. I. Sasa, and L.-H. Tang, Entropy production of nanosystems with time scale separation, Phys. Rev. Lett. **117**, 070601 (2016).
[11] G. Bisker, M. Polettini, T. R. Gingrich, and J. M. Horowitz, Hierarchical bounds on entropy production inferred from partial information, J. Stat. Mech. (2017) 093210.
[12] M. Polettini and M. Esposito, Effective thermodynamics for a marginal observer, Phys. Rev. Lett. **119**, 240601 (2017).
[13] J. Ehrich, Tightest bound on hidden entropy production from partially observed dynamics, J. Stat. Mech. (2021) P083214.
[14] E. Nitzan, A. Ghosal, and G. Bisker, Universal bounds on entropy production inferred from observed statistics, Phys. Rev. Res. **5**, 043251 (2023).
[15] A. C. Barato and U. Seifert, Thermodynamic uncertainty relation for biomolecular processes, Phys. Rev. Lett. **114**, 158101 (2015).
[16] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, Dissipation bounds all steady-state current fluctuations, Phys. Rev. Lett. **116**, 120601 (2016).
[17] P. Pietzonka, F. Ritort, and U. Seifert, Finite-time generalization of the thermodynamic uncertainty relation, Phys. Rev. E **96**, 012101 (2017).
[18] S. K. Manikandan, D. Gupta, and S. Krishnamurthy, Inferring entropy production from short experiments, Phys. Rev. Lett. **124**, 120603 (2020).
[19] D. J. Skinner and J. Dunkel, Estimating entropy production from waiting time distributions, Phys. Rev. Lett. **127**, 198101 (2021).
[20] J. Van der Meer, B. Ertel, and U. Seifert, Thermodynamic inference in partially accessible Markov networks: A unifying perspective from transition-based waiting time distributions, Phys. Rev. X **12**, 031025 (2022).
[21] P. E. Harunari, A. Dutta, M. Polettini, and É. Roldán, What to learn from a few visible transitions' statistics?, Phys. Rev. X **12**, 041026 (2022).
[22] I. A. Martínez, G. Bisker, J. M. Horowitz, and J. M. Parrondo, Inferring broken detailed balance in the absence of observable currents, Nat. Commun. **10**, 3542 (2019).
[23] D. J. Skinner and J. Dunkel, Improved bounds on entropy production in living systems, Proc. Natl. Acad. Sci. U.S.A. **118**, e2024300118 (2021).
[24] J. van der Meer, J. Degünther, and U. Seifert, Time-resolved statistics of snippets as general framework for model-free entropy estimators, Phys. Rev. Lett. **130**, 257101 (2023).
[25] P. Pietzonka and F. Coghi, Thermodynamic cost for precision of general counting observables, Phys. Rev. E **109**, 064128 (2024).
[26] M. Aguilera, M. Igarashi, and H. Shimazaki, Nonequilibrium thermodynamics of the asymmetric Sherrington-Kirkpatrick model, Nat. Commun. **14**, 3685 (2023).
[27] C. Battle, C. P. Broedersz, N. Fakhri, V. F. Geyer, J. Howard, C. F. Schmidt, and F. C. MacKintosh, Broken detailed balance at mesoscopic scales in active biological systems, Science **352**, 604 (2016).
[28] T. H. Tan, A. Mietke, J. Li, Y. Chen, H. Higinbotham, P. J. Foster, S. Gokhale, J. Dunkel, and N. Fakhri, Odd dynamics of living chiral crystals, Nature (London) **607**, 287 (2022).

[29] C. W. Lynn, E. J. Cornblath, L. Papadopoulos, M. A. Bertolero, and D. S. Bassett, Broken detailed balance and entropy production in the human brain, Proc. Natl. Acad. Sci. U.S.A. **118**, e2109889118 (2021).

[30] L. A. de la Fuente, F. Zamberlan, H. Bocaccio, M. Kringelbach, G. Deco, Y. S. Perl, C. Pallavicini, and E. Tagliazucchi, Temporal irreversibility of neural dynamics as a signature of consciousness, Cereb. Cortex **33**, 1856 (2023).

[31] D. Sekizawa, S. Ito, and M. Oizumi, Decomposing thermodynamic dissipation of linear Langevin systems via oscillatory modes and its application to neural dynamics, Phys. Rev. X **14**, 041003 (2024).

[32] S. M. Geli, C. W. Lynn, M. L. Kringelbach, G. Deco, and Y. S. Perl, Non-equilibrium whole-brain dynamics arise from pairwise interactions, Cell Rep. Phys. Sci. **6**, 102464 (2025).

[33] R. Nartallo-Kaluarachchi, M. L. Kringelbach, G. Deco, R. Lambiotte, and A. Goriely, Nonequilibrium physics of brain dynamics, Phys. Rep. **1152**, 1 (2026).

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

[35] A. Kolchinsky, A. Dechant, K. Yoshimura, and S. Ito, Generalized free energy and excess entropy production for active systems, arXiv:2412.08432.

[36] T. M. Cover, *Elements of Information Theory* (John Wiley & Sons, New York, 1999).

[37] See Supplemental Material at http://link.aps.org/supplemental/10.1103/xgkj-dxzh for a detailed description of the Newton-Raphson method results and the nonequilibrium spin model and the Neuropixels spike train examples. Includes additional Refs. [38–40].

[38] S. Boyd, *Convex Optimization* (Cambridge University Press, Cambridge, England, 2004).

[39] J. Barzilai and J. M. Borwein, Two-point step size gradient methods, IMA J. Numer. Anal. **8**, 141 (1988).

[40] C.-J. Lin and J. J. Moré, Newton's method for large bound-constrained optimization problems, SIAM J. Optim. **9**, 1100 (1999).

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

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

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

[44] S.-I. Amari, Information geometry on hierarchy of probability distributions, IEEE Trans. Inf. Theory **47**, 1701 (2001).

[45] N. Ay, Information geometry on complexity and stochastic interaction, Entropy **17**, 2432 (2015).

[46] Equation (9) is derived by using $\delta\sigma_{\boldsymbol{\theta}}(\boldsymbol{x}) = \ln p(\boldsymbol{x})/q_{\boldsymbol{\theta}}(\boldsymbol{x})$ and $\langle e^{-\delta\sigma_{\boldsymbol{\theta}}}\rangle_p = 1$.

[47] J. M. Horowitz and T. R. Gingrich, Thermodynamic uncertainty relations constrain non-equilibrium fluctuations, Nat. Phys. **16**, 15 (2020).

[48] J. Li, J. M. Horowitz, T. R. Gingrich, and N. Fakhri, Quantifying dissipation using fluctuating currents, Nat. Commun. **10**, 1666 (2019).

[49] A. Dechant, Multidimensional thermodynamic uncertainty relations, J. Phys. A **52**, 035001 (2018).

[50] M. I. Belghazi, A. Baratin, S. Rajeswar, S. Ozair, Y. Bengio, A. Courville, and R. D. Hjelm, Mutual Information Neural Estimation, International Conference on Machine Learning (2018).

[51] M. D. Donsker and S. S. Varadhan, Asymptotic evaluation of certain Markov process expectations for large time, I, Commun. Pure Appl. Math. **28**, 1 (1975).

[52] 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).

[53] A. Ruderman, M. Reid, D. García-García, and J. Petterson, Tighter variational representations of f-divergences via restriction to probability measures, *Proceedings of the 29th International Conference on Machine Learning* (2012).

[54] C. W. Lynn, C. M. Holmes, W. Bialek, and D. J. Schwab, Decomposing the local arrow of time in interacting systems, Phys. Rev. Lett. **129**, 118101 (2022).

[55] C. W. Lynn, C. M. Holmes, W. Bialek, and D. J. Schwab, Emergence of local irreversibility in complex interacting systems, Phys. Rev. E **106**, 034102 (2022).

[56] 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).

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

[58] A. Tang, D. Jackson, J. Hobbs, W. Chen, J. L. Smith, H. Patel, A. Prieto, D. Petrusca, M. I. Grivich, A. Sher *et al.*, A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro, J. Neurosci. **28**, 505 (2008).

[59] O. Marre, S. El Boustani, Y. Frégnac, and A. Destexhe, Prediction of spatiotemporal patterns of neural activity from pairwise correlations, Phys. Rev. Lett. **102**, 138101 (2009).

[60] H. C. Nguyen, R. Zecchina, and J. Berg, Inverse statistical problems: From the inverse Ising problem to data science, Adv. Phys. **66**, 197 (2017).

[61] R. Cofré and C. Maldonado, Information entropy production of maximum entropy Markov chains from spike trains, Entropy **20**, 34 (2018).

[62] K. Ghosh, P. D. Dixit, L. Agozzino, and K. A. Dill, The maximum caliber variational principle for nonequilibria, Annu. Rev. Phys. Chem. **71**, 213 (2020).

[63] K. Ishihara and H. Shimazaki, State-space kinetic ising model reveals task-dependent entropy flow in sparsely active nonequilibrium neuronal dynamics, Nat. Commun. **16**, 10852 (2025).

[64] M. Welling and Y. W. Teh, Approximate inference in Boltzmann machines, Artif. Intell. **143**, 19 (2003).

[65] F. Ricci-Tersenghi, The Bethe approximation for solving the inverse Ising problem: A comparison with other inference methods, J. Stat. Mech. (2012) P08015.

[66] É. Roldán and Juan M. R. Parrondo, Estimating dissipation from single stationary trajectories, Phys. Rev. Lett. **105**, 150607 (2010).

[67] É. Roldán and Juan M. R. Parrondo, Entropy production and Kullback-Leibler divergence between stationary trajectories of discrete systems, Phys. Rev. E **85**, 031129 (2012).
[68] S. Ro, B. Guo, A. Shih, T. V. Phan, R. H. Austin, D. Levine, P. M. Chaikin, and S. Martiniani, Model-free measurement of local entropy production and extractable work in active matter, Phys. Rev. Lett. **129**, 220601 (2022).
[69] N. M. Boffi and E. Vanden-Eijnden, Deep learning probability flows and entropy production rates in active matter, Proc. Natl. Acad. Sci. U.S.A. **121**, e2318106121 (2024).
[70] A. Crisanti and H. Sompolinsky, Dynamics of spin systems with randomly asymmetric bonds: Ising spins and Glauber dynamics, Phys. Rev. A **37**, 4865 (1988).
[71] H. Eissfeller and M. Opper, Mean-field Monte Carlo approach to the Sherrington-Kirkpatrick model with asymmetric couplings, Phys. Rev. E **50**, 709 (1994).
[72] Allen Institute for Brain Science, *Allen Brain Observatory: Visual Behavior Neuropixels Technical Whitepaper v1.0–8.16.22*, Whitepaper (Allen Institute for Brain Science, 2022).
[73] R. Fletcher, On the Barzilai-Borwein method, in *Optimization and Control with Applications* (Springer, New York, 2005), pp. 235–256.
[74] M. Aguilera and A. Kolchinsky, Inferring entropy production in many-body systems using nonequilibrium maximum entropy: Code repository (2026), https://github.com/MiguelAguilera/Entropy-Production-MaxEnt-Inference.
[75] A. Coolen, Statistical mechanics of recurrent neural networks II—Dynamics, in *Handbook of Biological Physics*, Vol. 4 (Elsevier, New York, 2001), pp. 619–684.
[76] E. Aurell and H. Mahmoudi, A message-passing scheme for non-equilibrium stationary states, J. Stat. Mech. (2011) P04014.
[77] P. Zhang, Inference of kinetic Ising model on sparse graphs, J. Stat. Phys. **148**, 502 (2012).
[78] M. Aguilera, S. A. Moosavi, and H. Shimazaki, A unifying framework for mean-field theories of asymmetric kinetic Ising systems, Nat. Commun. **12**, 1197 (2021).
[79] O. Sporns, Structure and function of complex brain networks, Dialogues Clin. Neurosci. **15**, 247 (2013).
[80] Y. Hasegawa and T. Van Vu, Fluctuation theorem uncertainty relation, Phys. Rev. Lett. **123**, 110602 (2019).
[81] L. Peliti and S. Pigolotti, *Stochastic Thermodynamics: An Introduction* (Princeton University Press, Princeton, NJ, 2021).

End Matter

Multipartite observables—We show that for multipartite observables, our optimization problem can be split into a set of simpler subproblems.

We say that observables $\boldsymbol{g}$ are multipartite if they can be decomposed as $\boldsymbol{g} = (\boldsymbol{g}_1, a_1, \ldots, \boldsymbol{g}_k, a_k)$ such that: (1) only a single "block" of observables $\boldsymbol{g}_i$ may be active (nonzero) under any forward or backward trajectory: $p(\boldsymbol{x}) = \tilde{p}(\boldsymbol{x}) = 0$ whenever $\boldsymbol{g}_i(\boldsymbol{x}) \neq \boldsymbol{0} \wedge \boldsymbol{g}_j(\boldsymbol{x}) \neq \boldsymbol{0}$ for some $i \neq j$, and (2) each $a_i(\boldsymbol{x}) \coloneqq 1 - \delta_{\boldsymbol{g}_i(\boldsymbol{x}),\boldsymbol{0}}$ is an indicator variable for activity of block i. For convenience, we use $a_0(\boldsymbol{x}) \coloneqq 1 - \sum_{i=1} a_i(\boldsymbol{x})$ to indicate that no block is active. [With minor changes, the derivations below generalize to the case where $a_i(\boldsymbol{x}) = \gamma_i(1 - \delta_{\boldsymbol{g}_i(\boldsymbol{x}),\boldsymbol{0}})$ for some constants $\gamma_i \neq 0$.]

Let $P_i \coloneqq \langle a_i \rangle_p$ be the forward probability that block i is active, and $p_i(\boldsymbol{x}) \coloneqq p(\boldsymbol{x}|\boldsymbol{g}_i \neq \boldsymbol{0}) \coloneqq a_i(\boldsymbol{x})p(\boldsymbol{x})/P_i$ the forward trajectory distribution conditioned on block i being active. $P_0 = 1 - \sum_i P_i$ is the probability that no block is active, and p_0 is the trajectory distribution conditioned on no block being active. $\tilde{P}_i$, and $\tilde{p}_i$ indicate the same quantities under the reverse process $\tilde{p}$. In general, P_i and $\tilde{P}_i$ can be estimated from empirical frequencies, as long as the number of blocks is not very large.

We now consider our variational expression (3) for $\Sigma_{\boldsymbol{g}}$, and we write the optimization variables as $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_k, \boldsymbol{\lambda})$, with $\boldsymbol{\theta}_i$ conjugate to $\boldsymbol{g}_i$ and λ_i conjugate to a_i for $i = 1 \ldots k$. For multipartite observables, Eq. (3) can be written as

$$\begin{aligned}\Sigma_{\boldsymbol{g}} &= \max_{\boldsymbol{\theta}} \sum_{i=1} (\boldsymbol{\theta}_i^\top \langle \boldsymbol{g}_i \rangle_p + \lambda_i \langle a_i \rangle_p) - \ln Z_{\boldsymbol{\theta},\boldsymbol{\lambda}} \\ &= \max_{\boldsymbol{\theta}} \sum_{i=1} P_i (\boldsymbol{\theta}_i^\top \langle \boldsymbol{g}_i \rangle_{p_i} + \lambda_i) - \ln Z_{\boldsymbol{\theta},\boldsymbol{\lambda}}, \end{aligned} \tag{A1}$$

where we introduced $Z_{\boldsymbol{\theta},\boldsymbol{\lambda}} \coloneqq \langle e^{\boldsymbol{\theta}^\top \boldsymbol{g}} \rangle_{\tilde{p}}$ for convenience. Defining $Z^i_{\boldsymbol{\theta}_i} \coloneqq \langle e^{\boldsymbol{\theta}_i^\top \boldsymbol{g}_i} \rangle_{\tilde{p}_i}$, we may write this term as

$$Z_{\boldsymbol{\theta},\boldsymbol{\lambda}} \coloneqq \langle e^{\sum_{i=1} \lambda_i a_i + \boldsymbol{\theta}_i^\top \boldsymbol{g}_i} \rangle_{\tilde{p}} = \tilde{P}_0 + \sum_{i=1} \tilde{P}_i e^{\lambda_i^*} Z^i_{\boldsymbol{\theta}_i}. \tag{A2}$$

We find the optimal $\boldsymbol{\lambda}^*$ by taking derivatives of the objective (A1) with respect to each λ_i:

$$0 = P_i - \partial_{\lambda_i} \ln Z_{\boldsymbol{\theta},\boldsymbol{\lambda}^*} \Rightarrow P_i = \tilde{P}_i e^{\lambda_i^*} Z^i_{\boldsymbol{\theta}_i} / Z_{\boldsymbol{\theta},\boldsymbol{\lambda}^*}, \tag{A3}$$

which gives $\lambda_i^* = \ln(P_i/\tilde{P}_i) - \ln Z^i_{\boldsymbol{\theta}_i} + \ln Z_{\boldsymbol{\theta},\boldsymbol{\lambda}^*}$. Plugging $\tilde{P}_i e^{\lambda_i^*} Z^i_{\boldsymbol{\theta}_i} = Z_{\boldsymbol{\theta},\boldsymbol{\lambda}^*} P_i$ into Eq. (A2) and rearranging gives $Z_{\boldsymbol{\theta},\boldsymbol{\lambda}^*} = \tilde{P}_0/P_0$. Combining with Eq. (A1) and simplifying shows that, at the optimal $\boldsymbol{\lambda}^*$, the objective can be written as

$$\Sigma_{\boldsymbol{g}} = D(P\|\tilde{P}) + \max_{\boldsymbol{\theta}_1,\ldots,\boldsymbol{\theta}_k} \left[\sum_{i=1} P_i (\boldsymbol{\theta}_i^\top \langle \boldsymbol{g}_i \rangle_{p_i} - \ln Z^i_{\boldsymbol{\theta}_i}) \right].$$

The sum now involves nonoverlapping parameter blocks that can be optimized independently.

We arrive at our main result for multipartite observables:

$$\Sigma_{\boldsymbol{g}} = D(P\|\tilde{P}) + \sum_{i=1} P_i \Sigma^{(i)}_{\boldsymbol{g}}, \tag{A4}$$

where $\Sigma^{(i)}_{\boldsymbol{g}}$ is the contribution from block $i = 1 \ldots k$,

$$\Sigma^{(i)}_{\boldsymbol{g}} \coloneqq \max_{\boldsymbol{\theta}_i} (\boldsymbol{\theta}_i^\top \langle \boldsymbol{g}_i \rangle_{p_i} - \ln \langle e^{\boldsymbol{\theta}_i^\top \boldsymbol{g}_i} \rangle_{\tilde{p}_i}). \tag{A5}$$

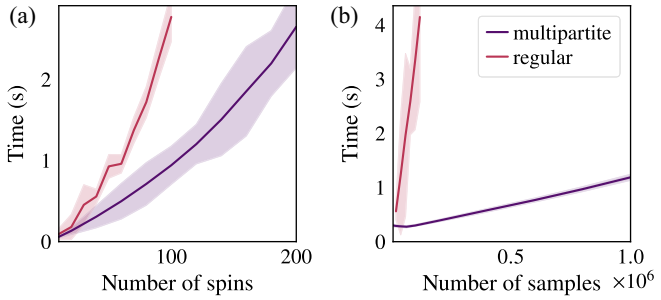

FIG. 3. Computation time for Σ_g using the regular (nonmultipartite) optimization (3) versus the multipartite decomposition (A4). (a) Runtime versus number of spins at fixed number of samples 2×10^4. (b) Runtime versus number of samples per spin (system with 40 spins). Shaded bands: standard deviations over 1000 trials. The end of the "regular" curve indicates the point where the GPU runs out of memory. Hardware: Intel Core i9-12900KF CPU, NVIDIA GeForce RTX 3050 GPU with 8 GB VRAM.

Although the decomposition (A4) increases the number of optimization problems, each optimization problem is (typically $\sim k$ times) smaller than Eq. (3), both in terms of the number of optimization variables and data points needed to estimate expectations under p_i and $\tilde{p}_i$. In many cases, we may also bound each $\Sigma_g^{(i)}$ using the optimization-free bound $\hat{\Sigma}_g^{(i)}$ using a smaller covariance matrix than in Eq. (11).

Figure 3 demonstrates the improved performance scaling of the multipartite optimization. Here we consider the nonequilibrium spin model across different system sizes (number of spins) and dataset sizes (number of samples). The multipartite method exhibits slower growth in computation time, particularly as the number of samples increases. In contrast, the nonmultipartite optimization rapidly exhausts the 8 GB of GPU memory available in our experiments.

We finish by considering the approximation of trajectory EP. For multipartite observables, we can decompose Eq. (8) as

$$\sigma_{\boldsymbol{\theta}^*}(\boldsymbol{x}) = \sum_i [\lambda_i^* a_i(\boldsymbol{x}) + \boldsymbol{\theta}_i^{*\top} \boldsymbol{g}_i(\boldsymbol{x})] - \ln Z_{\boldsymbol{\theta},\boldsymbol{\lambda}^*}$$
$$= a_0(\boldsymbol{x}) \ln \frac{P_0}{\tilde{P}_0} + \sum_i a_i(x) \left[\ln \frac{P_i}{\tilde{P}_i} + \boldsymbol{\theta}_i^{*\top} \boldsymbol{g}_i(\boldsymbol{x}) - \ln Z_{\boldsymbol{\theta}_i^*}^i \right]$$

where we used expressions of λ_i^* and $Z_{\boldsymbol{\theta}^*,\boldsymbol{\lambda}^*}$ derived above.

When each observable block $\boldsymbol{g}_i$ is antisymmetric and the system is in steady state and without odd variables, it can be shown that $P_i = \tilde{P}_i$. This leads to the simplified expressions for estimators of average and trajectory EP:

$$\Sigma_g = \sum_{i=1} P_i \Sigma_g^{(i)}$$
$$\sigma_{\boldsymbol{\theta}^*}(\boldsymbol{x}) = \sum_i a_i(x) \left[\boldsymbol{\theta}_i^{*\top} \boldsymbol{g}_i(\boldsymbol{x}) - \ln Z_{\boldsymbol{\theta}_i^*}^i \right]. \tag{A6}$$

Multidimensional TUR Eq. (12)—For a stationary system without odd variables, the "fluctuation theorem uncertainty relation" [80] states that $e^{\Sigma} - 1 \geq 2\langle \phi \rangle_p^2 / \mathrm{Var}_p(\phi)$ for any antisymmetric scalar observable $\phi(\boldsymbol{x})$. Then, for any antisymmetric $\boldsymbol{g}(\boldsymbol{x})$ and any $\boldsymbol{\alpha} \in \mathbb{R}^d$, we may define $\phi(\boldsymbol{x}) \coloneqq \boldsymbol{\alpha}^\top \boldsymbol{g}(\boldsymbol{x})$. After some rearranging, this gives

$$\Sigma \geq \ln[1 + 2(\boldsymbol{\alpha}^\top \langle \boldsymbol{g} \rangle_p)^2 / (\boldsymbol{\alpha}^\top \mathsf{K}_p \boldsymbol{\alpha})]. \tag{B1}$$

Setting $\boldsymbol{\alpha} = \mathsf{K}_p^{-1} \langle \boldsymbol{g} \rangle_p$ and simplifying gives Eq. (12).

A similar derivation may be found in Ref. [81] [Eq. (8.31)]. A continuous-time version of this bound appeared in Ref. [49].

Comparison with Σ_g^{KO} from Refs. [42,43]—Here we compare Σ_g (3) with the variational bound Σ_g^{KO} (14) from Refs. [42,43]. Using simple examples, we show that Σ_g sometimes gives an arbitrarily better bound on EP than Σ_g^{KO}.

We consider a three-state Markov chain measured at two time points $t \in \{0, 1\}$, with corresponding states $x_t \in \{0, 1, 2\}$. The system has uniform unicyclic transition probabilities. We parameterize the probability of moving up ($x_1 = x_0 + 1 \mod 3$), down ($x_1 = x_0 - 1 \mod 3$), and staying ($x_1 = x_0$) as

$$P_+ = \frac{\kappa}{e^{-\lambda} + 1}, \quad P_- = \frac{\kappa e^{-\lambda}}{e^{-\lambda} + 1}, \quad P_{\mathrm{stay}} = 1 - \kappa. \tag{C1}$$

The parameter κ controls dynamical activity while λ controls driving strength. The steady state is nonequilibrium if $\lambda \neq 0$.

The system has a uniform steady state with EP $\Sigma = (P_+ - P_-)\lambda$. For the estimators, we first consider a single observable $g(\boldsymbol{x}) = 1 + \delta_{x_0+1,x_1} - \delta_{x_0-1,x_1}$, whose expectation is $\langle g \rangle_p = 1 + (P_+ - P_-)$. The values of the two estimators are

$$\Sigma_g = \max_{\theta \in \mathbb{R}} [\theta \langle g \rangle_p - \ln(P_{\mathrm{stay}} e^{\theta} + P_- e^{2\theta} + P_+)] \tag{C2}$$

$$\Sigma_g^{\mathrm{KO}} = \max_{\theta \in \mathbb{R}} [\theta \langle g \rangle_p - P_{\mathrm{stay}} e^{\theta} - P_- e^{2\theta} - P_+ + 1]. \tag{C3}$$

Σ, Σ_g, and Σ_g^{KO} all vanish at $\lambda = 0$ and increase monotonically in $\lambda > 0$. The optimization problem (C2) can be solved in closed form to find $\Sigma_g = \Sigma$ with optimal parameter $\theta^* = \ln P_+ / P_- = \lambda$, thus $\Sigma = \Sigma_g \to \infty$ in the irreversible limit $\lambda \to \infty$. The expression for Σ_g^{KO} is more complicated, but it can be shown that it saturates at the finite value $\Sigma_g^{\mathrm{KO}} \to (1 + \kappa) 2 \tanh^{-1} \kappa - 2\kappa$ in the irreversible limit. Numerical results are shown in Fig. 4(a) for $\kappa = 9/10$ and a range of driving strengths. It is seen that Σ_g^{KO} remains finite while Σ_g diverges for large λ.

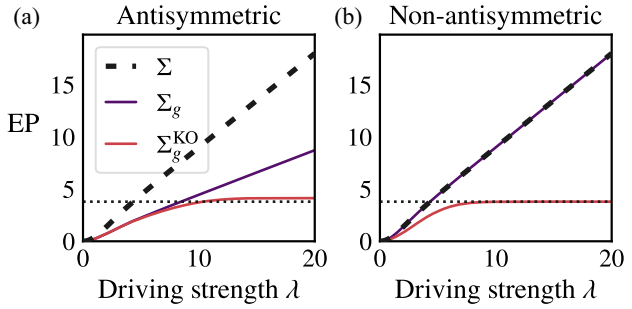

FIG. 4. Comparison between our EP bound $\Sigma_{\boldsymbol{g}}$ (3) and the variational bound $\Sigma_{\boldsymbol{g}}^{\mathrm{KO}}$ (14) on a three-state unicyclic system. (a) For the nonantisymmetric observable $g(\boldsymbol{x}) = 1 + \delta_{x_0+1,x_1} - \delta_{x_0-1,x_1}$. (b) For the antisymmetric observable $g(\boldsymbol{x}) = \delta_{x_0+1,x_1} - \delta_{x_0-1,x_1} - 1.2(\delta_{x_1,1}\delta_{x_0,0} - \delta_{x_0,1}\delta_{x_1,0})$. In both cases, $\Sigma_{\boldsymbol{g}}$ diverges in the irreversible limit $\lambda \to \infty$, while $\Sigma_{\boldsymbol{g}}^{\mathrm{KO}}$ saturates at a finite value.

The above example features a nonantisymmetric observable, but the effect also holds for some antisymmetric observables. As an example, we may consider the observable $g(\boldsymbol{x}) = \delta_{x_0+1,x_1} - \delta_{x_0-1,x_1} - 1.2(\delta_{x_1,1}\delta_{x_0,0} - \delta_{x_0,1}\delta_{x_1,0})$. In this case, $\Sigma_{\boldsymbol{g}} < \Sigma$, but our estimator still captures a diverging amount of EP in the irreversible limit, $\Sigma_{\boldsymbol{g}} \to \infty$ as $\lambda \to \infty$. On the other hand, $\Sigma_{\boldsymbol{g}}^{\mathrm{KO}}$ saturates at the finite value ≈ 4.13 as $\lambda \to \infty$ ($\kappa = 9/10$). Figure 4(b) displays the numerical values of the two estimators for this antisymmetric observable.

Comparison with Lynn et al. [55]—We consider the optimization problem that defines $\Sigma_{\boldsymbol{g}}^{\mathrm{L}}$ (15), our generalization of the estimator from Ref. [55]. The partial derivative of the objective with respect to $q(\boldsymbol{x})$ is

$$\partial_{q(\boldsymbol{x})} D(q\|\tilde{q}) = 1 + \ln\frac{q(\boldsymbol{x})}{\tilde{q}(\boldsymbol{x})} - \frac{\tilde{q}(\boldsymbol{x})}{q(\boldsymbol{x})},$$

where we used the condition $\tilde{q}(\boldsymbol{x}) = q(\tilde{\boldsymbol{x}})$. This expression does not diverge even when a pair of probabilities $q(\boldsymbol{x})$ and $q(\tilde{\boldsymbol{x}})$ approach zero. Therefore, in general, the optimal distribution in Eq. (15) may not have full support, instead lying on the boundary of the set of probability distributions. Using some numerical examples, we have verified that the optimization (15) often returns solutions without full support.

For comparison, the partial derivative of our MaxEnt objective (2), $\partial_{q(\boldsymbol{x})} D(q\|\tilde{p}) = 1 + \ln[q(\boldsymbol{x})/\tilde{p}(\boldsymbol{x})]$ diverges to $-\infty$ as $q(\boldsymbol{x}) \to 0$ for any $\boldsymbol{x}$ where $\tilde{p}(\boldsymbol{x}) > 0$. This implies that a point on the boundary cannot be the minimizer, thus the optimal solution will lie in the relative interior of the feasible set.

This difference has significant consequences for the tractability of the two optimization problems. The fact that strict positivity is enforced by our objective (2) allows us to restate the optimization in terms of a tractable dual problem (3). On the other hand, the dual formulation of Eq. (15) requires an exponential number of nonnegative Lagrange multipliers, one for each pair of trajectories $\boldsymbol{x}$ and $\tilde{\boldsymbol{x}}$ to guarantee non-negativity of $q(\boldsymbol{x})$ and $q(\tilde{\boldsymbol{x}})$. Therefore, the dual formulation of Eq. (15) involves an intractable constrained optimization problem over an exponential number of parameters.