


Utilizing time-series measurements for entropy-production estimation in partially observed systems

Uri Kapustin

School of Electrical Engineering, Faculty of Engineering, Tel Aviv University, Tel Aviv 6997801, Israel

Aishani Ghosal

*Department of Biomedical Engineering, Faculty of Engineering, Tel Aviv University, Tel Aviv 6997801, Israel*Gili Bisker^{✉*}*Department of Biomedical Engineering, Faculty of Engineering, Tel Aviv University, Tel Aviv 6997801, Israel;**Center for Physics and Chemistry of Living Systems, Tel Aviv University, Tel Aviv 6997801, Israel;**Center for Nanoscience and Nanotechnology, Tel Aviv University, Tel Aviv 6997801, Israel;**and Center for Light-Matter Interaction, Tel Aviv University, Tel Aviv 6997801, Israel* (Received 27 December 2022; revised 2 March 2024; accepted 26 March 2024; published 10 April 2024)

Estimating the dissipation, or the entropy-production rate (EPR), can provide insights into the underlying mechanisms of nonequilibrium-driven processes. However, in practical experimental settings, precise quantification of the EPR can be challenging, as only partial information is typically accessible. Here, we explore the relationship between the observed information and the accuracy of EPR estimation. We employ a range of coarse-grained time-series trajectory data, simulating scenarios where varying degrees of information are available. We discover a hierarchy of lower bounds on the total EPR, demonstrating that an increasing amount of information can be leveraged for obtaining tighter EPR estimation, underscoring the critical role of exploiting the available data. Moreover, we introduce a technique for utilizing waiting times within hidden states and tightening the lower bound on the total EPR for some cases. This approach highlights the potential of hidden features within the data to provide valuable insights into the dissipative dynamics of complex systems.

DOI: [10.1103/PhysRevResearch.6.023039](https://doi.org/10.1103/PhysRevResearch.6.023039)

I. INTRODUCTION

Entropy production, often referred to as energy dissipation, occupies a central role in the realm of nonequilibrium thermodynamics [1–3]. It serves as a fundamental quantity, providing crucial insights into the thermodynamic behavior of systems that operate far-from-equilibrium. In living systems, for example, the rate of dissipation is closely related to the consumption rate of chemical fuel molecules, such as adenosine triphosphate (ATP) by molecular motors [4] that drive numerous cellular processes, including cell division, locomotion, and growth [5].

The entropy production calculated along a single trajectory is a stochastic quantity, which has gathered significant experimental interest [6–16]. Estimating the total entropy-production rate (EPR), i.e., the average rate of the production of entropy, is straightforward, given complete information about the dissipative degrees of freedom [17]. However, practical scenarios pose a challenge on calculating the EPR since experimental data always have a finite resolution, and only a subset of the degrees of freedom are accessible [2]. Such

coarse-grained observation, where only some of the dissipative degrees of freedom are monitored or resolved, can only provide a lower bound on the total dissipation rate [18–25]. In some cases, only the transitions between states are observable, omitting information about the states themselves [26–31]. Nevertheless, exploiting the available data optimally, to gain the tightest possible lower bound on the EPR, is a subject of active research.

There are several estimators for partial EPR that do not require any prior information about the system, such as the number of states or the underlying topology, based on the deep connection between the dissipation and the breaking of time-reversal symmetry [32]. These estimators rely on the direct link between EPR and the difficulty of distinguishing forward and reverse processes, quantified by the relative entropy or Kullback–Leibler Divergence (KLD) [18,33–39]. Calculating the KLD between probability distributions of forward and reverse trajectories can be done using various techniques [33–36,40]. For instance, the plug-in method involves estimating the probabilities of data sequences while discarding information about transition times [35]. In contrast, the KLD can be calculated analytically when applied to semi-Markov processes, where it was shown to be the sum of two contributions, one of which captures irreversibility in the sequence of states, and the other quantifying irreversibility in waiting-time distributions (WTD) [33,41]. For second-order semi-Markov processes, this KLD estimator can detect and quantify entropy production even in the absence of observable currents [33,42,43].

*bisker@tauex.tau.ac.il

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

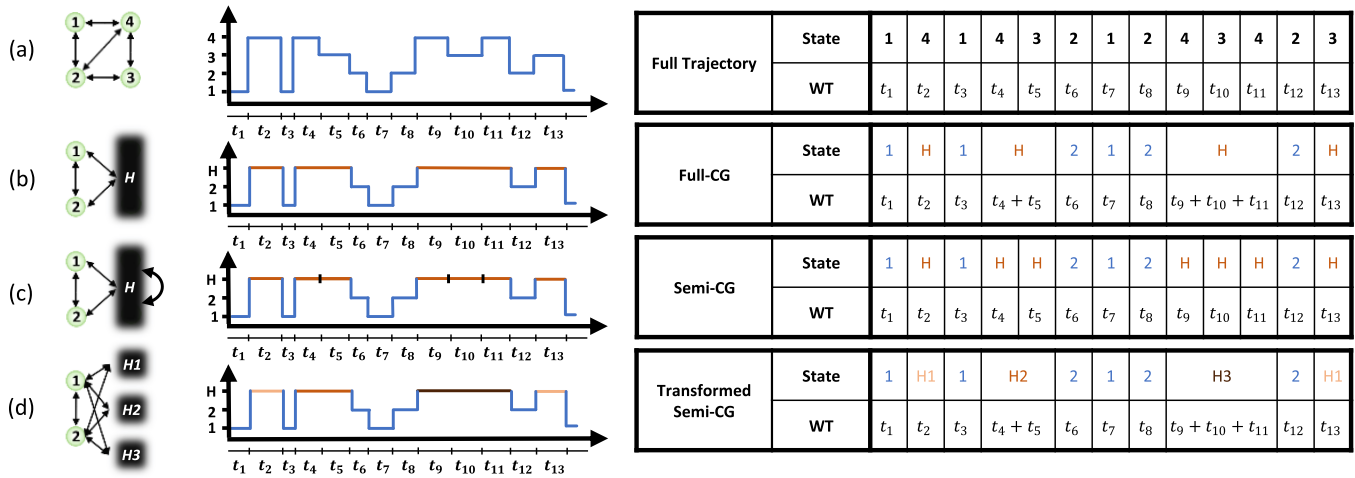


FIG. 1. Illustration of the partial information frameworks for arbitrary four-state system. (a) A fully observed four-states system. The trajectory (blue line) is described by the sequence of microstates and the corresponding waiting times (WT). (b) Full coarse-graining (*full-CG*). States 3 and 4 cannot be resolved, and are lumped together to a single macrostate H (orange line). (c) Semi-coarse-graining (*semi-CG*). States 3 and 4 cannot be resolved, but intratransitions between the hidden microstates can be recorded, where consecutive visits in the hidden microstates (orange line with black markers) are recorded as a sequence of H , with the corresponding WT of each hidden microstate between intratransition events. (d) Transformed semi-coarse-graining (*transformed semi-CG*). Each n consecutive visits to the hidden microstates in H are recorded as H_n (light orange, orange, and brown represent different sequence lengths) and the WT in H_n are the sum of the WT in consecutive visits in the hidden microstates 3 and 4.

Machine learning (ML) has also emerged as a valuable tool for entropy-production estimation by exploiting the irreversibility of data series [44–46]. The core idea is to optimize an objective function whose extremum is the KLD between the forward and reverse trajectories of sequences of states. For example, the recurrent neural network estimator for entropy-production (RNEEP) estimates the EPR from coarse-grained data of partially observed systems, using a recurrent neural network to solve the optimization problem [44]. Other estimators search over systems with the same statistics as the system at hand using different sets of observables [47–50], or provide a lower bound on the entropy production from the fluctuations of the transition fluxes or first passage times using the thermodynamic uncertainty relations (TUR) [51–60].

In this work, we focus on a continuous-time Markov chain (CTMC) model over a discrete set of states, in which a subset of the microstates is coarse-grained, or “lumped,” into a single macrostate. We consider different levels of observed statistics from different coarse-graining (CG) approaches and infer the EPR from the observed data using the KLD estimator [33], the plug-in estimator [35,36], and the RNEEP estimator [44], when applicable. These estimators do not require prior knowledge of the systems and only use the observed statistics to infer and quantify time-irreversibility. First, we use the sequence of observed microstates and coarse-grained macrostates, and the transitions between them. Then, we include information about transitions between the hidden microstates within the coarse-grained macrostates (intratransitions). In the latter case, the data do not follow semi-Markov statistics, so the KLD estimator cannot be directly applied. To overcome this challenge, we introduce a unique reformulation of the trajectory data of observed states and transitions, and intratransitions within macrostates, that enables us to exploit waiting-time information for a tighter EPR bound using

the KLD estimator. Specifically, we label the coarse-grained macrostates according to the number of times they are visited before jumping into an observed state. We apply the CG approaches to three model systems, namely, a four-state system in which two of the states are coarse-grained into a single hidden state (Fig. 1), a molecular motor model with two internal states that cannot be resolved by an external observer (Fig. 2), and the discrete Flashing Ratchet model with time-varying potential, whose values cannot be observed (Fig. 3). We provide a unifying comparison between the plug-in, RNEEP, and KLD estimators across various CG schemes, and emphasize how additional information can be exploited to infer tighter lower bounds on the total EPR.

II. THEORETICAL FRAMEWORK

A. Model systems

We explore three generic CTMC models as case studies, namely, a four-state system, a molecular motor, and a flashing




FIG. 2. A model of a molecular motor that can physically jump between spatial positions $i \leftrightarrow i + 1$, and switch between internal conformation $i \leftrightarrow i'$ corresponding to active (i') or passive (i) states. An external observer can only record the position of the motor, such that states of the same position, i and i' , regardless of the internal conformation, are coarse-grained into a single macrostate H_i . In the *semi-CG* framework, intratransitions $i \leftrightarrow i'$ are recorded.




FIG. 3. Discrete flashing ratchet of three states with periodic boundaries, and a potential that can be switched on (i) and off (i') but is not accessible to the observer. States of the same position i , regardless of the potential, are coarse-grained into a single macrostate H_a , H_b , and H_c . In the *semi-CG* framework, intratransitions $i \leftrightarrow i'$ are recorded.

ratchet. In each of these systems, some of the states cannot be resolved, giving rise to coarse-grained data of the system trajectories, rendering the observed statistics a second-order semi-Markov process [33].

In the four-state system, two of the states cannot be resolved and are effectively coarse-grained into a single hidden state (Fig. 1). Specifically, two states are observed (Markovian) states, while the other two states are lumped into a single macrostate. We allow an external control parameter to tune the transition rates over the observed link, thereby affecting the observed dynamics.

In the case of the molecular motor (Fig. 2), the motor is capable of both physical movement in space, transitioning between positions labeled as $i \leftrightarrow i + 1$ (upward or downward), and altering between two internal states, $i \leftrightarrow i'$ (passive or active). The motor movement is driven by an external source of chemical energy, $\Delta\mu$, facilitating upward spatial transitions from the active state, while a mechanical force F opposes this motion. The internal states can be thought of as two different conformations of the motor, and, therefore, cannot be distinguished by an external observer, who can only access the physical position of the motor.

The discrete flashing ratchet is commonly used to study the behavior of particles in the presence of an external potential, in which a series of barriers and wells are alternately flashed on and off in a periodic manner [36,61] (Fig. 3). This creates a series of potential energy barriers that particles must overcome to move in a preferred direction. Similar to the molecular motor, the system can jump between states $i \leftrightarrow i + 1$ in which the potential is on, and $i' \leftrightarrow (i + 1)'$ in which the potential is off, and can also switch between the on and off states, $i \leftrightarrow i'$. The system is inherently driven out of equilibrium by the external force that flashes the potential energy barriers.

B. Coarse-graining (CG)

Coarse-graining a system implies grouping together multiple states or variables into a smaller set of macroscopic variables. The coarse-grained trajectories of these systems represent the time-evolution of the system as it transitions between the coarse-grained states. Generally, the coarse-grained data cannot be treated with the Markovian approximation [62–65], and different levels of coarse-graining can be considered based on the partial information available about the system. In this section, we explain the different

coarse-graining approaches, taking the four-state model system as an example (Fig. 1).

1. Full-coarse-grain (full-CG)

In this CG approach, termed *full-CG*, we lump together a subset of the microstates into a single observed state, giving rise to a second-order semi-Markov process [Figs. 1(a) and 1(b)], since the waiting time in the hidden state depends on the state visited before [33]. In this example, states 1 and 2 are observed, whereas states 3 and 4 can not be distinguished and are recorded as a single state, H . Here, the waiting time in H is the sum of the corresponding waiting times in the microstates 3 and 4 before jumping to one of the observed states.

2. Semi-coarse-grain (semi-CG)

In *semi-CG*, we assume an observer can record intratransitions within the hidden states [Fig. 1(c)]. For example, a sequence of $1 \rightarrow 4 \rightarrow 3 \rightarrow 2$ is recorded as $1 \rightarrow H \rightarrow H \rightarrow 2$, with the corresponding waiting times, i.e., the time spent in the first visit to H and the time spent in the second visit to H are recorded separately. In this case of observed intratransitions within a coarse-grained state, the initial and final microstates are not known, as both are lumped together to the same macrostate. Still, the added information can be utilized to improve the lower bound on the total EPR.

3. Reformulation

The plug-in and the RNEEP estimators rely on the sequence of states, while discarding waiting-time information, and can, therefore, be directly applied to the *semi-CG* trajectory. However, to apply the KLD estimator to the *semi-CG* data, we first need to reformulate the trajectory to harness the information of the WTD, by labeling the coarse-grained macrostates according to the number of times they are visited before jumping into an observed state. The transformation, depicted in Fig. 1(d), consists of two steps. First, we look for all the consecutive sequences of the hidden state H , and record their length, i.e., the number of visits to the state H . Second, all sequences with the same length are considered a new state, so a sequence of n appearances of H is labeled H_n . The waiting time associated with H_n is now the sum of the individual waiting times in the n consecutive appearances of H . Using the summation of the waiting times associated with each of the visits in H , instead of keeping the data of the duration of each consecutive visit, is necessitated by numerical limitations (Appendix A 1). Following this reformulation, we lose information about the full trajectory, and the KLD estimator applied to the reformulated data is a lower bound on the total EPR (Appendix A 2). We stress that the new representation is only used for applying the KLD estimator, whereas the plug-in and RNEEP estimators do not require waiting-time information and are used as reference methods that only exploit sequences of states rather than the WTD.

C. Entropy-production rate estimators

1. Plug-in estimator

The plug-in estimator, σ_{plug} , was proposed for approximating the KLD rate between the forward and reverse sequences

of discrete stationary time series, by counting sequences of data and calculating their probabilities [35,36]. The approximated m th-order KLD between sequences of length m is

$$D_m^x = \sum_{x_1, x_2, \dots, x_m} p(x_{1 \rightarrow m}) \ln \left(\frac{p(x_{1 \rightarrow m})}{p(x_{m \rightarrow 1})} \right), \quad (1)$$

where $p(x_{1 \rightarrow m})$ and $p(x_{m \rightarrow 1})$ are the probabilities of a forward sequence $x_{1 \rightarrow m} = (x_1, \dots, x_m)$ and the backward one $x_{m \rightarrow 1} = (x_m, \dots, x_1)$. These probabilities can be estimated from the number of appearances of each sequence in a long trajectory. Based on the approach in Ref. [35], the slope of D_m^x as a function of m ,

$$\hat{d}_m^x = D_m^x - D_{m-1}^x, \quad (2)$$

gives the entropy production per step in the limit of large m . However, for a non-Markov process that cannot be described by a semi-Markov process of any order, calculating \hat{d}_m^x is challenging for large values of m . Therefore, the following ansatz [66] has been proposed:

$$\hat{d}_m^x \simeq \hat{d}_\infty^x - c \frac{\ln(m)}{m^\gamma}, \quad (3)$$

where \hat{d}_∞^x , c , and γ , are the fit parameters for \hat{d}_m^x as a function of m . Our plug-in estimator for the entropy-production rate *per time* is thus

$$\sigma_{\text{plug}} = \frac{1}{\tau} \hat{d}_\infty^x, \quad (4)$$

where τ is the mean waiting time in each step.

2. KLD estimator

The KLD estimator, σ_{KLD} , derived by calculating the KLD between forward and reverse trajectories in semi-Markov processes, has two contributions [33]:

$$\sigma_{\text{KLD}} = \sigma_{\text{aff}} + \sigma_{\text{WTD}}, \quad (5)$$

where the affinity, σ_{aff} , stems from observed currents, and the σ_{WTD} stems from time-asymmetries in WTD. To apply Eq. (5) to second-order semi-Markov processes, the observed states are reformulated as doublets, $[ij]$, where the first index is the previous state, and the second index is the current state [33].

The affinity contribution is

$$\sigma_{\text{aff}} = \frac{1}{\tau} \sum_{i,j,k} p_{(ijk)} \ln \left(\frac{p_{([ij] \rightarrow [jk])}}{p_{([kj] \rightarrow [ji])}} \right), \quad (6)$$

where $p_{(ijk)}$ is the probability to observe the sequence of state $i \rightarrow j \rightarrow k$, or $p_{(ijk)} = p_{([ij] \rightarrow [jk])} R_{[ij]}$, with $p_{([ij] \rightarrow [jk])}$ being the probability to jump to state k after jumping from i to j , and $R_{[ij]}$ being the fraction of visits to $[i, j]$. The affinity, σ_{aff} , is governed by the relation between the forward and reverse transition probabilities.

The WTD contribution stems from the Kullback-Leibler divergence between WTD associated with forward ($i \rightarrow j \rightarrow k$) and backward ($k \rightarrow j \rightarrow i$) transitions:

$$\sigma_{\text{WTD}} = \frac{1}{\tau} \sum_{i,j,k} p_{(ijk)} D[\psi(t|[ij] \rightarrow [jk]) || \psi(t|[kj] \rightarrow [ji])], \quad (7)$$

where $\psi(t|[ij] \rightarrow [jk])$ is the WTD in state j given that the previous state was i and the following is k , τ is the average waiting time per state, and $D[u(x)||v(x)]$ is the Kullback-Leibler Divergence between two probability distributions, $u(x)$ and $v(x)$, defined as $D[u(x)||v(x)] = \sum_x u(x) \ln(u(x)/v(x))$. See Appendix B 1 for details regarding WTD estimation.

3. RNEEP estimator

The RNEEP estimator, σ_{RNEEP} , is formulated as an optimization problem [44], with a specific objective function to be minimized using stochastic gradient descent. The input of the problem is the set of all sequences of length m from a single long trajectory, and the solution is the coarse-grained entropy-production rate per step along the input trajectory. Similar to the plug-in estimator, the RNEEP uses the discrete sequence of states and does not exploit the WTD data, so estimating the full probability distributions of the waiting times is not required. Intuitively, this estimator should yield similar results to the plug-in estimator, Eq. (4), and to σ_{aff} , Eq. (6), as it uses the same information (see Appendix B 2 for further discussion). Note that in Ref. [44], the RNEEP was compared to a semianalytical calculation of the KLD between trajectory distributions. However, since the results of the plug-in estimator were similar to the semi-analytical values for the *semi-CG* statistics [35], we focus on the plug-in estimator in our work.

The RNEEP can be implemented by different machine learning models, such as recurrent or convolutional neural networks [44,45]. Following the approach of Ref. [44], we use a recurrent neural network, whose input is a sequence of some length m , $x_t^m = (x_t, x_{t+1}, \dots, x_{t+m-1})$, and its output is $h_\theta(x_t^m)$, where θ represents the learnable weights of the network. The output of the RNEEP is [44]

$$\Delta S_\theta(x_t^m) \equiv h_\theta(x_t^m) - h_\theta(\tilde{x}_t^m), \quad (8)$$

where \tilde{x}_t^m is the time-reversed sequence of x_t^m . The RNEEP estimator is the solution of the optimization problem of minimizing the following objective function over $\Delta S_\theta(x_t^m)$ for all possible sequences of length m :

$$\sigma_{\text{RNEEP},m} = \frac{1}{\tau} \min_{\Delta S_\theta} E_t E_{(x_t^m)} [\Delta S_\theta(x_t^m) - e^{-\Delta S_\theta(x_t^m)}], \quad (9)$$

where E_t is the expectation over t , and $E_{(x_t^m)}$ is the expectation over the observed sequences x_t^m . See Appendix B 2 for a detailed explanation regarding the implementation of the RNEEP estimator and the corresponding numerical considerations.

III. RESULTS

We have evaluated the three EPR estimators, plug-in, RNEEP, and KLD, across the three coarse-grained systems, including the four-state system, the molecular motor, and the flashing ratchet, subjected to the different CG schemes. Both *full-CG* and *semi-CG* approaches were applied to trajectories comprising approximately $N = 10^7$ states, simulated using the Gillespie algorithm [67], where the reformulation (*transformed-semi-CG*) was used for calculating σ_{KLD} for the




FIG. 4. Entropy-production rates for the 4-state system. Total EPR, σ_{tot} (dark red), KLD estimator, σ_{KLD} , for the *semi-CG* (dark blue) and *full-CG* (light blue) data, plug-in estimator, σ_{plug} , for the *semi-CG* (dark orange) and *full-CG* (light orange), RNEEP estimator, $\sigma_{\text{RNEEP},m}$, for the *semi-CG* (light to dark purple for increasing sequence length m) and *full-CG* (light to dark green for increasing sequence length m) data, and the affinity contribution, σ_{aff} , for the *full-CG* data (red). Rates: $\bar{\omega}_{12} = 2s^{-1}$, $\bar{\omega}_{21} = 3s^{-1}$, $\omega_{13} = 0s^{-1}$, $\omega_{14} = 1s^{-1}$, $\omega_{23} = 2s^{-1}$, $\omega_{24} = 35s^{-1}$, $\omega_{31} = 0s^{-1}$, $\omega_{32} = 50s^{-1}$, $\omega_{34} = 0.7s^{-1}$, $\omega_{41} = 8s^{-1}$, $\omega_{42} = 0.2s^{-1}$, $\omega_{43} = 75s^{-1}$. Error bars represent standard deviation. In cases where error bars are not visible, they are smaller than the size of the marker.

semi-CG statistics. Each reported result for any system originates from 10 distinct randomly generated trajectories with the specified parameter along the x axis.

For all tested systems, we find the results of the EPR estimators calculated on the *semi-CG* data better than the ones calculated on the *full-CG* data. This comparison serves as compelling evidence for the importance of appropriate utilization of the available data for extracting tight bounds on the total EPR.

A. Four-state system

The four-state system (Fig. 1) has two observed states, 1 and 2, where the states 3 and 4 are coarse-grained into a single state H . The rates of the observed link, 1 – 2, between the two observed microstates are tuned according to $\omega_{12} = \bar{\omega}_{12}e^x$ and $\omega_{21} = \bar{\omega}_{21}e^{-x}$, where ω_{ij} is the transition rate from j to i , to mimic an external forcing, where the range of x was chosen to include the stalling force in which there is no observable current over the 1-2 link [18]. The results for the four-state system under the two CG schemes, *semi-CG* and *full-CG*, and the three estimators, RNEEP, plug-in, and KLD, are presented in Fig. 4, where the stalling force corresponds to the x value for which the affinity contribution, σ_{aff} , vanishes.

As expected, the bounds on the total EPR obtained from estimators applied to the *semi-CG* statistics are respectively better compared with the same estimators applied to the *full-CG* trajectories. In the *full-CG* case, $\sigma_{\text{RNEEP},m}$ converges with increasing sequence length, m to the affinity, σ_{aff} , as both estimators use the same data. The plug-in estimator, σ_{plug} , which also uses the same data of the *full-CG* trajectory, provides similar results as σ_{aff} away from the stalling force. However, close to the stalling force, where σ_{aff} vanishes, σ_{plug} provides nonzero values that stem from the inherent bias of the method,

which assigns positive values to all the probabilities [36] (see Appendix B 2). The KLD estimator, σ_{KLD} , provides the tightest lower bound for the *full-CG* data, as it is the only one that utilizes information of the irreversibility in WTD.

Calculated on the *semi-CG* data, the RNEEP estimator, $\sigma_{\text{RNEEP},m}$, provides a tighter bound for increasing sequence lengths m , converging to approximately similar values as the KLD estimator, applied to the *transformed semi-CG* trajectories (Appendix B 3). Above sequence length of $m \approx 20$, the values of σ_{KLD} are mostly within the standard deviation of the $\sigma_{\text{RNEEP},m}$ results. In this case, both the KLD and RNEEP surpass the plug-in estimator, where the values of σ_{plug} are similar to $\sigma_{\text{RNEEP},8}$, indicating that $\sigma_{\text{RNEEP},m}$ for larger m , and σ_{KLD} calculated on the post-reformulation trajectory, effectively utilize longer sequence data. The relationship between the *semi-CG* and the *transformed semi-CG* data is considerably complex. While in the case of the *transformed semi-CG*, we gain information from the WTD, the KLD estimator interprets each H_n state post-reformulation as distinct, despite their origination from the same subset of states with varying numbers of visits. Conversely, in the case of the *semi-CG* data, the labeling of hidden states remains invariant to the number of their successive appearances, underlining a fundamental difference in the treatment of states between the two models.

The difference in σ_{KLD} between the *transformed semi-CG* and the *full-CG* schemes reflects the additional information regarding irreversibility encoded in the intratransitions between microstates in the hidden macrostate. Moreover, the irreversibility encoded in these intratransitions in the *semi-CG* data is also reflected in the values of the three estimators that do not vary significantly for x values near the stall force, in contrast to the estimators applied to the *full-CG* trajectories that strongly depend on the deviation of x from stalling conditions.




FIG. 5. Entropy-production rates for the molecular motor model. Total EPR, σ_{tot} (dark red) of the full trajectory, KLD estimator, σ_{KLD} (dark blue), plug-in estimator, σ_{plug} (dark orange), and RNEEP estimator, $\sigma_{\text{RNEEP},m}$ (light to dark purple for increasing sequence length m) for *semi-CG* data. The transition rates for $\forall i \neq j$ are $\omega_{ij} = \omega_{i'i} = \omega_{i'j} = \omega_{ij'} = 1$, and $\omega_{ij} = 0.01$, where states i (i') represent the passive (active) states of the motor. The driving force affects the rates by the following rule: $\frac{\omega_{ij}}{\omega_{ji}} = e^F$ and $\frac{\omega_{i'j}}{\omega_{j'i'}} = e^{F-\Delta\mu}$. The three presented branches, from left to right, are for $\Delta\mu = 1, 2, 3$, respectively. Error bars represent standard deviation. In cases where error bars are not visible, they are smaller than the size of the marker.

It is crucial to highlight not just the precision of the estimator but also the complexity and robustness of the results. While the KLD estimator provided a lower bound on the total EPR by construction (Appendix A 2), the RNEEP estimator requires dedicated hyperparameter tuning for each system, as improper adjustments of these parameters could lead to an overestimation of the EPR (Appendix B 3). Furthermore, computing $\sigma_{\text{RNEEP},m}$ is considerably more resource-intensive and time-consuming from a computational standpoint, compared to the calculation of σ_{KLD} .

B. Molecular motor

The molecular motor moves “up” or “down” along a one-dimensional line in discrete steps. For each physical position, the motor can be either active or passive, where in the former case, the upward jump is favored by a chemical potential $\Delta\mu$. An external force F acts downward, against the preferred direction. We assume an external observer cannot distinguish between the active and passive states and can only record the position of the motor. The results for the molecular motor system under the *full-CG* and the *semi-CG* schemes with the three EPR estimators, RNEEP, plug-in, and KLD, are presented in Fig. 5 for three values of $\Delta\mu$ and a range of F values which include the respective stalling conditions, in which there is no net movement of the motor and the affinity estimator, σ_{aff} , vanishes.

Similar to the four-state system, the estimators calculated on the *semi-CG* data provide tighter bounds compared to their respective counterparts calculated on the *full-CG* data, underscoring the contribution of the additional information of the intratransition within the macrostates exploited in the *semi-CG* scheme. In the *full-CG* trajectories, the KLD, plug-in, and RNEEP estimators provide similar values, except near

the stalling force, where σ_{KLD} [33,49] is consistently a tighter estimator than σ_{plug} and $\sigma_{\text{RNEEP},m}$, as we saw in the four-state system.

In the case of the *semi-CG* scheme, the σ_{KLD} , calculated on the *transformed semi-CG* trajectory, provides the tightest lower bound on the total EPR compared to the plug-in [35,36] and the RNEEP [44]. While this result is not trivial, as we saw in the four-state system in which the RNEEP estimator was comparable to the KLD values, it emphasizes that our method of reformulation can potentially exploit viable information and produce tighter lower bounds in some cases. The RNEEP estimator, $\sigma_{\text{RNEEP},m}$, provides a tighter bound for increasing sequence lengths m , similar to the four-state system (Appendix B 3). Above sequence length of $m \approx 64$, the values of σ_{plug} are mostly within the standard deviation of $\sigma_{\text{RNEEP},m}$.

The difference in σ_{KLD} between the *transformed semi-CG* and the *full-CG* frameworks reaffirms the unique encoding of irreversibility within the hidden transitions. This distinction highlights how irreversibility is reflected in the *semi-CG* data, evident in the consistent behavior of the estimators near the stalling force. In contrast, the estimation on the *full-CG* trajectories strongly relies on the proximity to stalling conditions.

C. Flashing ratchet

The flashing ratchet models a Brownian particle moving along a periodic one-dimensional line under the influence of a linear potential V that can be switched on and off at a constant rate [35,36]. The particle is described by its position in the “on,” i , or “off,” i' , states. The information regarding the potential is not accessible, and both the on and off states, i and i' , are lumped into a single macrostate H_i . The results for the discrete flashing ratchet system under the *semi-CG* scheme, which includes the information of the intratransition within




FIG. 6. Entropy-production rate estimators for the flashing ratchet calculated on the *semi-CG* trajectory data as a function of the potential V . Total EPR, σ_{tot} (dark red) of the full trajectory, KLD estimator, σ_{KLD} (dark blue), plug-in estimator, σ_{plug} (dark orange), and RNEEP estimator, $\sigma_{\text{RNEEP},m}$ (light to dark purple for increasing sequence length m) for *semi-CG* data. The transition rates for $\forall i \neq j$ are: $\omega_{i'i'} = \omega_{j'j} = \omega_{j'i} = 1$, $\omega_{ij} = e^{(V_j - V_i)/2}$, where states i (i') correspond to the turned-on (off) potential. Error bars represent standard deviation. In cases where error bars are not visible, they are smaller than the size of the marker.

the macrostates, and the three estimators, RNEEP, plug-in, and KLD, are presented in Fig. 6 as a function of the potential value, V . Since the estimators applied to the *full-CG* scheme are several orders of magnitude smaller, we present them separately in Appendix C. Here, too, we validate the importance of the amount of available information in the context of EPR estimation.

As expected, all estimators increase monotonically with increasing values of V , corresponding to larger total EPR values. However, for the case of the flashing ratchet system with the *semi-CG* data, in contrast to the previous examples, we find that the plug-in estimator, σ_{plug} , provides the tightest lower bound on the total EPR, suppressing the KLD estimator, σ_{KLD} . The RNEEP estimator, $\sigma_{\text{RNEEP},m}$, yields a tighter bound for increasing sequence lengths m , converging to the σ_{plug} value for $m \approx 128$.

The difference in the relation between the estimators across the various systems tested could be attributed to the different topologies. For example, while both the molecular motor and the flashing ratchet represent cyclic processes, upward transitions in the molecular motor system, either from an active or inactive state, can only arrive at an inactive state, whereas in the flashing ratchet system, transitions arrive at either on or off states of the potential from a neighboring position with a similar potential status. These distinctions might yield different contributions of the WTD data extracted from the *transformed semi-CG* information.

D. Statistical comparison

To underscore the performance of the three estimators, and specifically the contribution of our proposed data reformulation to include the WTD of intratransitions, we randomly sampled 50 fully connected four-state systems. We sampled the rates independently from a uniform distribution in the range $[0,50]$, and calculated the plug-in and RNEEP

estimators for the *semi-CG* framework, and the KLD estimator for the *transformed semi-CG* data. The results are summarized in Fig. 7.

Evidently, the $\sigma_{\text{RNEEP},128}$ resulted in an overestimation of the total EPR in 10 (20%) systems. In these cases, the plug-in estimator yielded tighter EPR bounds compared to the KLD in 6 of the 10 systems. We stress that in this analysis, the hyperparameters of the RNEEP estimator were not separately adjusted for each of the 50 randomly sampled systems, resulting in the overestimation of the total EPR for some of the tested systems. Out of the remaining 40 systems, the tightest lower bound on the total EPR was obtained by the $\sigma_{\text{RNEEP},128}$, σ_{KLD} , and σ_{plug} estimators for 20, 17, and 3 of the cases, respectively. Comparing the results of the KLD and plug-in estimators for all systems tested, we find that σ_{KLD} resulted in a better estimation in 40 (80%) cases, compared to the σ_{plug} values.

While the RNEEP estimator resulted in the tightest lower bound in some of the cases, the requirement for meticulous calibration per system, and susceptibility to overestimation, render its general applicability challenging. Nevertheless, our proposed data reformulation, for cases in which intratransitions are available has proven to be robust, independent from system specifics, and a preferable tool for bounding the total EPR.

IV. CONCLUSIONS

In conclusion, we performed a rigorous comparison of time-irreversibility-based EPR estimators, with a primary focus on KLD-based approaches, encompassing both KLD estimators with WTD statistics and those without (specifically, plug-in and RNEEP), while employing two distinct coarse-graining methodologies. Our findings affirm that the *semi-coarse-graining* framework, which incorporates intratransitions data, consistently yields tighter EPR bounds




FIG. 7. EPR estimation for 50 fully connected four-state systems, whose rates were randomly sampled from a uniform distribution in the range $[0,50]$. The estimators were evaluated on trajectory realizations with $N = 10^7$ steps, coarse-grained under the *semi-CG* scheme. The results of the σ_{KLD} (blue), σ_{plug} (orange), and $\sigma_{\text{RNEEP},128}$, are shown in comparison to the hotel EPR, σ_{tot} (black line). In some cases, $\sigma_{\text{RNEEP},128}$ overestimates σ_{tot} .

compared to the *full-coarse-graining* framework. This outcome underscores the capacity of the former to harness a richer source of time-irreversibility information. Additionally, a direct comparison of the EPR bounds derived from *full-CG* and *semi-CG* statistics offers a quantitative assessment of the time-irreversibility embedded within the intratransitions as captured by each estimator.

Moreover, we have introduced an innovative approach for the reformulation of *semi-coarse-grained* trajectories, a method previously tailored for the plug-in and RNEEP estimators, to accommodate the KLD estimator. The application of this *transformed-semi-CG* approach, together with the KLD estimator, has exhibited varying degrees of success. While it has proven effective in distilling time-irreversibility information within the intratransitions of lumped microstates for certain cases, the RNEEP or the plug-in estimators outperformed it in others. This could be attributed to the treatment of newly introduced states in the reformulation corresponding to the number of consecutive visits in the hidden states, in which the partial utilization of the WTD statistics may not suffice, as we exclude information about individual residence times between the recorded intratransitions.

It is essential to reiterate the intricate relationship between the *semi-CG* and *transformed-semi-CG* datasets. The incorporation of the WTD data in the *transformed-semi-CG* trajectories contributes to the ability to detect time-irreversibility. However, the KLD estimator recognizes each post-transformation H_n state as unique, even though they all stem from an identical group of states, differing only in the number of consecutive visits. This is in contrast to the *semi-CG* approach, where the classification of hidden states is not influenced by the number of their appearances. This key distinction underpins a fundamental difference in state treatment across the two models, possibly leading to a varying tightness of the resulting bound on the total EPR. Consequently, in the *semi-CG* case, we recommend opting for the maximum

of the three estimators, bearing in mind the need for careful calibration of the RNEEP calculation to avoid overestimation in this case.

It is worth noting that the proposed transformations and EPR estimation techniques showcased here hold broader applicability, extending to other continuous-time systems when only partial information is available, and the observations are coarse-grained. Further exploration of the governing factors influencing post-reformulation results is a promising avenue for future research.

ACKNOWLEDGMENTS

G.B. acknowledges the Zuckerman STEM Leadership Program, and the Tel Aviv University Center for AI and Data Science (TAD). This work was supported by the ERC NanoNonEq 101039127, the Air Force Office of Scientific Research (AFOSR) under Award No. FA9550-20-1-0426, and by the Army Research Office (ARO) under Grant No. W911NF-21-1-0101. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Office or the U.S. Government.

APPENDIX A: REFORMULATION

1. Implementation limitations

As depicted in Fig. 1, our reformulation method aims to manipulate trajectory data to enhance the utilization of available information, thereby refining the estimation of EPR. Our approach seeks to leverage the waiting-time distribution (WTD) within the *semi-CG* framework for employing the KLD estimator. By creating a new, effective, hidden state for each consecutive sequence length of hidden states within the original *semi-CG* trajectory, we assign an effective waiting time for the new state, which is the cumulative sum of all

component waiting times. Attempting to handle vectors of the individual waiting time becomes impractical when the new hidden state comprises a sequence of more than four jumps, as the WTD becomes at least four-dimensional in this scenario. Consequently, obtaining adequate samples to sufficiently reconstruct the distribution becomes challenging. Furthermore, varying states within the reformulation would exhibit distinct dimensionalities in their WTD, necessitating tailored adjustments. Therefore, we note that while we gain information about the WTD by applying such reformulation, some pertinent information present in the *semi-CG* states is lost in the new effective hidden states.

2. Lower bound

Here, we detail the proof that calculating the KLD between forward and reverse trajectories under coarse-graining results in a lower bound on the total EPR, being directly linked to the KLD between forward and reverse trajectories with complete information [17,33]. Note that in our tested systems, the assumption that the coarse-graining scheme and the time-reversal commute [42], holds. While this proof aims to support our new reformulation approach for the *semi-CG* data, termed *transformed-semi-CG*, to allow for the calculation of the KLD estimator when intratransitions information is available, it is valid for both the *semi-CG* and *full-CG* schemes.

Let γ denote the full observable trajectory with complete information of states and WTD, and let γ_{cg} denote the coarse-grained trajectory, $\gamma_{cg} = CG(\gamma)$. Moreover, $\tilde{\gamma}$, $\tilde{\gamma}_{cg}$ represent the reversed full and coarse-grained trajectories, respectively. We now apply the chain rule for the relative entropy on the KLD between the joint distributions of γ and γ_{cg} [68,69]:

$$\begin{aligned} D[P(\gamma, \gamma_{cg})||P(\tilde{\gamma}, \tilde{\gamma}_{cg})] \\ &= D[P(\gamma)||P(\tilde{\gamma})] + D[P(\gamma_{cg}|\gamma)||P(\tilde{\gamma}_{cg}|\tilde{\gamma})] \\ &= D[P(\gamma_{cg})||P(\tilde{\gamma}_{cg})] + D[P(\gamma|\gamma_{cg})||P(\tilde{\gamma}|\tilde{\gamma}_{cg})]. \end{aligned} \quad (A1)$$

The crucial aspect of our reformulation approach is that the transformation of a trajectory into a new representation, based on the number of consecutive visits in a hidden state, is uniquely defined by the original trajectory. This means that for any given initial trajectory, there is a specific and unambiguous way to represent its *transformed-semi-CG* counterpart using our suggested scheme. Therefore, as the coarse-grained trajectory, γ_{cg} , is uniquely determined given the full trajectory, γ , we have

$$P(\gamma_{cg}|\gamma) = \begin{cases} 1, & \gamma_{cg} = CG(\gamma), \\ 0, & \text{otherwise,} \end{cases} \quad (A2)$$

and similarly for $\tilde{\gamma}_{cg}$ and $\tilde{\gamma}$,

$$P(\tilde{\gamma}_{cg}|\tilde{\gamma}) = \begin{cases} 1, & \tilde{\gamma}_{cg} = CG(\tilde{\gamma}), \\ 0, & \text{otherwise,} \end{cases} \quad (A3)$$

which implies

$$D[P(\gamma|\gamma_{cg})||P(\tilde{\gamma}|\tilde{\gamma}_{cg})] = 0. \quad (A4)$$

Plugging this result into Eq. (A1), and invoking the non-negativity of the relative entropy, we arrive at the required result:

$$D[P(\gamma)||P(\tilde{\gamma})] \geq D[P(\gamma_{cg})||P(\tilde{\gamma}_{cg})]. \quad (A5)$$

APPENDIX B: NUMERICAL CONSIDERATIONS

1. Waiting-time distributions estimation

The KLD estimator, σ_{KLD} , has two contributions, namely, σ_{aff} and σ_{WTD} . While σ_{aff} can be directly calculated by counting second-order transitions, σ_{WTD} requires the estimation of continuous functions. To numerically evaluate continuous probability density functions, we use the kernel density estimation (KDE) method [70]. In this approach, the estimated function depends on the bandwidth of the kernel, and the optimal bandwidth is correlated to the sample size. We used

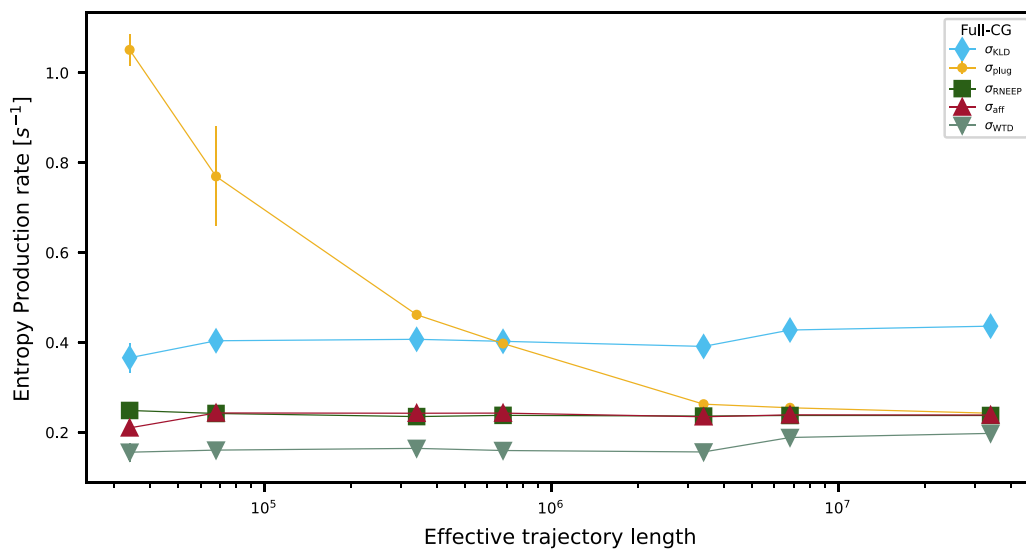


FIG. 8. Convergence of EPR estimators, σ_{KLD} (blue diamonds), σ_{plug} (light orange dots), $\sigma_{\text{RNEEP},128}$ (dark green squares), σ_{aff} (red triangles), and σ_{WTD} (gray triangles), as a function of effective trajectory length, calculated for the four-states system and $x = -0.15$. The effective trajectory length is calculated for the *full-CG* data. Each data point of the KLD, plug-in, and RNEEP estimators is the mean of 10 realizations of trajectories of varying lengths between $N = 5 \times 10^4$ to $N = 5 \times 10^7$ states. Error bars represent the standard deviation.

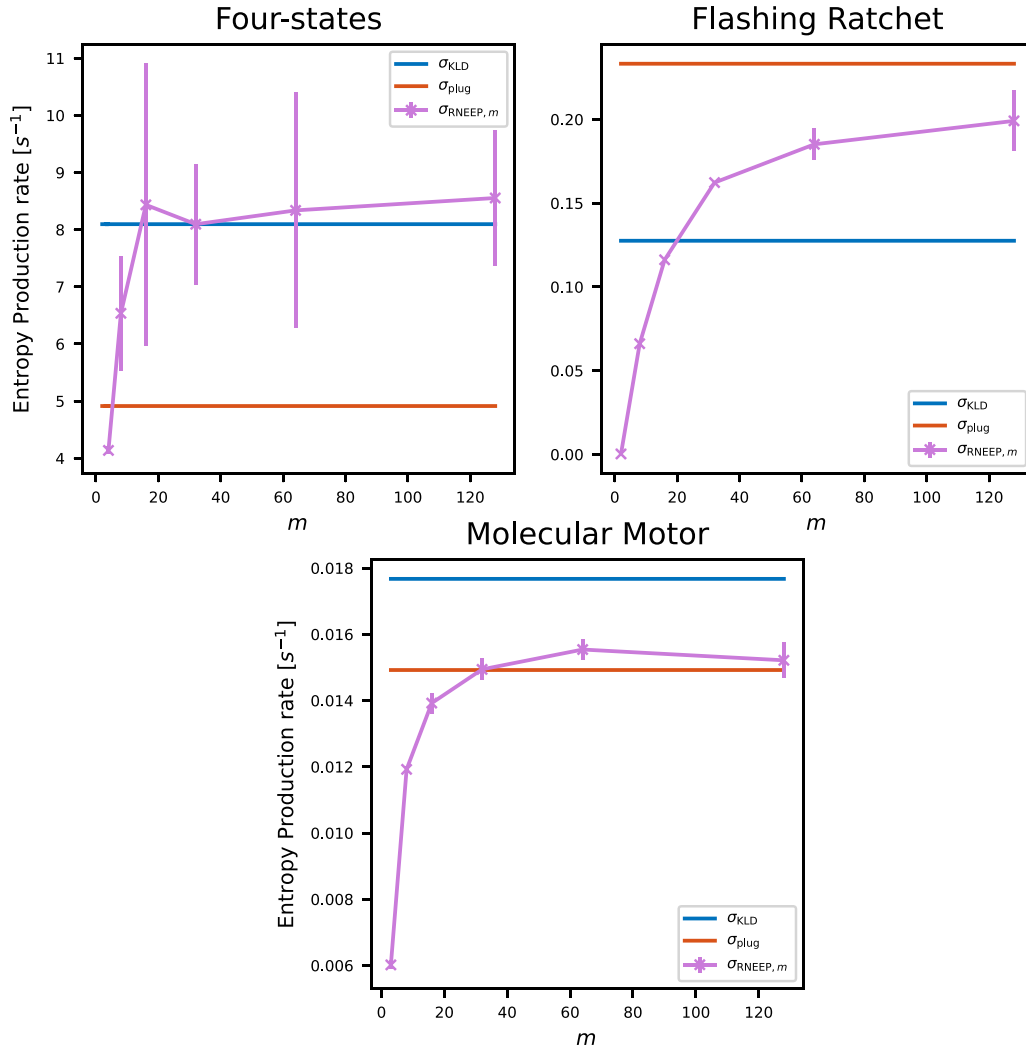


FIG. 9. Convergence of $\sigma_{\text{RNEEP},m}$ as a function of the input sequence length m (purple) for *semi-CG* data, compared to the values of the σ_{KLD} (blue), and σ_{plug} (orange), which do not depend of m , for the four-state system with $x = 0.45$ (top left), the flashing ratchet with $V = 2$ (top right), and the molecular motor with $F = 2$ and $\Delta\mu = 2$ (bottom). Each data point of the RNEEP estimator as a function of the input sequence length, $\sigma_{\text{RNEEP},m}$, is the mean of the RNEEP calculation for 10 trajectory realizations of length $N = 10^7$ for all systems. The error bars correspond to the standard deviation.

Silverman's rule of thumb [71] to choose the bandwidth. In our case, the sample size is the number of observed second-order jumps for each WTD, which can lead to a large variation between sample sizes for different transitions. We chose three different kernels, each for a different range of sample size, $2 \times 10^2 - 5 \times 10^3$, $5 \times 10^3 - 10^5$, and $> 10^5$. A sample size of less than 2×10^2 was not considered due to the lack of statistics.

The grid size of the KDE was also chosen empirically for optimized convergence to the WTD, considering the computational cost. The estimation of the WTD is better for longer trajectories, whereas short trajectories can result in an inaccurate estimation of the EPR due to insufficient statistics (Fig. 8). Moreover, the required trajectory length for a desired tolerance of the estimation depends on the system parameters [33].

2. Plug-in estimator implementation

The calculation of the plug-in estimator, σ_{plug} , implemented according to Ref. [35], involves log ratios of

probabilities for forward and reverse sequences. However, in cases where a forward sequence occurs while its reverse is missing or lacks sufficient statistics, computing the log ratio becomes impossible. We must, therefore, resort to numerical methods to handle these instances. As proposed in Ref. [35], a small bias could be added to the number of observations of each sequence, preventing the probability from being zero. However, this approach may introduce overall biases and potentially overestimate the results if not meticulously calibrated. Instead, we chose to exclude pairs of forward and reverse sequences from the analysis if any second-order transitions in either direction were not observed, to ensure a reliable and stable lower bound for our calculations.

We note the observed difference between σ_{plug} and σ_{aff} for the *full-CG* scheme (Figs. 4 and 5), despite them relying on the same underlying information. This divergence stems from the distinct computational methods employed, where σ_{plug} is derived from a fitting function, whereas σ_{aff} is calculated directly from the log-ratio. This fundamental difference accounts for

the variations seen between σ_{plug} and σ_{aff} values, especially noticeable around the stalling force.

Figure 8 presents σ_{plug} estimator results as a function of the effective trajectory length, emphasizing the role of sufficient statistics for robust analysis. Longer trajectories are important for the convergence of σ_{plug} , thereby providing more reliable data for the calculation and capturing the true behavior of the system.

3. RNEEP convergence

The RNEEP estimator, implemented according to Ref. [44], was calculated on 12 training trajectory realizations of length $N = 10^7$ run in parallel on eight *Geforce RTX 2080 Ti* GPUs. Evaluating the RNEEP estimator for a trajectory input required approximately 20 minutes, where the entire data set was generated in a few hours.

The $\sigma_{\text{RNEEP},m}$ estimator gives a tighter bound on the total EPR for increasing sequence length, m , up to saturation for $m \approx 20$, $m \approx 64$, and $m \approx 128$, for the *semi-CG* of the four-state system, the molecular motor, and the flashing ratchet, respectively (Fig. 9).

A significant limitation of the RNEEP estimator is the need for recalibration of hyperparameters for each system under study. This recalibration is crucial since improper management of the training process might lead to unreliable results, when the training parameters are not meticulously adjusted for each system. In our statistical analysis of the 50 randomly sampled four-state systems (Fig. 7), the RNEEP model was trained for the various systems without making any modifications to its hyperparameters, resulting in an overestimate of the total EPR in a few cases. This approach revealed that the RNEEP estimator required careful and system-specific tuning to ensure accurate results.

APPENDIX C: EPR ESTIMATORS ON *full-CG* FLASHING RATCHET TRAJECTORIES

The results of the KLD, plug-in, and RNEEP estimators for the flashing ratchet calculated on the *full-CG* trajectory data are presented in Fig. 10. Noticeably, these results are approximately three orders of magnitude smaller compared to the same estimators calculated on the *semi-CG* data (see Fig. 6). This substantial difference further underscores the




FIG. 10. Entropy-production rate estimators for the flashing ratchet calculated on the *full-CG* trajectory data. KLD estimator, σ_{KLD} (light blue), plug-in estimator, σ_{plug} (light orange), and RNEEP estimator, $\sigma_{\text{RNEEP},m}$ (light to dark green for increasing sequence length m). The transition rates for $\forall i \neq j$ are: $\omega_{i'i'} = \omega_{i'i} = \omega_{i'j} = 1$, $\omega_{ij} = e^{(V_j - V_i)/2}$, where states i (i') correspond to the turned-on (off) potential.

sensitivity of EPR estimations to the level of available information, suggesting that the finer details captured in the *semi-CG* scheme, which are absent in the *full-CG* framework, can be exploited to provide a tighter lower bound on the total EPR.

Among the tested estimators on the *full-CG* data of the flashing ratchet, the KLD estimator, σ_{KLD} , provides the tightest bound, followed by the plug-in estimator, σ_{plug} , and the RNEEP estimator, $\sigma_{\text{RNEEP},m}$, whose values increase for an increasing sequence length, m . These results, and the relation between the values of the various estimators, align with our findings for the four-state system and the molecular motor for the *full-CG* scheme.

- [1] S. K. Manikandan, T. Ghosh, T. Mandal, A. Biswas, B. Sinha, and D. Mitra, Estimate of entropy generation rate can spatiotemporally resolve the active nature of cell flickering, [arXiv:2205.12849](#).
- [2] F. S. Gnesotto, F. Mura, J. Gladrow, and C. P. Broedersz, Broken detailed balance and non-equilibrium dynamics in living systems: a review, *Rep. Prog. Phys.* **81**, 066601 (2018).
- [3] C. Jarzynski, Equalities and inequalities: Irreversibility and the second law of thermodynamics at the nanoscale, *Annu. Rev. Condens. Matter Phys.* **2**, 329 (2011).
- [4] J. M. R. Parrondo and B. J. de Cisneros, Energetics of Brownian motors: a review, *Appl. Phys. A* **75**, 179 (2002).
- [5] M. Schliwa and G. Woehlke, Molecular motors, *Nature (London)* **422**, 759 (2003).
- [6] S. Ciliberto, A. Imparato, A. Naert, and M. Tanase, Heat flux and entropy produced by thermal fluctuations, *Phys. Rev. Lett.* **110**, 180601 (2013).
- [7] M. Esposito and C. Van den Broeck, Three detailed fluctuation theorems, *Phys. Rev. Lett.* **104**, 090601 (2010).
- [8] U. Seifert, Stochastic thermodynamics, fluctuation theorems and molecular machines, *Rep. Prog. Phys.* **75**, 126001 (2012).
- [9] D. J. Evans and D. J. Searles, The fluctuation theorem, *Adv. Phys.* **51**, 1529 (2002).
- [10] E. M. Sevick, R. Prabhakar, S. R. Williams, and D. J. Searles, Fluctuation theorems, *Annu. Rev. Phys. Chem.* **59**, 603 (2008).

- [11] N. Shiraishi and T. Sagawa, Fluctuation theorem for partially masked nonequilibrium dynamics, *Phys. Rev. E* **91**, 012130 (2015).
- [12] M. Poletini and M. Esposito, Effective fluctuation and response theory, *J. Stat. Phys.* **176**, 94 (2019).
- [13] F. Zamponi, Is it possible to experimentally verify the fluctuation relation? A review of theoretical motivations and numerical evidence, *J. Stat. Mech.: Theory Exp.* (2007) P02008.
- [14] G. M. Wang, E. M. Sevick, E. Mittag, D. J. Searles, and D. J. Evans, Experimental demonstration of violations of the second law of thermodynamics for small systems and short time scales, *Phys. Rev. Lett.* **89**, 050601 (2002).
- [15] D. Collin, F. Ritort, C. Jarzynski, S. Smith, I. Tinôco, Jr., and C. Bustamante, Verification of the Crooks fluctuation theorem and recovery of RNA folding free energies, *Nature (London)* **437**, 231 (2005).
- [16] I. D. Terlizzi, M. Gironella, D. Herráez-Aguilar, T. Betz, F. Monroy, M. Baiesi, and F. Ritort, Variance sum rule for entropy production, *Science* **383**, 971 (2024).
- [17] R. Kawai, J. M. R. Parrondo, and C. Van den Broeck, Dissipation: The phase-space perspective, *Phys. Rev. Lett.* **98**, 080602 (2007).
- [18] G. Bisker, M. Poletini, T. R. Gingrich, and J. M. Horowitz, Hierarchical bounds on entropy production inferred from partial information, *J. Stat. Mech.: Theory Exp.* (2017) 093210.
- [19] D. Seiferth, P. Sollich, and S. Klumpp, Coarse graining of biochemical systems described by discrete stochastic dynamics, *Phys. Rev. E* **102**, 062149 (2020).
- [20] P. Bilotto, L. Caprini, and A. Vulpiani, Excess and loss of entropy production for different levels of coarse graining, *Phys. Rev. E* **104**, 024140 (2021).
- [21] A. Puglisi, S. Pigolotti, L. Rondoni, and A. Vulpiani, Entropy production and coarse graining in Markov processes, *J. Stat. Mech.: Theory Exp.* (2010) P05015.
- [22] G. Teza and A. L. Stella, Exact coarse graining preserves entropy production out of equilibrium, *Phys. Rev. Lett.* **125**, 110601 (2020).
- [23] M. Esposito, Stochastic thermodynamics under coarse graining, *Phys. Rev. E* **85**, 041125 (2012).
- [24] M. Poletini and M. Esposito, Effective thermodynamics for a marginal observer, *Phys. Rev. Lett.* **119**, 240601 (2017).
- [25] A. Ghosal and G. Bisker, Entropy-production rates for different notions of partial information, *J. Phys. D* **56**, 254001 (2023).
- [26] P. E. Harunari, A. Dutta, M. Poletini, and E. Roldán, What to learn from a few visible transitions' statistics? *Phys. Rev. X* **12**, 041026 (2022).
- [27] 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).
- [28] 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).
- [29] B. Ertel and U. Seifert, An estimator of entropy production for partially accessible Markov networks based on the observation of blurred transitions, [arXiv:2312.08246](https://arxiv.org/abs/2312.08246).
- [30] A. Gomez-Marin, J. M. R. Parrondo, and C. Van den Broeck, Lower bounds on dissipation upon coarse graining, *Phys. Rev. E* **78**, 011107 (2008).
- [31] S. Bo and A. Celani, Entropy production in stochastic systems with fast and slow timescales, *J. Stat. Phys.* **154**, 1325 (2014).
- [32] J. M. Parrondo, C. Broeck, and R. Kawai, Entropy production and the arrow of time, *New J. Phys.* **11**, 073008 (2009).
- [33] I. A. Martínez, G. Bisker, J. M. Horowitz, and J. M. R. Parrondo, Inferring broken detailed balance in the absence of observable currents., *Nat. Commun.* **10**, 3542 (2019).
- [34] Z. Rached, F. Alajaji, and L. L. Campbell, The Kullback-Leibler divergence rate between Markov sources, *IEEE Trans. Inf. Theory* **50**, 917 (2004).
- [35] E. Roldán and J. M. R. Parrondo, Entropy production and Kullback-Leibler divergence between stationary trajectories of discrete systems, *Phys. Rev. E* **85**, 031129 (2012).
- [36] E. Roldán and J. M. R. Parrondo, Estimating dissipation from single stationary trajectories, *Phys. Rev. Lett.* **105**, 150607 (2010).
- [37] A. Ghosal and G. Bisker, Inferring entropy-production rate from partially observed Langevin dynamics under coarse-graining, *Phys. Chem. Chem. Phys.* **24**, 24021 (2022).
- [38] 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).
- [39] E. Meyberg, J. Degünther, and U. Seifert, Entropy production from waiting-time distributions for overdamped Langevin dynamics, [arXiv:2402.18155](https://arxiv.org/abs/2402.18155).
- [40] Q. Wang, S. R. Kulkarni, and S. Verdú, Divergence estimation of continuous distributions based on data-dependent partitions, *IEEE Trans. Inf. Theory* **51**, 3064 (2005).
- [41] A. M. Berezhkovskii and D. E. Makarov, On the forward/backward symmetry of transition path time distributions in nonequilibrium systems, *J. Chem. Phys.* **151**, 065102 (2019).
- [42] G. Bisker, I. A. Martinez, J. M. Horowitz, and J. M. R. Parrondo, Comment on "Inferring broken detailed balance in the absence of observable currents", [arXiv:2202.02064](https://arxiv.org/abs/2202.02064).
- [43] D. Hartich and A. Godec, Comment on "Inferring broken detailed balance in the absence of observable currents", [arXiv:2112.08978](https://arxiv.org/abs/2112.08978).
- [44] D.-K. Kim, Y. Bae, S. Lee, and H. Jeong, Learning entropy production via neural networks, *Phys. Rev. Lett.* **125**, 140604 (2020).
- [45] Y. Bae, D.-K. Kim, and H. Jeong, Inferring dissipation maps from videos using convolutional neural networks, *Phys. Rev. Res.* **4**, 033094 (2022).
- [46] S. Otsubo, S. Manikandan, and T. Sagawa, Estimate non-equilibrium trajectories, *Commun. Phys.* **5**, 11 (2022).
- [47] D. J. Skinner and J. Dunkel, Improved bounds on entropy production in living systems, *Proc. Natl. Acad. Sci. USA* **118**, e2024300118 (2021).
- [48] D. J. Skinner and J. Dunkel, Estimating entropy production from waiting-time distributions, *Phys. Rev. Lett.* **127**, 198101 (2021).
- [49] E. Nitzan, A. Ghosal, and G. Bisker, Universal bounds on entropy production inferred from observed statistics, *Phys. Rev. Res.* **5**, 043251 (2023).
- [50] J. Ehrich, Tightest bound on hidden entropy production from partially observed dynamics, *J. Stat. Mech.: Theory Exp.* (2021) 083214.

- [51] N. Shiraishi, Optimal thermodynamic uncertainty relation in Markov jump processes, *J. Stat. Phys.* **185**, 19 (2021).
- [52] J. M. Horowitz and T. R. Gingrich, Thermodynamic uncertainty relations constrain non-equilibrium fluctuations, *Nat. Phys.* **16**, 15 (2020).
- [53] G. Falasco, M. Esposito, and J.-C. Delvenne, Unifying thermodynamic uncertainty relations, *New J. Phys.* **22**, 053046 (2020).
- [54] A. C. Barato and U. Seifert, Thermodynamic uncertainty relation for biomolecular processes, *Phys. Rev. Lett.* **114**, 158101 (2015).
- [55] 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).
- [56] U. Seifert, From stochastic thermodynamics to thermodynamic inference, *Annu. Rev. Condens. Matter Phys.* **10**, 171 (2019).
- [57] T. V. Vu and Y. Hasegawa, Generalized uncertainty relations for semi-Markov processes, *J. Phys.: Conf. Ser.* **1593**, 012006 (2020).
- [58] B. Ertel, J. van der Meer, and U. Seifert, Operationally accessible uncertainty relations for thermodynamically consistent semi-Markov processes, *Phys. Rev. E* **105**, 044113 (2022).
- [59] T. Kamijima, S. Otsubo, Y. Ashida, and T. Sagawa, Higher-order efficiency bound and its application to nonlinear nanothermoelectrics, *Phys. Rev. E* **104**, 044115 (2021).
- [60] S. K. Manikandan, S. Ghosh, A. Kundu, B. Das, V. Agrawal, D. Mitra, A. Banerjee, and S. Krishnamurthy, Quantitative analysis of non-equilibrium systems from short-time experimental data, *Commun. Phys.* **4**, 258 (2021).
- [61] S. Borsley, J. M. Gallagher, D. A. Leigh, and B. M. W. Roberts, Ratcheting synthesis, *Nat. Rev. Chem.* **8**, 8 (2024).
- [62] H. Wang and H. Qian, On detailed balance and reversibility of semi-Markov processes and single-molecule enzyme kinetics, *J. Math. Phys.* **48**, 013303 (2007).
- [63] C. Maes, K. Netočný, and B. Wynants, Dynamical fluctuations for semi-Markov processes, *J. Phys. A: Math. Theor.* **42**, 365002 (2009).
- [64] S. Rahav and C. Jarzynski, Fluctuation relations and coarse-graining, *J. Stat. Mech.: Theory Exp.* (2007) P09012.
- [65] D. Hartich and A. Godec, Emergent memory and kinetic hysteresis in strongly driven networks, *Phys. Rev. X* **11**, 041047 (2021).
- [66] T. Schürmann and P. Grassberger, Entropy estimation of symbol sequences, *Chaos* **6**, 414 (1996).
- [67] D. T. Gillespie, Exact stochastic simulation of coupled chemical reactions, *J. Phys. Chem.* **81**, 2340 (1977).
- [68] E. T. Jaynes, Information theory and statistical mechanics, *Phys. Rev.* **106**, 620 (1957).
- [69] N. Shiraishi, *An Introduction to Stochastic Thermodynamics: From Basic to Advanced*, Fundamental Theories of Physics, Vol. 212 (Springer Nature, Singapore, 2023).
- [70] E. Parzen, On estimation of a probability density function and mode, *Ann. Math. Stat.* **33**, 1065 (1962).
- [71] B. W. Silverman, *Density Estimation for Statistics and Data Analysis*, Vol. 26 (CRC Press, 1986).