



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Direct evaluation of dynamical large-deviation rate functions using a variational ansatz

Daniel Jacobson and Stephen Whitelam

Phys. Rev. E **100**, 052139 — Published 25 November 2019

DOI: [10.1103/PhysRevE.100.052139](https://doi.org/10.1103/PhysRevE.100.052139)

Direct evaluation of dynamical large-deviation rate functions using a variational ansatz

Daniel Jacobson^{1*} and Stephen Whitelam^{2†}

¹*Division of Chemistry and Chemical Engineering,*

California Institute of Technology, Pasadena, California 91125, USA

²*Molecular Foundry, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA*

We describe a simple form of importance sampling designed to bound and compute large-deviation rate functions for time-extensive dynamical observables in continuous-time Markov chains. We start with a model, defined by a set of rates, and a time-extensive dynamical observable. We construct a reference model, a variational ansatz for the behavior of the original model conditioned on atypical values of the observable. Direct simulation of the reference model provides an upper bound on the large-deviation rate function associated with the original model, an estimate of the tightness of the bound, and, if the ansatz is chosen well, the exact rate function. The exact rare behavior of the original model does not need to be known in advance. We use this method to calculate rate functions for currents and counting observables in a set of network- and lattice models taken from the literature. Straightforward ansätze yield bounds that are tighter than bounds obtained from Level 2.5 of large deviations via approximations that involve uniform scalings of rates. We show how to correct these bounds in order to recover the rate functions exactly. Our approach is complementary to more specialized methods, and offers a physically transparent framework for approximating and calculating the likelihood of dynamical large deviations.

I. INTRODUCTION

Dynamical systems, such as chemical networks [1], biochemical and molecular machines [2–4], and models of driven [5–8] and glassy [9–11] systems, exhibit fluctuations, departures from typical behavior [12]. Fluctuations of time-extensive observables – which can be work, entropy production [13, 14], other currents [8, 15, 16], or dynamical activity [11, 17] – characterize the behavior of these systems, much as fluctuations of size-extensive quantities, such as energy or magnetization, characterize the static behavior of equilibrium systems [18, 19]. The probability distributions that control dynamical fluctuations satisfy certain requirements, known as fluctuation relations [3, 13, 14, 20–26], which impose constraints on their symmetries. The precise form of these distributions, however, must be obtained by explicit calculation.

Here we focus on calculating probability distributions $\rho_T(A)$ for models with a discrete state space, for stochastic dynamical trajectories of elapsed time T and time-extensive observables A . Time-extensive observables are those that can be built from a sum of values of individual pieces of a trajectory. For large values of T these distributions often adopt the large-deviation form [20–25, 27, 28]

$$\rho_T(A) \approx e^{-TJ(a)}, \quad (1)$$

in which $a = A/T$ is the time-intensive value of the observable. $J(a)$ is the large-deviation *rate function*, which quantifies the likelihood of observing particular values of the observable a [29]. The symbol \approx denotes equality of the logarithms of both sides of (1), to leading order in T ,

for all values of a . In the physics literature, most numerical methods for calculating $J(a)$ aim to first compute its Legendre transform, the scaled cumulant-generating function (SCGF) [11, 28, 30–35]. It is possible to recover $J(a)$ from the SCGF if the former is convex [36]. A common way to calculate the SCGF is to use cloning methods [30, 37], which duplicate or eliminate trajectories according to their time-integrated weights. Often cloning is supplemented by other importance-sampling methods [32–34], some of which make use of a modified dynamics in order to produce trajectories that more closely resemble the rare dynamics of the original model.

Determining $J(a)$ *solely* by reweighting trajectories of a modified dynamics, without prior knowledge of the rare dynamics of the model of interest, is not widely done (see, however, Refs. [38–41]). Standard arguments suggest that determining the probability distribution of a within one dynamics by reweighting against a second dynamics requires, in general, the evaluation of random quantities whose variance is exponential in the trajectory length [42–45] (see Section II F). Such observations are sometimes taken to mean that trajectory reweighting, without advance knowledge of the rare dynamics to be sampled, is little better than direct sampling using the original model [46]. Here we argue that more optimism is warranted, and show that the conditions under which meaningful results can be extracted from trajectory reweighting are much less restrictive than has been recognized. Moreover, trajectory reweighting presents few technical complications beyond the requirement to simulate the original model with modified rates, and allows the reconstruction of $J(a)$ directly, without first calculating the SCGF.

To compute the large-deviation rate function $J(a)$ for a given model and dynamical observable a , we use a simple form of importance sampling [42–44, 47–50]. We begin with a modification of the model dynamics. This

* Daniel.Jacobson@caltech.edu

† swhitelam@lbl.gov

modified or *reference* model is a microscopic ansatz for the original model’s behavior, conditioned on particular values of the observable a . The ansatz is characterized by a set of parameters whose values we determine variationally. In practical terms we simply guess a reference dynamics that is able to generate more or less of a than the normal dynamics. Let the typical value of a produced by original and reference models be a_0 and \tilde{a}_0 , respectively [51]. Reweighting trajectories of the reference model produces an upper bound on the rate function J associated with the original model at the point \tilde{a}_0 , an estimate of the tightness of the bound, and, if the ansatz is chosen well, the exact rate function (to within statistical error). That is, the reference dynamics is a true ansatz, a guess whose accuracy can be determined by subsequent calculation. Repeating the calculation for a set of reference models possessing a set of distinct values $\{\tilde{a}_0\}$ allows us to attempt reconstruction of $J(a)$ at the set of points $a \in \{\tilde{a}_0\}$. In this respect the procedure is similar to umbrella sampling of equilibrium systems. We show that the conditions under which the exact rate function can be recovered are less restrictive than usually assumed.

Any reference dynamics can be reweighted to produce *some* upper bound on $J(a)$, simply by making the desired value of a typical [52]. Good choices of reference dynamics, leading to tight bounds, render the fluctuations of the reweighting factor or likelihood ratio (the ratio of path probability of new and old dynamics) small. We show here that relatively simple reference-model choices produce meaningful (i.e. tight) bounds, for a set of models and observables taken from the literature. We compare the bounds produced by our method with universal bounds on currents [53, 54] and non-decreasing counting variables [55]. Those bounds can be obtained from Level 2.5 of large deviations [56, 57], the exact rate function for the empirical flow (jumps between states) and measure (state occupation times), via a uniform rescaling of rates. That approach provides important physical insight into the quantities that constrain fluctuations of time-integrated observables, and also provides numerical bounds on rate functions. Our approach, which uses a microscopic ansatz within the exact path integral for the dynamics, produces tighter bounds, particularly far into the tails of rate functions. The extent to which bounds vary as we change the nature of the ansatz provides physical insight into how much certain types of microscopic processes contribute to the rare dynamics of a model. Microscopic ansätze, even relatively simple ones, are capable of capturing a wide range of behavior, including regimes of anomalous fluctuations in which the usual central-limit theorem breaks down [39]. Computing a correction to these bounds, by measuring fluctuations of the likelihood ratio, allows the recovery of the exact rate function. Importantly, fluctuations of the likelihood ratio do not need to be zero for $J(a)$ to be calculated

The approach described here is variational, in the sense that we vary the parameters of the reference model in order to identify the dynamics that best approximates

the rare dynamics of interest. Variational principles underpin the study of large deviations, embodied by the notion that “any large deviation is done in the least unlikely of all the unlikely ways” [27]. Variational ideas are central to different representations of rare processes – see e.g. Section 5 of Ref. [31] – and have been widely used in analytic and numerical work [33, 35, 58–60]. The aim of this paper is to present a simple, physically-motivated approach to bounding and calculating rate functions using a variational principle enacted by (only) direct simulations, and to present a set of convergence criteria, adapted from Ref. [61], that reveal when bounds can be corrected to produce the exact rate function. We have provided a GitHub script [62] that computes the correction term automatically, this being the most involved step of the calculation. These results extend our previous work [39–41] by a) showing how different forms of physically-motivated reference dynamics can be used to treat different models, and b) by providing a set of criteria that identify when the exact rate function can be recovered. One point we emphasize is that considerable progress can be made using physical intuition and basic knowledge of the properties of a model, without the application of other forms of importance sampling (such as cloning or transition-path sampling). Our method requires only continuous-time Monte Carlo simulation, and so can be applied to any set of circumstances in which that method can be used, including to models with unbounded state spaces [39]. In addition, it can be used to reconstruct families of large-deviation rate functions from a single set of simulations, using the principle that the dynamics of one model can be reweighted to examine the dynamics of many others [41]. Reference models represent a form of importance sampling similar in spirit but different in detail to the umbrella potentials used in equilibrium sampling [63, 64].

In Section II we describe our approach, which we refer to as VARD (for Variational Ansatz for Rare Dynamics). In general terms there are many forms of VARD that have been used in the literature (see above); we use the term to convey the specific notion of doing (only) direct simulations of a family of modified models. In Section III we apply the method to four models taken from the literature. We have chosen models from the literature that display a variety of interesting behavior: two lattice models (the asymmetric simple exclusion process [6, 65] and the Fredrickson-Andersen model [66]) and two network models, and we sample both currents and non-decreasing counting variables to show that the method works the same way for each. In the cases described in Section III B–Section III E, relatively simple choices of reference model allow the computation of the exact rate function, and we gain physical insight into the nature of the dynamics that contributes to particular pieces of $J(a)$. We also show, in Section III F, that bounds that are descriptive in small systems remain so in larger systems. We conclude in Section IV.

II. A VARIATIONAL ANSATZ FOR RARE DYNAMICS

A. Continuous-time Markov chains and large deviations

Consider a continuous-time dynamics [67] on a set of discrete states, defined by the master equation

$$\partial_t P_x(t) = \sum_{y \neq x} W_{yx} P_y(t) - R_x P_x(t). \quad (2)$$

Here $P_x(t)$ is the probability that a system resides in (micro)state x at time t . W_{xy} is the rate for passing from state x to state y , and $R_x = \sum_{y \neq x} W_{xy}$ is the escape rate from x (Table 1 provides a reference for some of the more frequently-used symbols in this paper). The standard algorithm for simulating the dynamics (2) is as follows [68]. From state x , choose a destination state y with probability

$$p_{xy} = \frac{W_{xy}}{R_x}. \quad (3)$$

The time increment Δt required to make this move is a random number drawn from the exponential distribution with mean $1/R_x$,

$$p_x(\Delta t) = R_x e^{-R_x \Delta t}. \quad (4)$$

Given an initial state x_0 , the dynamics defined by (3) and (4) generates a trajectory $\omega = x_0 \rightarrow x_1 \rightarrow \dots \rightarrow x_{K(\omega)}$, which consists of a sequence of $K(\omega)$ jumps $x_k \rightarrow x_{k+1}$ and associated jump times Δt_k . In this paper we are concerned with calculating

$$\rho_T(A) = \sum_{\omega} p(\omega) \delta(T(\omega) - T) \delta(A(\omega) - A), \quad (5)$$

the probability distribution, taken over all trajectories of elapsed time T , of a time-extensive dynamical observable

$$A(\omega) = \sum_{k=0}^{K(\omega)-1} \alpha_{x_k x_{k+1}}. \quad (6)$$

Here α_{xy} is the change of the observable A upon moving from x to y , and $A(\omega)$ is the sum of these quantities over a single trajectory ω . We define $a(\omega) \equiv A(\omega)/T(\omega)$ as the time-intensive version of A . $T(\omega)$ is the elapsed time of trajectory ω , and is equal to T when $T_{K(\omega)} \leq T < T_{K(\omega)+1}$, where $T_{K(\omega)} = \sum_{k=0}^{K(\omega)-1} \Delta t_k$. The symbol $p(\omega)$ is the probability of a trajectory ω . Given an initial state, this term is equal to a product of factors (3) and (4) for all jumps of the trajectory, multiplied by the probability of not exiting state $x_{K(\omega)}$ between times $T_{K(\omega)}$ and T .

In (5), the delta functions denote the conditions of fixed A and fixed T that we wish to impose on the trajectory ensemble. This conditioning defines the *microcanonical path ensemble* [31], of which (5) is the normalization constant.

TABLE I. Glossary of frequently-used symbols

$A = aT$	path-extensive order parameter
T	elapsed time of a trajectory
K	number of jumps of a trajectory
α_{xy}	change of A upon making the jump $x \rightarrow y$
W_{xy}	“original” model rates for the jump $x \rightarrow y$
R_x	“original” model escape rate $\sum_{y \neq x} W_{xy}$
a_0	typical value of a under the original dynamics
$J(a)$	rate function for a under the original dynamics
$J_0(a)$	upper bound on $J(a)$
$J_1(a)$	correction to the bound: $J(a) = J_0(a) + J_1(a)$
\tilde{W}_{xy}	reference-model rates for jumps $x \rightarrow y$
s	one parameter of \tilde{W}_{xy} : see (8)
β_{xy}	remaining parameters of \tilde{W}_{xy} : see (8)
\tilde{R}_x	reference-model escape rate $\sum_{y \neq x} \tilde{W}_{xy}$
λ	clock bias: see (9)
\tilde{a}_0	typical value of a under the reference dynamics
$J_0[\{x\}]$	$J_0(a)$ determined from a scan of the reference-model parameter set $\{x\}$

We focus on models for which, for large values of T , the probability distribution (5) adopts the large-deviation form (1). Our aim is to calculate the rate function $J(a)$ (sometimes the notation $I(a)$ or $\varphi(a)$ is used to denote a rate function [28, 55]). This function quantifies the rate of decay of atypical values of a . For many models $J(a)$ has a unique minimum at a point $a = a_0$, where $J(a_0) = 0$. This point defines the typical value of a : the distribution $\rho_T(A)$ concentrates on a_0 in the long-time limit, an expression of the law of large numbers [27, 28]. In general, a rate function can have more than one point at which it is zero, defining multiple typical values of an observable [28, 39, 69]. Often, the rate function is quadratic in the neighborhood of its minimum, an expression of the central-limit theorem [27, 28]. Exceptions to this norm occur at phase transitions, where the usual central-limit theorem does not hold [70]. Far from their minima, rate functions display a variety of behaviors [28]. However, direct simulation of the dynamics (3) and (4) leads to poor sampling of $J(a)$ anywhere other than in the neighborhood $a \approx a_0$.

B. Quantifying rare events

To remedy the sampling problem for values of the observable a far from a_0 , we introduce a *reference model*. We wish to reweight the trajectories of the reference model in order to approximate or calculate $J(a)$ for values of a potentially far from a_0 . The reference model must satisfy certain requirements. It needs to be able to generate all trajectories possible in the original model, but no trajectories not possible in the original model (otherwise the reweighting factor, discussed below, can be infinity or zero). We want the reference model to be able to generate trajectories possessing values of a that

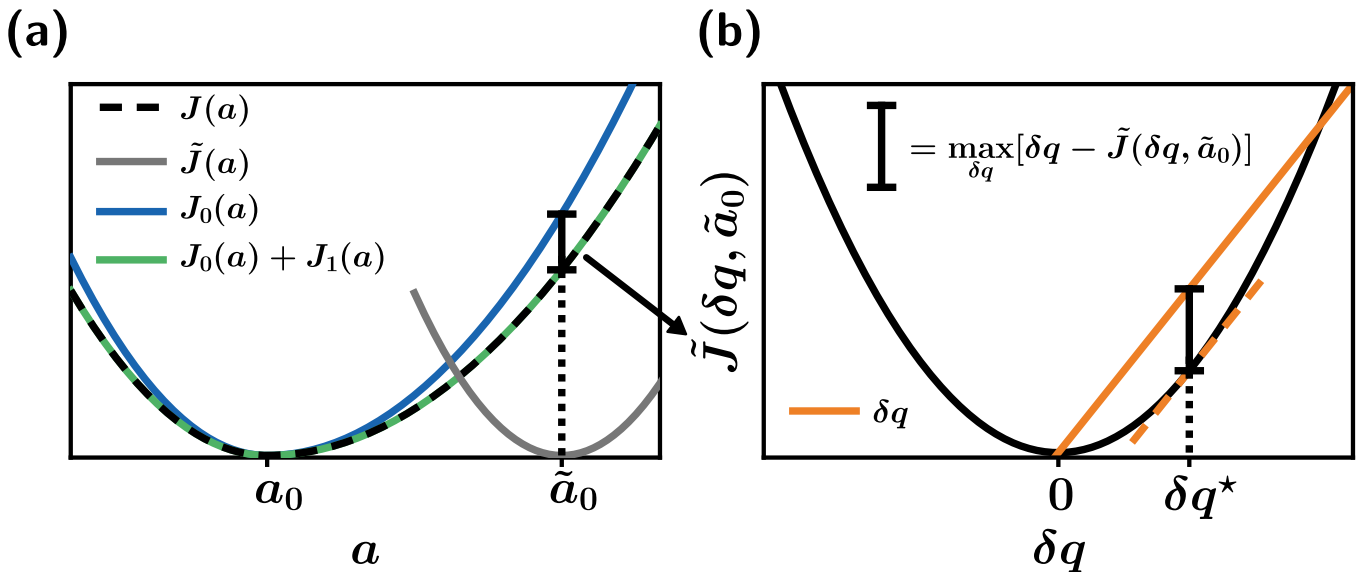


FIG. 1. Large deviations from a variational ansatz for rare dynamics (VARD). (a) We aim to bound or calculate the large-deviation rate function $J(a)$ (black dashed line) for a given model and observable a , under the continuous-time dynamics (3) and (4). We introduce a reference model (7) and (9), a variational ansatz for the rare dynamics of the original model conditioned upon a . The reference model has rate function $\tilde{J}(a)$ (gray line; we do not aim to compute this function). The typical value of a generated by the reference model is \tilde{a}_0 ; this is potentially far from the typical value a_0 generated by the original model. Evaluation of (17) from the sample mean of a single trajectory of the reference model produces one point $J_0(\tilde{a}_0)$ on the blue line, an upper bound on $J(a)$. (b) If we can evaluate the auxiliary rate function $\tilde{J}(\delta q, \tilde{a}_0)$ at the point δq^* at which its gradient is unity, then we can calculate the correction term (18) and determine one point on the green curve in panel (a), the exact rate function of the original model. If not, then we gain information about the quality of the bound $J_0(\tilde{a}_0)$. Variation of the parameters of the reference model allows reconstruction of the entire blue and (potentially) green curves in panel (a). VARD thus reduces a single nonlocal problem, the computation of $J(a)$ arbitrarily far from a_0 , to a series of local problems, each requiring the evaluation of an auxiliary rate function $\tilde{J}(\delta q, \tilde{a}_0)$ at the point δq^* at which its gradient is unity. As we show in this paper, this procedure can be carried out for models commonly found in the literature using relatively simple choices of reference model.

are rarely generated by the original model, which is relatively easy to arrange, but we also need to be able to recover from reference-model trajectories the probability with which such trajectories *would* have been generated by the original model. This second requirement is harder to arrange, but not prohibitively so. As we show, if the reference model generates trajectories possessing values of a in a manner completely unlike the original model, then we have to do prohibitively heavy sampling of reference-model trajectories in order to calculate $J(a)$. If, however, the reference model generates trajectories possessing values of a in a manner similar to the original model, then $J(a)$ can be reconstructed from trajectories of the reference model with relatively little effort. Importantly, the method tells us when this is so: we do not need to know in advance the precise nature of the rare dynamics of the original model in order to recover $J(a)$.

For a trajectory ω of the reference model we want to be able to influence how much of the dynamical observable is produced per jump, $A(\omega)/K(\omega)$, and the number of jumps per unit time, $K(\omega)/T$. To control the former we use a reference-model dynamics that selects destination

states with probability

$$\tilde{p}_{xy} = \frac{\tilde{W}_{xy}}{\tilde{R}_x}, \quad (7)$$

where \tilde{W}_{xy} is an effective rate, and $\tilde{R}_x = \sum_{y \neq x} \tilde{W}_{xy}$. (The true rates of the reference model are, from (7) and (9), $\tilde{W}_{xy}(R_x + \lambda)/\tilde{R}_x$.) Here we use the parameterization

$$\tilde{W}_{xy} = e^{-s\alpha_{xy} - \beta_{xy}} W_{xy}, \quad (8)$$

which is a modification of (3). The factor $e^{-s\alpha_{xy}}$ is chosen in order to guide the jump destination according to the change of the observable α_{xy} weighted by a parameter s . In general such a bias is not sufficient to control $A(\omega)/K(\omega)$, and so we also consider an additional arbitrary bias, β_{xy} . For the models considered here a simple and physically-motivated guess for what β_{xy} should be is sufficient to produce a good reference model. We shall return to this point.

To control the number of jumps per unit time, $K(\omega)/T$, we draw times between jumps of the reference model from the distribution

$$\tilde{p}_x(\Delta t) = (R_x + \lambda)e^{-(R_x + \lambda)\Delta t}, \quad (9)$$

where $\lambda > -\min_x R_x$ serves to make the jump time from a given state unusually large or small by the reckoning of the original model. This ‘‘clock trick’’ provides a simple way of sampling jump times without having those times appear explicitly in the reweighting factor [41]. (The parameter β_{xy} can also affect jump times indirectly, if, for example, it is chosen to be proportional to R_y , the escape rate from the destination state.)

Next observe that the path weight $p(\omega)$ in (5) can be written $\tilde{p}(\omega)\phi(\omega)$, where $\phi(\omega) = p(\omega)/\tilde{p}(\omega)$ is the reweighting factor, the ratio of weights of a trajectory ω in the original and reference models. $\phi(\omega)$ is also known as the likelihood ratio or the Radon-Nikodym derivative [31, 50]. For a jump $x \rightarrow y$ in time Δt , the reweighting factor is the product of (3) and (4), divided by the product of (7) and (9); for the entire trajectory ω we have

$$\phi(\omega) = e^{sA(\omega) + \lambda T(\omega) + Tq(\omega)}, \quad (10)$$

where

$$q(\omega) = \frac{1}{T} \sum_{k=0}^{K(\omega)-1} \left(\beta_{x_k x_{k+1}} + \ln \frac{\tilde{R}_{x_k}}{R_{x_k} + \lambda} \right) \quad (11)$$

is the piece of ϕ that is not fixed by the delta-function constraints in (5). (The time-dependent piece of (10) produced by $K(\omega)$ jumps is $e^{\lambda T K(\omega)}$; the contribution from the final entry in the path weight, the probability of not jumping between time $T_{K(\omega)}$ and $T(\omega)$, leads to the factor of $e^{\lambda T(\omega)}$ shown in (10).)

We can then write (5) in the form

$$\frac{\rho_T(A)}{\tilde{\rho}_T(A)} = e^{sA + \lambda T} \frac{\sum_{\omega} \tilde{p}(\omega) e^{Tq(\omega)} \delta(T) \delta(A)}{\sum_{\omega} \tilde{p}(\omega) \delta(T) \delta(A)}, \quad (12)$$

where $\tilde{\rho}_T(A) \approx e^{-T\tilde{J}(A)}$ is the analog of (5) for the reference model, and we have used the shorthand $\delta(X) \equiv \delta(X(\omega) - X)$. Replacing the sums over trajectories with an integral over trajectory weights gives

$$\frac{\rho_T(A)}{\tilde{\rho}_T(A)} = e^{sA + \lambda T} \int dq \tilde{p}_T(q|a) e^{Tq}, \quad (13)$$

where $\tilde{p}_T(q|a)$ is the normalized probability distribution of $q(\omega)$ for trajectories of the reference model that have specified values of a and T . For large T we assume that this distribution will obey a large-deviation principle of its own. If so we can write, using the rules of conditional probability,

$$\tilde{p}_T(q|a) \approx e^{-T\tilde{J}(q|a)} = e^{-T[\tilde{J}(q,a) - \tilde{J}(a)]}, \quad (14)$$

where $\tilde{J}(q, a)$ is the joint rate function for q and a within the reference model.

We next take the large- T limit in (13), replace all probability distributions with their large-deviation forms, and set $a = \tilde{a}_0$, the value typical of the reference model (such that $\tilde{J}(\tilde{a}_0) = 0$). The result, upon taking logarithms, is

$$J(\tilde{a}_0) = -s\tilde{a}_0 - \lambda - \lim_{T \rightarrow \infty} T^{-1} \ln \int dq e^{T[q - \tilde{J}(q, \tilde{a}_0)]}. \quad (15)$$

Finally, we introduce $\delta q \equiv q - \tilde{q}_0$, where \tilde{q}_0 is the value of q typical of the reference model. This value can be computed from a single reference-model trajectory (for a given set of parameters s, λ, β_{xy}). Extracting $e^{T\tilde{q}_0}$ from the exponential in (15) and evaluating the integral using the saddle-point method yields

$$J(\tilde{a}_0) = J_0(\tilde{a}_0) + J_1(\tilde{a}_0), \quad (16)$$

where

$$J_0(\tilde{a}_0) = -s\tilde{a}_0 - \lambda - \tilde{q}_0 \quad (17)$$

and

$$J_1(\tilde{a}_0) = -\max_{\delta q} [\delta q - \tilde{J}(\delta q, \tilde{a}_0)]. \quad (18)$$

Eqs.(16)–(18) provide an exact representation of the rate function, if the probability distributions (5) and (14) adopt large-deviation forms [71]. Fig. 1 illustrates the relationship between Equations (16), (17), and (18), which are central to the sampling scheme discussed in this paper.

C. We can compute $J(a)$ as the sum of a bound and a correction

The piece $J_0(\tilde{a}_0) \geq J(\tilde{a}_0)$, Eq. (17), is an upper bound on the rate function at the point $a = \tilde{a}_0$, by Jensen’s inequality applied to (13), and can be obtained from the sample mean of single trajectory of the reference model. It is always possible to calculate this term. If $\tilde{J}(a)$ is locally quadratic about \tilde{a}_0 , meaning that the usual central-limit theorem holds [72], then errors in the computation of \tilde{a}_0 go as $\sqrt{\langle (a - \tilde{a}_0)^2 \rangle} \sim T^{-1/2}$. The same is true for the computation of \tilde{q}_0 . Thus statistical errors in the computation of the bound can be made negligible simply by computing (17) for a sufficiently long trajectory.

The term $J_1(\tilde{a}_0)$, Eq. (18), is a correction to the bound, and can be calculated by sampling values of q , Eq. (11), of trajectories of the reference model that have $a = \tilde{a}_0$. It is possible to calculate this term if the reference model is chosen well, but not if it is chosen badly.

The two terms in (18) describe a competition between the logarithmic weight δq associated with reference-model trajectories that have atypical values of q , and the logarithmic probability $\tilde{J}(\delta q, \tilde{a}_0)$ of observing such trajectories. When $\tilde{J}(\delta q, \tilde{a}_0)$ is differentiable, (18) can be written

$$J_1(\tilde{a}_0) = -\delta q^* + \tilde{J}(\delta q^*, \tilde{a}_0), \quad (19)$$

where

$$\partial_{\delta q} \tilde{J}(\delta q, \tilde{a}_0)|_{\delta q = \delta q^*} = 1. \quad (20)$$

Thus we need to measure the value of $\tilde{J}(\delta q, \tilde{a}_0)$ at the point δq^* at which its gradient is unity, which will be a unique point when $\tilde{J}(\delta q, \tilde{a}_0)$ is convex. The sampling

problem is now *localized*: instead of sampling $J(a)$ arbitrarily far from a_0 (using the original model), we need only sample a specific piece δq^* of an auxiliary rate function, $\tilde{J}(\delta q, \tilde{a}_0)$ (using the reference model). This fact, summarized in Fig. 1, shows why the present scheme has the potential to be much more efficient than unbiased simulation, if the reference model is chosen well.

This sampling problem is still formidable in general. If the reference model is chosen badly, meaning that its typical trajectories have very different character to trajectories of the original model that have $a = \tilde{a}_0$, then the bound will be slack, meaning that $J_0(\tilde{a}_0) \gg J(\tilde{a}_0)$, and so $J_1(\tilde{a}_0)$ will be large. In this case $\tilde{J}(\delta q, \tilde{a}_0)$ will be broad around its minimum $\delta q = 0$ (the variance of δq will be large) and unreasonably heavy sampling using the reference model will be required to determine the point δq^* (because this corresponds to a rare event within the reference model).

However, for good choices of the reference model the opposite situation arises: the bound will be tight, meaning that $J_0(\tilde{a}_0) \approx J(\tilde{a}_0)$, and so $J_1(\tilde{a}_0)$ will be small. In this case the latter can be evaluated with reasonable numerical effort (in the examples that follow we can gather the required statistics of q by sub-sampling a single trajectory of the reference model.) If we can reconstruct $\tilde{J}(\delta q^*, \tilde{a}_0)$ then we can calculate $J_1(\tilde{a}_0)$ and we have obtained the exact rate function.

D. Computing the correction

Given a model and an observable a , we construct a reference model (7) and (9) so as to approximate or calculate $J(a)$. In Section III we provide a set of worked examples of this procedure. In general terms we simply guess which rates of the original model can be modified so as to produce more or less of a than usual, and introduce a parameter (s , λ , or β_{xy}) able to control the rate in question. We do not know in advance which combination of these modified rates best approximates the way in which the original model produces rare values of a , but by running short trajectories of the reference model for different values of its parameters we can identify how this is done within the space of possibilities defined by the reference model. From the sample mean of each reference-model trajectory we obtain the values \tilde{a}_0 and \tilde{q}_0 ; plotting \tilde{a}_0 against $-s\tilde{a}_0 - \lambda - \tilde{q}_0$ for a range of values of reference-model parameters, and identifying the lower envelope of these points (conveniently calculated using a union of convex hull constructions), gives the bound $J_0(a)$ associated with that choice of reference model.

This bound is the starting point for our attempt to calculate the correction $J_1(a)$. The correction term can be interpreted as a measure of how close the typical dynamics of the reference model is to the desired rare dynamics of the original model. If $J_1(a) = 0$ then typical trajectories of the reference model correspond exactly to trajectories of the original model conditioned on the rel-

evant value a of the order parameter. If J_1 is small then (slightly) atypical versions of reference-model trajectories correspond to the desired rare dynamics; and if J_1 is large (or cannot be calculated) then it is the very rare trajectories of the reference model that correspond to the desired rare dynamics of the original model.

In previous versions of our sampling method [39–41] we used a cumulant expansion to evaluate the integral in (15), giving, in place of (18),

$$J_1^{\text{approx}}(\tilde{a}_0) = \frac{T}{2}\sigma_{\text{ref}}^2 + \frac{T^2}{6}\kappa_{\text{ref}} + \dots \quad (21)$$

Here $\sigma_{\text{ref}}^2 \propto 1/T$ is the variance of δq over typical trajectories of the reference model (those having $a = \tilde{a}_0$), i.e. $\sigma_{\text{ref}}^2 = \langle (\delta q)^2 \rangle_{\text{ref}}$, and $\kappa_{\text{ref}} = \langle (\delta q)^3 \rangle_{\text{ref}} \propto 1/T^2$. Eq. (21) can give accurate results for the rate-function correction [39–41], but does not provide a self-consistent way of determining *when* the correction is accurate. At best we can determine that the first omitted term in the expansion (21) is small, but this does not provide a proof of convergence.

In this paper we present an alternative way to calculate the correction term (18), which builds upon methods designed to compute rate functions (or their SCGF Legendre duals) empirically [61, 73]. This process is more involved than the computations required to evaluate (21), but has the advantage of providing a set of clear convergence criteria and statistical error bars. This information reveals when we have converged (18), and so have the exact rate function, and when we do not, thus turning the reference-model guess into a true ansatz. In the remainder of this section we describe the method we use to compute (18). We have provided a GitHub script [62] for calculating the correction automatically.

To obtain the correction we first assume that $\tilde{J}(\delta q, a)$ is differentiable, and so work with (19) instead of (18). We then introduce the two-dimensional scaled cumulant-generating function (SCGF),

$$\tilde{\theta}(k_{\delta q}, k_a) \equiv \lim_{T \rightarrow \infty} T^{-1} \ln \langle e^{T(k_{\delta q} \delta q + k_a a)} \rangle_{\text{ref}}, \quad (22)$$

where the angle brackets denote a trajectory ensemble average within the reference model. The SCGF (22) is related to $\tilde{J}(\delta q, a)$ through the double Legendre transform

$$\tilde{J}(\delta q, a) = k_{\delta q} \delta q + k_a a - \tilde{\theta}(k_{\delta q}, k_a) \quad (23)$$

where $k_{\delta q}$ and k_a are conjugate to δq and a respectively. As a result, if we want to calculate $\tilde{J}(\delta q, a)$ at a single point, we can calculate the right-hand side of (23). Doing so is much more efficient than attempting to reconstruct $\tilde{J}(\delta q, a)$ directly, for the reasons given in Section IIF.

The formula for the correction (19) depends on $\tilde{J}(\delta q^*, \tilde{a}_0)$, which we can get from (23):

$$\tilde{J}(\delta q^*, \tilde{a}_0) = k_{\delta q^*} \delta q^* + k_{\tilde{a}_0} \tilde{a}_0 - \tilde{\theta}(k_{\delta q^*}, k_{\tilde{a}_0}). \quad (24)$$

We can simplify this relationship by combining the implicit definition of $k_{\delta q}$ in the Legendre transform with (20) to get

$$k_{\delta q^*} = \partial_{\delta q} \tilde{J}(\delta q, \tilde{a}_0)|_{\delta q = \delta q^*} = 1. \quad (25)$$

Inserting (25) into (24) yields

$$\tilde{J}(\delta q^*, \tilde{a}_0) = \delta q^* + k_{\tilde{a}_0} \tilde{a}_0 - \tilde{\theta}(1, k_{\tilde{a}_0}). \quad (26)$$

The quantity \tilde{a}_0 is the typical value of the observable in the reference model, and can be obtained from a single reference-model trajectory. The three other unknown quantities on the right-hand side of (26) that are needed for the correction are δq^* , $k_{\tilde{a}_0}$, and $\tilde{\theta}(1, k_{\tilde{a}_0})$.

To calculate these quantities we have to compute the value of the two-dimensional SCGF (22) at various points $(k_{\delta q}, k_a)$. We can do this using a simple extension of existing techniques developed to sample points on 1D SCGFs [61, 73]. Following Ref. [61] we generate a single long trajectory of the reference model and sub-sample it into M approximately independent blocks ω_i of length $T(\omega_i) = B$. Within each block we record the sample mean of δq and a , which we write as δq_i and a_i . $\tilde{\theta}(k_{\delta q}, k_a)$ can be calculated from this data set using the estimator

$$\hat{\theta}(k_{\delta q}, k_a) = \frac{1}{B} \ln \left(\frac{1}{M} \sum_{i=1}^M e^{B(k_{\delta q} \delta q_i + k_a a_i)} \right). \quad (27)$$

Eq. (27) is guaranteed to converge to the exact value of the SCGF, $\tilde{\theta}(k_{\delta q}, k_a)$, in the limit $M \rightarrow \infty$ and $B \rightarrow \infty$. The convergence properties of this estimator for finite M and B will be addressed in the next section, II E. For now we assume that we can obtain convergence as needed. Finally, note that by changing the values of $k_{\delta q}$ and k_a in (27) a single data set consisting of values of a_i and δq_i , generated from a single long trajectory, can be used to recover many points $(k_{\delta q}, k_a)$ on $\tilde{\theta}(k_{\delta q}, k_a)$.

We now turn to the calculation of the three unknowns in (26), δq^* , $k_{\tilde{a}_0}$ and $\tilde{\theta}(1, k_{\tilde{a}_0})$. First we use the relation

$$\tilde{a}_0 = \partial_{k_a} \tilde{\theta}(k_{\delta q^*} = 1, k_a)|_{k_a = k_{\tilde{a}_0}} \quad (28)$$

to find $k_{\tilde{a}_0}$. We do so by calculating the SCGF, $\tilde{\theta}(k_{\delta q}, k_a)$, along a 1D slice through its 2D domain using the estimator (27). This slice is defined by fixing $k_{\delta q} = k_{\delta q^*} = 1$ and varying k_a . We then use the method of finite differences to get $\partial_{k_a} \tilde{\theta}(k_{\delta q^*} = 1, k_a)$ at each point k_a , and find the point that fulfills (28). Once we know the value of $k_{\tilde{a}_0}$ we can calculate $\tilde{\theta}(1, k_{\tilde{a}_0})$, again using (27). Finally we can compute δq^* using the analog of (28) for δq ,

$$\delta q^* = \partial_{k_{\delta q}} \tilde{\theta}(k_{\delta q}, k_{\tilde{a}_0})|_{k_{\delta q} = 1}. \quad (29)$$

Inserting δq^* , $k_{\tilde{a}_0}$, $\tilde{\theta}(1, k_{\tilde{a}_0})$ and \tilde{a}_0 into (26) yields $\tilde{J}(\delta q^*, \tilde{a}_0)$. The correction to the bound, $J_1(\tilde{a}_0)$, then follows from (19).

E. Convergence of the SCGF Estimator

When used with only a finite number M of blocks of finite length B , the estimator $\hat{\theta}(k_{\delta q}, k_a)$, defined in (27), can exhibit statistical and systematic errors. In this section, we analyze these errors to understand when the estimate of $J_1(\tilde{a}_0)$ calculated through the SCGF (22) will be accurate. As in the previous subsection, this analysis closely follows Ref. [61].

To quantify the statistical error associated with our calculated value of $J_1(\tilde{a}_0)$ we repeat the calculation procedure using R independent trajectories. Each of these trajectories is split up into M blocks of length B , and used to calculate $J_1(\tilde{a}_0)$. Our final estimate for the correction is then

$$\hat{J}_1(\tilde{a}_0) = \frac{1}{R} \sum_{j=1}^R J_1^{(j)}(\tilde{a}_0), \quad (30)$$

where $J_1^{(j)}(\tilde{a}_0)$ is the value of the correction calculated from the j^{th} trajectory. The statistical error of (30) can be estimated using

$$\text{Err}[\hat{J}_1(\tilde{a}_0)] = \sqrt{\frac{\text{Var}[\hat{J}_1(\tilde{a}_0)]}{R}}. \quad (31)$$

This statistical error is only meaningful if we know that the systematic error in the calculation is comparatively small. There are two sources of systematic error that arise when using (27): correlation error and linearization error. Correlation error results from the fact that the derivation of the estimator assumes that the trajectory blocks are long enough to be approximately independent. This will be true if $B > T_{\text{corr}}$ where T_{corr} is the correlation time of the reference model. If, however, the sub-sampled blocks of a trajectory are correlated, meaning that $B < T_{\text{corr}}$, then we will not obtain an accurate estimate of $\tilde{\theta}(k_{\delta q}, k_a)$ even as $M, R \rightarrow \infty$.

One way to resolve this correlation issue is to increase the block time B , but this also increases the magnitude of the other systematic source of error, linearization error. Linearization error is a manifestation of the fact that trajectories that contribute most to the average in the SCGF (22) for $k_a, k_{\delta q} \neq 0$ have atypical values of δq and a . Using larger B creates more self-averaging within a single trajectory and, as a result, makes sampling these atypical values more difficult. Linearization error also increases for fixed B with increasing $|k_a|$ and $|k_{\delta q}|$, because larger values of these parameters weight rare trajectories' contributions to the SCGF more heavily. Ref. [61] contains a detailed discussion of these problems. The authors of that work describe a method for checking to see if linearization error will substantially influence the estimate of the SCGF at a point $(k_{\delta q}, k_a)$ for some fixed B . We use this check, modified to account for the fact that our SCGF is two dimensional, as follows.

First, we calculate the SCGF (27) along another 1D slice through its 2D domain. This slice is defined by fixing

$k_{\delta q} = 0$ and increasing k_a , starting from $k_a = 0$ [74]. By using finite difference along this slice we can calculate how a varies with k_a . To compute the statistical error, we generate R independent trajectories (usually the same R trajectories we used to get (31)), split each one into M blocks of length B , and use each data set to calculate $a(k_a)$ along the slice. Our final estimate for the value of a at each k_a is

$$\hat{a}(k_a) = \frac{1}{R} \sum_{j=1}^R a^{(j)}(k_a), \quad (32)$$

where $a^{(j)}(k_a)$ is calculated from the j^{th} trajectory.

We will not end up using the $\hat{a}(k_a)$ values themselves. Instead, we focus on the associated values of the statistical error, calculated in the same way as the error of (27),

$$\text{Err}[\hat{a}(k_a)] = \sqrt{\frac{\text{Var}[\hat{a}(k_a)]}{R}}. \quad (33)$$

A plot of $\text{Err}[\hat{a}(k_a)]$ as a function of k_a will peak at some point $\hat{k}'_a(B)$ and then decline. Again $k_{\delta q} = 0$ is fixed during this entire calculation. As discussed in Ref. [61], $\hat{k}'_a(B)$ is an estimate for the maximum value $k'_a(B)$ at which the calculation of the SCGF will converge without being overwhelmed by linearization error. $k'_a(B)$ is a decreasing function of B , because linearization error grows as B is increased.

Next we note that $\tilde{\theta}(k_{\delta q} = 0, k'_a)$ corresponds to a point on the rate function $\tilde{J}(a')$, via the Legendre transform (23). If we can converge the value of $\tilde{J}(a')$ then we can, with the same data set, also converge the value of $\tilde{J}(\delta q, a)$ at any point for which

$$\tilde{J}(\delta q, a) < \tilde{J}(a'). \quad (34)$$

This statement is intuitive in the context of probabilities and rate functions. However, it also applies when working with the SCGF, provided that $\tilde{J}(\delta q, a)$ is convex [36]. Thus the value of a point $\hat{\theta}(k_{\delta q}, k_a)$ estimated using (27) will be unaffected by linearization error if the associated point on the rate function, $\hat{\tilde{J}}(\delta q, a)$, satisfies

$$\hat{\tilde{J}}(\delta q, a) < \xi \hat{\tilde{J}}(a'), \quad (35)$$

where $\xi < 1$ is an empirical constant (we set $\xi = 0.8$). The terms $\hat{\tilde{J}}$ in (35) are averages over R independent data sets of the corresponding Legendre transform (23). This formula allows the convergence criteria derived in [61] for estimating one-dimensional SCGFs and rate functions to be applied in an arbitrary number of dimensions.

We now discuss the procedure we use to converge the correction (19) while accounting for correlation and linearization errors. For fixed block time B we first increase M and R until the error bars for $\hat{J}_1(\tilde{a}_0)$, (33),

are smaller than a desired value. We repeat this process for larger and larger B until $\hat{J}_1(\tilde{a}_0)$ becomes independent of B . This is equivalent to increasing B until it becomes larger than the reference-model correlation time T_{corr} . If this happens while the convergence criterion (35) holds at $\tilde{J}(\delta q^*, \tilde{a}_0)$ then the calculation has succeeded, and we have computed (to within statistical error) the exact value of the correction $J_1(\tilde{a}_0)$.

If, however, the convergence criterion (35) fails to hold in the regime in which (27) still changes rapidly with B , then the bound $J_0(\tilde{a}_0)$ is too far from the exact answer for us to effectively sample $J_1(\tilde{a}_0)$ using direct simulation of the reference model. In this case the chosen ansatz is probably missing a crucial piece of the physics of the rare trajectories of the original model. On several occasions our failure to converge $J_1(\tilde{a}_0)$ based on an initial guess led us to construct a modified ansatz from which we could converge the exact correction. In the cases described in this paper we were able to reconstruct $J(a)$ using physically transparent ansätze containing only a modest number of parameters.

A quick way to estimate the scale of the correction is to compute the first term in (21), which requires computing only the variance of δq within the reference model. By the central limit theorem we will have $\tilde{J}(\delta q, \tilde{a}_0) \approx (\delta q)^2 / (2T\sigma_{\text{ref}}^2)$ close enough to the origin, where $\sigma_{\text{ref}}^2 \propto 1/T$ is the variance of q within the reference model. If σ_{ref}^2 is small (which is the case when the ansatz is very good) then the correction $J_1(\tilde{a}_0) \approx T\sigma_{\text{ref}}^2/2$. Thus if $T\sigma_{\text{ref}}^2/2$ looks small when plotted in the form of Fig. 1 it might be worth attempting to compute the correction (18). If not, a better reference model ansatz is probably required.

F. Efficiency of the correction calculation

Standard arguments are often used to suggest that computing the exact value of $J(a)$ is not possible without knowledge of the exact rare dynamics, or the use of methods such as cloning or transition-path sampling. This claim is based on the fact that computing $J(a)$ requires computation of the integral in (15), and assumes that because the integrand grows exponentially with T , the number of trajectories required to converge this expression as $T \rightarrow \infty$ is prohibitively large. While this latter statement is correct, it does not speak directly to the task at hand. To compute the integral we do not need to take $T \rightarrow \infty$. Instead, we need $T \gtrsim T'$ where T' is smallest time at which the large deviation principle applies. Taking $T \gg T'$ makes sampling more difficult and is unnecessary. It is possible for T' to be large, but if the reference-model dynamics are close enough to the rare dynamics of the original model then the variance of the q will be small and it will be possible to converge this integral term without using an unreasonable number of trajectories.

Moreover, computing the correction is numerically

cheaper than inspection of the integral alone might suggest. If we switch to the SCGF representation, (22), we can instead work with trajectories of length $T = T_{\text{corr}}$, where T_{corr} is the correlation time of the reference model (in practice these trajectories are constructed by sub-sampling a single longer trajectory). Generally, T_{corr} is much smaller than T' , the time required for the large-deviation principle to apply, and so sampling the moderately rare events required to reconstruct the auxiliary rate function $\tilde{J}(\delta q, a)$ near its minimum is cheaper in the SCGF representation. This property is ideal for the present method because we have reduced the problem of sampling $J(a)$ arbitrarily far from its minimum to one of sampling $\tilde{J}(\delta q, \tilde{a}_0)$ (potentially) close to its minimum. Working with the SCGF also removes the constraint $a = \tilde{a}_0$ present in (18). Finally, we note that we are calculating the SCGF that is Legendre dual to the correction term (18), and not the SCGF that is Legendre dual to the original rate function $J(a)$. Thus our method can in principle reproduce rate functions $J(a)$ that are not convex (J_0 is not required to be convex).

G. Summary – A variational ansatz for rare dynamics (VARD)

1. Given a continuous-time dynamics with rates W_{xy} and a path-extensive dynamical observable a , we wish to determine $J(a)$, the large-deviation rate function for a for trajectories of the model W_{xy} of fixed time (assuming that $J(a)$ exists). We use a reference dynamics to calculate $J(a)$ as the sum of an upper bound $J_0(a)$ and a correction $J_1(a)$. The bound can always be calculated, and the correction can be calculated if the criteria described in Section II E hold. If so then we succeed in calculating $J(a)$; if not, then the method returns an upper bound $J_0(a) \geq J(a)$.
2. Determine a reference-model dynamics (7) and (9) able to produce more or less of a than the original model. In this paper we set the arbitrary bias β_{xy} using physical intuition.
3. Run a set of reference-model trajectories for different values of the reference-model parameters (s, λ, β_{xy}). For each trajectory, evaluate \tilde{a}_0 and \tilde{q}_0 , using Eqs. (6) and (11), and then use Eq. (17) to plot the point $(\tilde{a}_0, J(\tilde{a}_0))$; see Fig. 1(a). The lower envelope of these points over values of the reference-model parameter set is the tightest upper bound on $J(a)$ associated with the ansatz chosen in Step 1.
4. Attempt to calculate the correction $J_1(a)$ at points on the bound by running a few (~ 5) trajectories for each reference model. With the data from each trajectory, calculate $\tilde{J}(\delta q^*, \tilde{a}_0)$, Eq. (26), using Eqs. (27), (28) and (29). Insert the resulting values into Eq. (19). Calculate the averaged correction $\hat{J}_1(a)$

and associated statistical error using Eqs. (30) and (31).

5. To verify convergence, repeat the calculation of Step 3 for increasing values of the block time B , until the averaged estimate for the correction Eq. (30) no longer changes with B , and the convergence criterion (35), with $\delta q = \delta q^*$ and $a = \tilde{a}_0$, holds [75]. The accompanying GitHub script [62] performs steps 4 and 5 automatically.

III. VARD APPLIED TO FOUR EXAMPLES

A. Summary of the section

We now apply the method to four models taken from the literature. In each case, a simple and physically-motivated ansatz for the modified dynamics allows computation of the exact rate function $J(a)$. We focus on models whose state space is small enough that the exact rate function can be computed by standard methods – Legendre transform of the SCGF calculated using the largest eigenvalue of the tilted rate matrix [28] – in order to validate our method (at the end of the section we also use VARD to compute descriptive rate-function bounds for two systems too large to solve by matrix diagonalization). In figures, the exact rate function is shown as a black dashed line. Absent the exact answer we would apply the method in exactly the same way. For a given reference model the fluctuations of the quantity q reveal whether the bound $J_0(a)$ is tight, and whether we can compute $J(a)$ exactly.

B. Entropy production in a 4-state model

We start by sampling entropy production in the 4-state model of Ref. [54], shown in the inset of Fig. 2(a). This is a fully-connected network model with transition rates

$$\begin{aligned} W_{12} &= 3, & W_{13} &= 10, & W_{14} &= 9, \\ W_{21} &= 10, & W_{23} &= 1, & W_{24} &= 2, \\ W_{31} &= 6, & W_{32} &= 4, & W_{34} &= 1, \\ W_{41} &= 7, & W_{42} &= 9, & W_{43} &= 5. \end{aligned} \tag{36}$$

These rates do not satisfy detailed balance, and so the model produces nonzero entropy on average. To quantify the fluctuations of entropy production for trajectories of fixed time we construct a reference model as follows. The dynamical observable is $a = \sigma = T^{-1} \sum \alpha_{xy}$, where the sum is taken over all jumps $x \rightarrow y$ of a trajectory, and $\alpha_{xy} = \ln(p_{xy}/p_{yx})$, where $p_{xy} = W_{xy} / \sum_y W_{xy}$. Our basic reference-model parameterization is defined by the parameters s and λ appearing in (7) and (9), together with any additional set of biases β_{xy} suggested by the physics of the problem under study. Here we reason that none is necessary: the bias λ is always required, in order

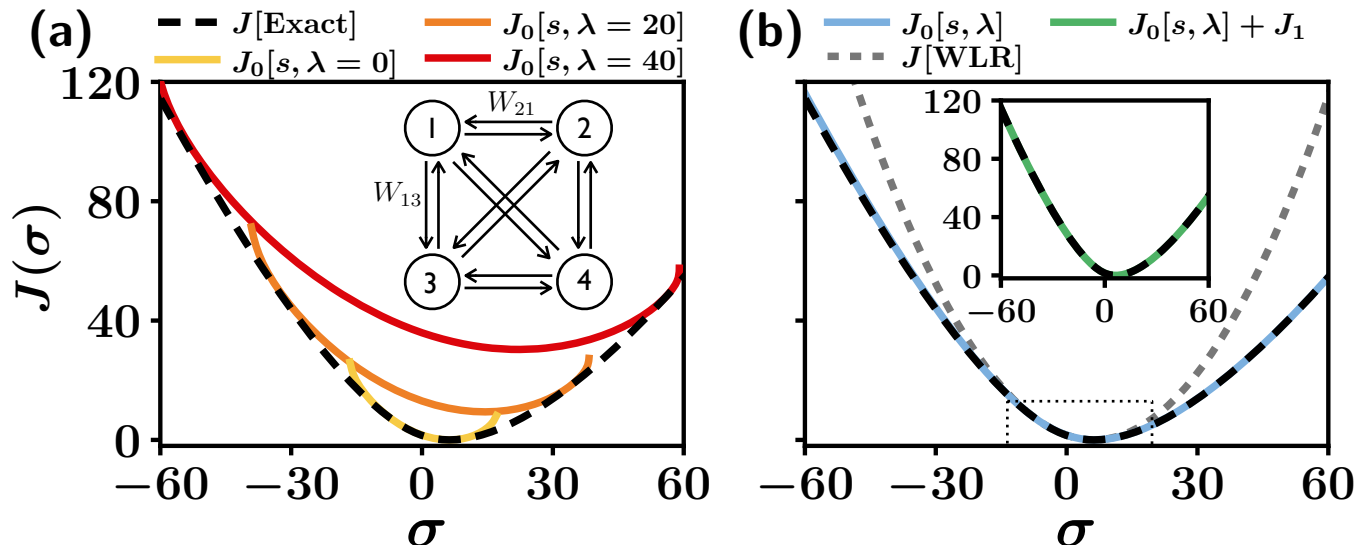


FIG. 2. Entropy production rate σ for the 4-state model of Ref. [54], the rates W_{xy} of which are given by (36). (a) Bounds (17) for the large-deviation rate function $J(\sigma)$. Each point on the colored lines results from a reference-model dynamics [Eqns. (7) and (9)] obtained by scanning the indicated parameters. The black dashed line is the exact answer. (b) Bound (blue line) resulting from a scan of the parameters s and λ compared with the WLR bound of Ref. [54], Eq. (37). Inset: The sum (green line) of the bound and the correction (18) equals the exact rate function, arbitrarily far into its tails. The boxed region in panel (b) indicates the scale of Fig. 2 of Ref. [54]. Error bars for both axes are smaller than the thickness of the lines.

to sample jump times, and the bias s is sufficient to influence the entropy produced per jump, A/K (a fact that is easy to guess, and to confirm with some short simulations). We therefore impose no additional bias, and set $\beta_{xy} = 0$.

We ran trajectories for a fixed number $K = 1.5 \times 10^8$ of events, roughly equivalent to a time of $T = 10^7$ in the unbiased model. We simulate in the constant-event ensemble for convenience, because there all simulations, regardless of the value of λ , take approximately the same amount of processor time. The equations of Section II then allow us to compute the rate function for the original model in the constant-time ensemble (see Ref. [76] for more on the relationship between the constant-event and constant-time ensembles).

The bounds $J_0(a)$ resulting from a scan of s , for three values of λ , are shown as colored lines in Fig. 2(a). In figures we use the compact notation $J_0[x]$ to indicate the bound $J_0(a)$ swept out by scanning the set of parameters $\{x\}$. We also show the exact rate function (black dashed line), obtained by matrix diagonalization. Different combinations of s and λ produce the best (lowest) bound at different values of a , so indicating the “least unlikely way” of realizing each value of a within the manifold of dynamics accessible to the reference model. The bound produced by the scan $\lambda = 0$ provides the best bound close to the typical value a_0 , but not far from it, indicating that very rare values of a are produced by the original model using a combination of rare jump types ($s \neq 0$) and rare jump times ($\lambda \neq 0$).

The bound swept out by scanning both s and λ is

shown in blue in Fig. 2(b). We used 201 equally spaced s values on the interval $[-5, 5]$, and 51 equally spaced λ values on $[0, 50]$. This bound lies close to the exact answer, even far into the tails of the rate function. For comparison we show the weakened linear response (WLR) universal current bound of Ref. [54] (gray dashed line); the dotted box in the center of the figure indicates the scale of Fig. 2 of that paper. The WLR bound is

$$J[\text{WLR}] = \frac{\sigma_0}{4c_0^2}(c - c_0)^2, \quad (37)$$

where c is a current, c_0 is its typical value (in the original model), and σ_0 is the typical value of the rate of entropy production. c_0 and σ_0 must be determined by running a single trajectory of the original model, and (37) then provides a bound on the probability of observing an atypical value of c . The bound is tightest in the case $c \propto \sigma$.

The WLR bound can be derived from Level 2.5 of large deviations [56, 57] using a mean-field ansatz that assumes all currents scale with time in the same way (for both forward and time-reversed versions of the model). By contrast, the (s, λ) -bound results from a microscopic ansatz, (7) and (9), inserted into the exact result (5) for the dynamical partition sum, and does not assume that all currents scale in the same way. Inspection of the tails of the bounds reveals that the microscopic ansatz captures the rare behavior of the model more accurately than does the homogeneous ansatz. Thus we learn that the rare behavior of even this very simple model does not simply resemble a speeded-up or slowed-down version of its typical behavior. The bound (37) is a universal statement

about the physics that constrain fluctuations, and is not designed to be a means of accurate numerical sampling. Nonetheless, it is meaningful and instructive to compare the bounds produced by different types of ansätze.

For each of the reference models that lie on the bound $J_0[s, \lambda]$ we calculate the correction (18) using the procedure described in Section IID. For all points we obtain convergence of the correction. The result, $J_0(a) + J_1(a)$, is shown as a green line in the inset of Fig. 2(b), and matches the exact answer (black), obtained by matrix diagonalization, as it should. We used 10^4 blocks, each of length of $50T_{\text{event}}$, where T_{event} is the time per event for each reference model. The average correction (30) and statistical error are obtained from 10 independent data sets. Error bars on the rate function are calculated by combining the error from the correction and the error from the bound according to

$$\text{Err}[\hat{J}(\tilde{a}_0)] = \sqrt{(\text{Err}[\hat{J}_0(\tilde{a}_0)])^2 + (\text{Err}[\hat{J}_1(\tilde{a}_0)])^2}. \quad (38)$$

The error in the bound is estimated by running 10 additional trajectories at each point and calculating the standard deviation of $J_0(\tilde{a}_0)$. The standard deviation of \tilde{a}_0 , calculated in the same way, yields error bars for the horizontal axis.

From this example it is clear that VARD is numerically much more efficient than direct simulation (of the original model): accurate calculation of the rate function at values of order 100 indicates accurate calculation of probabilities of order e^{-100T} , where T is the elapsed time of the trajectory. We do not know in advance which values of reference-model parameters constitute good choices, for particular values of a , but it is a simple matter to scan parameters and pick the smallest value of J_0 given a . Additional sampling then allows us to determine if we can calculate the correction (18), and therefore the exact rate function. We were able to do this here with little additional numerical effort. In this example, the state space of the model is small enough that its dynamics can be solved by matrix diagonalization, and so we possess the exact answer in advance. We made this choice because we wish to benchmark the method. However, our procedure would be identical if we did not know the exact answer ahead of time: define the reference model, pick the best bound, and attempt to calculate the correction term. The results of the latter calculation tell us if we have the exact answer, or, if not, roughly how close we are to obtaining it. If we are not close at all then we need a better reference-model guess. Inspection of the bounds produced by different reference models is also physically instructive, indicating the extent to which certain dynamical mechanisms contribute to the rate function at particular values of the order parameter.

C. Current in the ASEP

We next sample current in the asymmetric simple exclusion process (ASEP), a model of hard particles that

hop between lattice sites [6, 65] (an interesting alternative would be to consider the symmetric simple exclusion process, which has fewer parameters but also shows complex scaling behavior [77]). We consider the version of the model studied in Fig. 3 of [54], shown in Fig. 3, with open boundaries and a lattice of $L = 15$ sites. The rate constants are $\alpha = 1.25, \beta = 0.5, \gamma = 0.5, \delta = 1.5, p = 1, q = 0.5$, placing the model in the high-density region of the ASEP phase diagram [78, 79]. The dynamical observable is $a = c = T^{-1} \sum \alpha_{xy}$, where the sum is taken over all jumps $x \rightarrow y$ of a trajectory, and $\alpha_{xy} = \pm 1$ if the jump $x \rightarrow y$ sees a particle move to the right (upper sign) or left (lower sign).

In Fig. 3(a) we show the bound swept out by our default (s, λ) reference-model parameterization (dark blue), which provides a meaningful numerical bound on the exact rate function (black) even far into the tails. We ran trajectories for $K = 5 \times 10^5$ events, roughly equivalent to a time $T = 10^5$ in the unbiased model. We scanned 81 equally spaced s values on the interval $[-2, 2]$, and 41 equally spaced λ values on $[0, 10]$. Also shown is the WLR universal current bound [53, 54] (gray). The WLR bound describes accurately the moderately rare behavior of the model, but not the very rare behavior, which is quantified by the tails of the rate function. Comparison of bounds indicates, as in the previous subsection, that very rare currents are generated by trajectories that do not resemble speeded-up or slowed-down versions of the forward- or backward-running typical dynamics: the configurations visited in the tails of the rate function are different to those visited near the center.

While the (s, λ) -bound is already meaningful, it is possible to produce tighter numerical bounds by guessing additional ways in which the very rare high- or low-current behavior might be achieved. Inspection of the way in which s couples to the rate constants (here any rate involving a hop to the right is multiplied by e^{-s} and any rate involving a hop to the left is multiplied by e^s) reveals that varying s moves the reference model around the ASEP phase diagram [78, 79]. The original model sits in the high-density region of phase space, but the reference model need not. Inspection of the phase diagram indicates that the end rates $\alpha, \beta, \gamma, \delta$, separate from the bulk rates p and q , play a key role in determining the ASEP's typical behavior: if particles are fed in relatively quickly or slowly then we reside in the high- or low-density region of phase space, respectively, and if input- and output rates are balanced then we can access the maximum-current region. Returning to (8) we introduce a set of parameters β_{xy} that couple to the end rates, such that the rate α in the original model becomes $e^{-u_\alpha \alpha}$ in the reference model (u_α being a parameter), and similarly for the three other end rates. We also include a contribution to β_{xy} that biases trajectories toward or away from creating particle-particle contacts (i.e. particles on adjacent sites), reasoning that controlling such contacts can help control the escape rate of visited configurations, so helping cause or prevent traffic jams. A simple way




FIG. 3. Large-deviation rate function $J(c)$ of particle current, c , for the version of the ASEP studied in Ref. [54]; rate constants are given in the text. (a) We show the exact answer (black), the WLR universal current bound [53, 54] (gray), and the bound produced using our default (s, λ) reference-model parameterization (dark blue). (b) With additional physical insight used to refine the reference model it is possible to tighten the bound (compare dark blue and light blue lines). The green line in (a) is the corrected version of the 6D bound. The boxed region in panel (a) indicates the scale of Fig. 3 of Ref. [54] (not the scale of Fig. 3(b)). Error bars for both axes are smaller than the thickness of the lines.

to do this is to add to β_{xy} a bias $-\mu\Delta_{xy}$, where μ is a parameter and Δ_{xy} is the change in the number of particle-particle contacts when moving from x to y .

With the new bias determined we can generate an improved bound for the ASEP. We split the calculation into two parts and focus separately on the piece of the rate function for values of the observable greater than the mean, $a > a_0$, and less than the mean, $a < a_0$. Since the dynamics in these two different regimes are qualitatively different, generating an effective set of reference models for each requires scanning over different regions of the ansatz parameter space. For $a < a_0$ we scanned s, λ, u_γ and μ . For $a > a_0$ we scanned $s, \lambda, u_\alpha, u_\beta$ and μ , making 6 parameters in total. Combining these calculations produces the 6D bound shown in light blue in Fig. 3(b). This bound is tighter than the default (s, λ) -bound. The 6-parameter scan can be carried out with reasonable numerical effort: for a given set of parameters we need only a short single trajectory to compute the averages required for the bound, and there is no requirement for communication between the calculations. On the left side we scanned 11 equally spaced s values on the interval $[0, 2]$; 21 equally spaced λ values on the interval $[-0.6, 0]$, and 17 more on the interval $[0, 8]$; 29 u_γ values so that γ takes on equally spaced values on the interval $[0.1, 1.5]$; and 31 equally spaced μ values on the interval $[-1, 0.5]$. On the right side we scanned 11 equally spaced s values on the interval $[-2, 0]$; 17 equally spaced λ values on the interval $[0, 8]$; 21 u_α values and 21 u_β values so that α and β each take on values equally spaced on the interval $[0.5, 1.25]$;

and 11 equally spaced μ values on the interval $[0, 0.5]$. The lower envelope of the values of (17) constitutes the improved bound.

Correcting the 6D bound by calculating the correction J_1 at points along the bound gives the green line shown in Fig. 3(a), which agrees with the exact rate function even far into the tails. For this calculation we used 10^4 blocks of length $100T_{\text{event}}$, where T_{event} is the time per event in each reference model. Errors are computed as in Section III B.

D. Activity in a 3-state model

We consider the three-state model of Fig. 3 of Ref. [55], shown in Fig. 4. The rate constants are $\gamma = 1$ and $\kappa = 1/2$. Our chosen dynamical observable, a , is the number of jumps from states $1 \rightarrow 0$ per unit time. The parameter s in our default (s, λ) reference-model parameterization (8) has no role to play here: s controls the probability of the $1 \rightarrow 0$ process, but once in state 1 there is nowhere to go but state 0. Thus s cannot influence the number of counted events per jump, A/K , and so we set $s = 0$.

At this point we need to apply our physical intuition in order to create a reference-model ansatz able to control A/K . Inspection of the network reveals that controlling the jump destination from state 0 is sufficient for this purpose: if we jump $0 \rightarrow 1$ then we must subsequently jump $1 \rightarrow 0$, whereas if we jump $0 \rightarrow 2$ we will return to 0 without making the counted jump. In Eq. (8) we therefore set




FIG. 4. Large-deviation rate function $J(a)$ for the number of jumps $1 \rightarrow 0$ per unit time, a , for the 3-state model of Ref. [55]. We compare with the exact answer (black) the bounds produced by reference models in which we control jump times (blue), or jump times and jumps from $0 \rightarrow 1$ (green). Error bars for both axes are smaller than the thickness of the lines. Also shown is the CMP universal activity bound [55], Eq. (39) (gray).

$\beta_{01} = \eta$ (a parameter) such that the reference-model rate for the process $0 \rightarrow 1$ is $\tilde{W}_{01} = e^{-\eta}$. We set all other $\beta_{xy} = 0$. Scanning η and λ (our usual jump-time bias) produces the bound shown in green in Fig. 4. Bounds were calculated using 11 equally spaced λ values on the interval $[-0.5, 0]$ and 101 values on the interval $[0, 25]$, and 51 equally spaced η values on the interval $[-5, 5]$. All trajectories were run for $K = 10^7$ events, roughly equivalent to a time of $T = 10^7$ in the unbiased model. Error bars are computed as in Section III B.

The (η, λ) -bound is effectively exact, as we can deduce by measuring the fluctuations of q (which here are nonexistent). In this case the model is simple enough that each reference model used to compute the bound enacts the exact rare dynamics of the original model, conditioned on a particular value of a . As a result, the correction term J_1 vanishes, and the bound J_0 is exact. (This exactness is reasonable on account of the fact that the system has relatively few ways of realizing values of A/K and K/T , but it is not obvious in advance that the chosen parameterization would require no additional correction.) Recall that the correction term can be interpreted as a measure of how close the typical dynamics of the reference model is to the desired rare dynamics of the original model; here, typical trajectories of the reference model correspond exactly to trajectories of the original model conditioned on the relevant value of the order parameter.

The chosen observable is a non-decreasing counting variable, not a current, and so the universal bound of Refs. [53, 54] does not apply. One that does apply is the Conway-Maxwell-Poisson (CMP) bound of Ref. [55],

$$J[\text{CMP}] = \frac{k_0}{a_0} \left(a \ln \frac{a}{a_0} + a_0 - a \right), \quad (39)$$

where a is the dynamical observable, a_0 is its typical value (in the original model), and k_0 is the typical dynamical activity (the total number of events per unit time) of the original model (note that there is an a missing in front of the logarithm in Eq. (17) of [55]).

The CMP bound is shown in gray in Fig. 4. Similar to the universal current bound, the CMP bound is derived from Level 2.5 of large deviations using an ansatz that assumes the rare behavior of the system to be a speeded-up or slowed-down version of its typical behavior. It therefore has similar properties to our λ -bound, shown in blue in Fig. 4 (the λ bound is constructed from the pieces of the (η, λ) -bound with $\eta = 0$). Comparison of this bound and the (λ, η) -bound shows the extent to which the very rare behavior of this model is dominated by trajectories comprising rare jump times *and* an atypical propensity to jump left from state 0. Analogous to its current counterpart, the CMP bound is a general statement about the physics controlling the fluctuations of counting variables, and is not intended to be a means of numerical sampling. Nonetheless, comparison of its properties with bounds obtained by the microscopic ansatz used here is instructive, and addresses the point raised in Ref. [55]: “It would be interesting to find alternative yet simple variational ansatzes that can capture [the] strong fluctuation behavior [of the 3-state model]”.

E. Activity in the FA model

We consider the one-dimensional Fredrickson-Andersen (FA) model with periodic boundary conditions [66]. This is a lattice model with simple thermodynamics and with dynamical rules that give rise to slow relaxation and complex spatiotemporal behavior [80]. On each site of a lattice (here of length $L = 15$) lives a spin, which can be up or down. Up-spins (resp. down-spins) can flip down (resp. up) with rate $1 - c$ (resp. c) if at least one of their neighboring spins is up; if not, then they cannot flip (the rate c here should not be confused with the symbol for current in previous sections). In Fig. 5(a), top, we show an example FA model configuration, with periodic boundary conditions; the spins in red cannot flip. Our chosen dynamical observable, $a = k = T^{-1} \sum \alpha_{xy}$, is the number of jumps per unit time, where $\alpha_{xy} = 1$ for all jumps $x \rightarrow y$. The large-deviation properties of k have been studied in detail, and give rise, in certain limits, to singular behavior in the SCGF that is Legendre dual to $J(k)$ [10, 11].

In Fig. 5(a) we show the CMP universal activity bound [55] on $J(k)$ and the bound produced by our

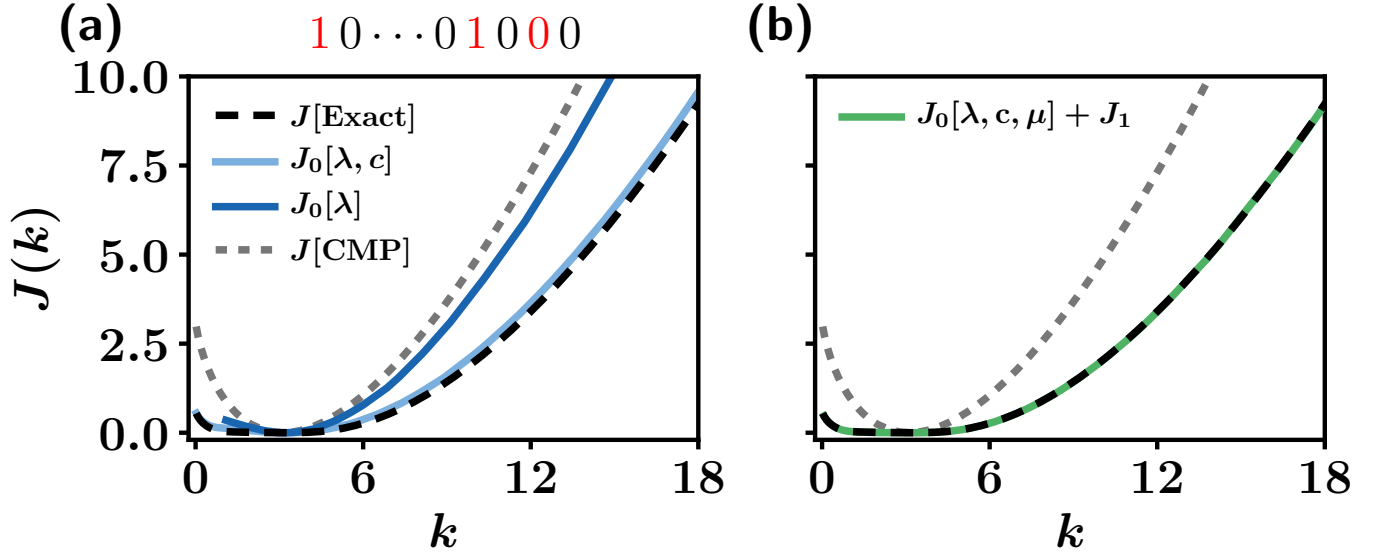


FIG. 5. Large-deviation rate function $J(k)$ for the number of jumps per unit time, k , for the one-dimensional Fredrickson-Andersen model. (a) We compare with the exact answer (black) the bounds produced by reference models in which we control jump times (dark blue) or jump times and the mean up-spin fraction (light blue). Also shown is the CMP universal activity bound [55], Eq. (39) (gray). (b) The sum (green) of the bound (produced using a reference model in which we control jump times, mean up-spin fraction, and pair correlations) and the correction matches the exact answer. Error bars for both axes are smaller than the thickness of the lines.




FIG. 6. Large-deviation rate function $J(k)$ for the number of jumps per unit time, k , for the Fredrickson-Andersen model of Ref. [81] with $c = 0.1$. Lattice sizes are (a) $L = 40$ sites and (b) $L = 100$ sites. We compare the matrix product state (MPS) calculation of Ref. [81] (black) with the three-parameter VARD bound of Fig. 5 (blue). Also shown is the CMP universal activity bound [55], Eq. (39) (gray).

reference-model λ -scan. These are of similar character, because they assume that the rare behavior of the model is a speeded-up or slowed-down version of its typical behavior. All trajectories were run for $K = 3 \times 10^5$ events, roughly equivalent to a time of $T = 10^5$ in the unbiased model. We used 61 equally spaced λ values on the inter-

val $[-0.6, 0]$, and 61 more on the interval $[0, 12]$. To produce a tighter bound we need to assume that the configurations visited by rare trajectories are different to those visited by typical ones (the CMP bound assumes that they are the same). The parameter s in our default (s, λ) reference-model parameterization (8) again has no role to

play, because biasing all jumps equally is equivalent to biasing none. A simple alternative is to choose the bias β_{xy} so that the reference model can generate a larger or smaller number of up-spins than is typical in the original model. We choose the parameters β_{xy} so that the parameter c in the original model becomes $e^{-\eta}c$ in the reference model, with η being a parameter. A (λ, η) -scan of the reference model produces the bound $J_0[\lambda, c]$ shown in light blue in Fig. 5(a). This bound provides a reasonable approximation of the exact answer over the whole range of k , indicating that much of the physics of rare activity fluctuations of the FA model can be accounted for by considering the typical behavior of versions of the model with different values of the parameter c . Here we chose η so that c takes on 100 equally spaced values on the interval $[0.02, 2]$.

It is possible to tighten this bound by reasoning that there must exist spatial correlations between up-spins if we condition trajectories upon activity per unit time k . For instance, given an up-spin fraction of exactly $1/2$, the escape rate (the sum of rates leading out of a given state, a quantity relevant to the number of jumps per unit time) is maximized by having pairs of up-spins separated by pairs of down-spins, and minimized by having up-spins and down-spins alternate. We induce these types of spatial correlations using the same μ -bias used for the ASEP in Section III C, favoring more or fewer contacts between up spins. Scanning c, η , and μ produces a bound slightly tighter than $J_0[\lambda, \eta]$ (not shown). We used 41 equally spaced μ values on the interval $[-1, 1]$. From this bound we compute the correction J_1 , and the sum of the bound and correction matches the exact answer: see Fig. 5(b). We computed the correction by splitting the calculation into two pieces, one on either side of the mean value $a = a_0$. For $a < a_0$ we used 10^4 blocks of length $50T_{\text{event}}$, where T_{event} is the average time per event in each reference model. For $a > a_0$ we used 10^5 blocks of length $600T_{\text{event}}$. The different block lengths generated by the convergence procedure (see Section II E) for $a < a_0$ and $a > a_0$ signal that the correlations present in the dynamics are qualitatively different in each of these regimes. Understanding the nature of these correlations is of physical interest [11, 80]. Errors are calculated as in Section III B.

F. Toward large-scale calculations

In this paper we have demonstrated proof-of-principle of VARD using network systems or lattice models whose state space is small enough that their rate functions can be obtained by matrix diagonalization, so providing a benchmark for the method. VARD can also be applied to systems too large for matrix diagonalization to be feasible, in order to produce bounds or (if the ansatz is good enough and convergence of the correction is obtained) exact rate functions. An active line of research is to study large versions of certain lattice models in order to de-

termine how their large-deviation properties change with system size [11, 34, 81]. In these regimes, specialized techniques are necessary. For instance, in Ref. [34], a cloning procedure combined with feedback control was used to calculate large-deviation functions for an FA model of $L = 36$ sites. In Ref. [81], a matrix product state (MPS) calculation was used to compute large-deviation functions for an FA model of sizes of order $L = 100$ sites (these results show some differences with the results of Ref. [34], indicating that this is a technically challenging regime). Note that the FA model of Ref. [81] has open boundaries and slightly different facilitation rules than used in the previous section: spins facilitated by two spins flip at twice the rate of spins facilitated by one spin.

In Fig. 6 we compare the MPS calculation of Ref. [81] with the three-parameter VARD bound used in Fig. 5 for two lattice sizes that are considered large by current standards (the parameter $c = 0.1$). In both cases the VARD bound is descriptive, capturing the main features and the trends with k of the MPS result. The bound is less tight for the larger system size, suggesting that more terms in the ansatz are required as the system becomes larger. However, the bound quality, even using an ansatz containing only 3 parameters, remains reasonable. As for the ASEP, the natural next step is to include additional parameters in the ansatz in order to tighten the bound and calculate the correction (compare dark blue and light blue lines in Fig. 3(b)). A natural way to develop improved bounds is to use Monte Carlo learning procedures in order to optimize reference models containing a potentially large number of parameters [82].

IV. CONCLUSIONS

We have described how direct simulation of a variational ansatz for rare dynamics (VARD) can be used to compute bounds for large-deviation rate functions in continuous-time Markov chains [83]. This approach requires only direct simulation of versions of the original model with modified rates, and so is technically simple and easy to implement. It is also physically instructive, in the sense that the quality of the bounds produced by different physical ansätze reveal the extent to which different types of dynamical processes contribute to the rare behavior of the model of interest.

If the ansatz is chosen well then bounds can be corrected to produce the exact rate function, arbitrarily far into its tails; in the literature it is often assumed or stated that such precision is not accessible via direct reweighting of trajectories, and requires the use of specialized numerical techniques such as cloning or path sampling. For the models studied in Figs. 2–5, two network models and two lattice models taken from the literature, it is possible to calculate the exact rate function using only simple and approximate guesses about the nature of the rare dynamics. Although this rare behavior can be complex,

we are rarely working in the dark: the model itself can exhibit different behavior in different parameter regions, and often its rare behavior at one point in parameter space is similar to its typical behavior at another point in parameter space. For example, we have studied the ASEP in its high-density region, where (typically) the lattice is crowded and particles move slowly. The ansatz we used to calculate its current rate function is equivalent to guessing that the rare, high-current behavior in the high-density region is similar to the typical behavior in the maximum-current region, where particles move quickly and possess spatial anticorrelations. Similarly, the FA model is complex, but the likelihood of its rare behavior at one value of the parameter c can be well approximated by looking at the typical behavior of models at different values of c . We have also shown that bounds that are descriptive in small systems remain so in systems too large to solve by matrix diagonalization; we will discuss this regime further in forthcoming work.

VARD is similar to classical umbrella sampling [63, 64] in the sense that the rate function $\tilde{J}(a)$ of the reference model can be regarded as a nonequilibrium umbrella potential, concentrating sampling at a desired point: see Fig. 1. It is different, however, in that VARD does not require overlapping sampling windows – reference models are used independently – and we compute an absolute rate-function value $J(a)$, as opposed to a free-energy difference. This latter distinction results from the fact that the path weight appearing in (5) is known exactly, and at the sampling point \tilde{a}_0 we know that the rate function of the reference model vanishes; by contrast, in the equilibrium case we know the probability of visiting a certain state only up to a normalization constant, and we do not know the absolute free energy of the reference model (unless it is particularly simple [84]).

VARD provides insight into the approaches used to produce universal rate-function bounds from Level 2.5 of large deviations, via homogeneous ansätze [53–55], by showing how relaxing such assumptions leads to the tightening of bounds in different sectors of parameter space. It is also complementary to numerical large-deviation methods that use path-sampling, cloning, or adaptive methods to calculate the SCGF [85] that is Legendre dual to $J(a)$ [11, 28, 30–35]. Sometimes path sam-

pling or cloning are used in isolation, and sometimes they are combined with a modified dynamics. VARD lies at the other extreme of the methods spectrum in the sense that it uses *only* a modified dynamics. The bounds that result provide a natural starting point for those specialized methods, because the set of reference models that live on the bounds already resemble the rare dynamics of interest. Indeed, direct sampling of those reference models is, in the cases described in Figs. 2–5, sufficient to recover the statistics (the constituent configurations and jump times) required to compute $J(a)$ exactly.

There are several possible variants of the present method. The λ -scan accesses roughly the same information as the universal bounds of Refs. [53–55], and one possible numerical simplification would be to eliminate the λ -scan in favor of the universal bounds (37) or (39). We have also simply scanned parameters in order to identify the best bounds associated with a given ansatz, but as the complexity of an ansatz grows it is natural to replace the scan with an evolutionary Monte Carlo procedure [82]. For instance, consider a set of reference models having N parameters, and construct an initial (s, λ) -bound. Take models at various points on this bound, and define a window Δa for the observable. For each model, perturb the N parameters, generate a short trajectory, and calculate Eq. (17). If this value is less than the current bound (and a lies within the designated window) accept the new reference model; otherwise, retain the original.

V. ACKNOWLEDGMENTS

We thank Hugo Touchette, Tom Ouldridge, and Juan Garrahan for discussions and comments on the manuscript, and thank Juan Garrahan, Todd Gingrich, and Mari Carmen Bañuls for providing data from Refs. [55], [16], and [81], respectively. This work was performed as part of a user project at the Molecular Foundry, Lawrence Berkeley National Laboratory, supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. DJ acknowledges support from the Department of Energy Computational Science Graduate Fellowship.

-
- [1] D. T. Gillespie, *Annu. Rev. Phys. Chem.* **58**, 35 (2007).
 - [2] T. McGrath, N. S. Jones, P. R. ten Wolde, and T. E. Ouldridge, *Physical Review Letters* **118**, 028101 (2017).
 - [3] U. Seifert, *Reports on progress in Physics* **75**, 126001 (2012).
 - [4] A. I. Brown and D. A. Sivak, *Proceedings of the National Academy of Sciences* **114**, 11057 (2017).
 - [5] P. Visco, A. Puglisi, A. Barrat, E. Trizac, and F. van Wijland, *Journal of statistical Physics* **125**, 533 (2006).
 - [6] T. Chou, K. Mallick, and R. Zia, *Reports on progress in Physics* **74**, 116601 (2011).
 - [7] S. Vaikuntanathan, T. R. Gingrich, and P. L. Geissler, *Physical Review E* **89**, 062108 (2014).
 - [8] R. J. Harris, *Journal of Statistical Mechanics: Theory and Experiment* **2015**, P07021 (2015).
 - [9] L. Berthier, *Physical Review Letters* **112**, 220602 (2014).
 - [10] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, *Physical Review Letters* **98**, 195702 (2007).

- [11] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, *Journal of Physics A: Mathematical and Theoretical* **42**, 075007 (2009).
- [12] F. Ritort, *Advances in Chemical Physics* **137**, 31 (2008).
- [13] U. Seifert, *Physical Review Letters* **95**, 040602 (2005).
- [14] T. Speck, A. Engel, and U. Seifert, *Journal of Statistical Mechanics: Theory and Experiment* **2012**, P12001 (2012).
- [15] V. Lecomte, A. Imparato, and F. v. Wijland, *Progress of Theoretical Physics Supplement* **184**, 276 (2010).
- [16] T. R. Gingrich and J. M. Horowitz, *Physical Review Letters* **119**, 170601 (2017).
- [17] É. Fodor, M. Guo, N. Gov, P. Visco, D. Weitz, and F. van Wijland, *EPL (EuroPhysics Letters)* **110**, 48005 (2015).
- [18] D. Chandler, *Introduction to Modern Statistical Mechanics*, by David Chandler, pp. 288. Foreword by David Chandler. Oxford University Press, Sep 1987. ISBN-10: 0195042778. ISBN-13: 9780195042771, 288 (1987).
- [19] J. J. Binney, N. Dowrick, A. Fisher, and M. Newman, *The theory of critical phenomena: an introduction to the renormalization group* (Oxford University Press, Inc., 1992).
- [20] G. Gallavotti and E. G. D. Cohen, *Physical Review Letters* **74**, 2694 (1995).
- [21] C. Maes, *Journal of Statistical Physics* **95**, 367 (1999).
- [22] C. Jarzynski, *Physical Review Letters* **78**, 2690 (1997).
- [23] J. Kurchan, *Journal of Physics A: Mathematical and General* **31**, 3719 (1998).
- [24] J. L. Lebowitz and H. Spohn, *Journal of Statistical Physics* **95**, 333 (1999).
- [25] G. E. Crooks, *Physical Review E* **60**, 2721 (1999).
- [26] R. J. Harris and G. M. Schütz, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P07020 (2007).
- [27] F. Den Hollander, *Large Deviations*, Vol. 14 (American Mathematical Soc., 2008).
- [28] H. Touchette, *Physics Reports* **478**, 1 (2009).
- [29] More directly it quantifies the rate of decay of a fluctuation a , which depends both on the likelihood of a and the basic timescale governing the establishment and decay of fluctuations. This latter piece plays a key role for certain models [86–88].
- [30] C. Giardinà, J. Kurchan, and L. Peliti, *Physical Review Letters* **96**, 120603 (2006).
- [31] R. Chetrite and H. Touchette, *Journal of Statistical Mechanics: Theory and Experiment* **2015**, P12001 (2015).
- [32] R. L. Jack and P. Sollich, *The European Physical Journal Special Topics* **224**, 2351 (2015).
- [33] U. Ray, G. K.-L. Chan, and D. T. Limmer, *Physical Review Letters* **120**, 210602 (2018).
- [34] T. Nemoto, R. L. Jack, and V. Lecomte, *Physical Review Letters* **118**, 115702 (2017).
- [35] G. Ferré and H. Touchette, arXiv preprint arXiv:1803.11117 (2018).
- [36] R. Chetrite and H. Touchette, *Physical Review Letters* **111**, 120601 (2013).
- [37] V. Lecomte and J. Tailleur, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P03004 (2007).
- [38] A. Kundu, S. Sabhapandit, and A. Dhar, *Physical Review E* **83**, 031119 (2011).
- [39] K. Klymko, P. L. Geissler, J. P. Garrahan, and S. Whitelam, *Phys. Rev. E* **97**, 032123 (2018).
- [40] S. Whitelam, *Physical Review E* **97**, 032122 (2018).
- [41] S. Whitelam, *Journal of Statistical Mechanics: Theory and Experiment* **2018**, 063211 (2018).
- [42] J. A. Bucklew, P. Ney, and J. S. Sadowsky, *Journal of Applied Probability* **27**, 44 (1990).
- [43] P. W. Glynn and D. L. Iglehart, *Management Science* **35**, 1367 (1989).
- [44] J. S. Sadowsky and J. A. Bucklew, *IEEE transactions on Information Theory* **36**, 579 (1990).
- [45] P. Glasserman, Y. Wang, *et al.*, *The Annals of Applied Probability* **7**, 731 (1997).
- [46] P. B. Warren and R. J. Allen, *Molecular Physics* **116**, 3104 (2018).
- [47] J. A. Bucklew, *Large deviation techniques in decision, simulation, and estimation* (Wiley New York, 1990).
- [48] S. Asmussen and P. W. Glynn, *Stochastic simulation: algorithms and analysis*, Vol. 57 (Springer Science & Business Media, 2007).
- [49] S. Juneja and P. Shahabuddin, *Handbooks in operations research and management science* **13**, 291 (2006).
- [50] J. Bucklew, *Introduction to rare event simulation* (Springer Science & Business Media, 2013).
- [51] The “typical” value of a time-integrated observable is the value to which the sample mean of an unbiased estimator of that quantity converges at long times, provided that a large-deviation principle exists and that the associated rate function has a unique zero [28].
- [52] S. S. Varadhan, in *Proceedings of the International Congress of Mathematicians 2010 (ICM 2010) (In 4 Volumes) Vol. I: Plenary Lectures and Ceremonies Vols. II–IV: Invited Lectures* (World Scientific, 2010) pp. 622–639.
- [53] P. Pietzonka, A. C. Barato, and U. Seifert, *Physical Review E* **93**, 052145 (2016).
- [54] T. R. Gingrich, J. M. Horowitz, N. Perunov, and J. L. England, *Physical Review Letters* **116**, 120601 (2016).
- [55] J. P. Garrahan, *Physical Review E* **95**, 032134 (2017).
- [56] C. Maes and K. Netočný, *EPL (EuroPhysics Letters)* **82**, 30003 (2008).
- [57] L. Bertini, A. Faggionato, D. Gabrielli, *et al.*, in *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, Vol. 51 (Institut Henri Poincaré, 2015) pp. 867–900.
- [58] P. Tsoibgni Nyawo and H. Touchette, *Phys. Rev. E* **94**, 032101 (2016).
- [59] R. L. Jack and P. Sollich, *Journal of Physics A: Mathematical and Theoretical* **47**, 015003 (2013).
- [60] U. Ray, G. K. Chan, and D. T. Limmer, arXiv preprint arXiv:1708.00459 (2017).
- [61] C. M. Rohwer, F. Angeletti, and H. Touchette, *Physical Review E* **92**, 052104 (2015).
- [62] D. Jacobson and S. Whitelam, “vard_correction,” https://github.com/drdrjacobs/vard_correction.git (2019).
- [63] G. Torrie and J. Valleau, *Journal of Computational Physics* **23**, 187 (1977).
- [64] P. R. ten Wolde and D. Frenkel, *The Journal of Chemical Physics* **109**, 9901 (1998).
- [65] B. Derrida, *Physics Reports* **301**, 65 (1998).
- [66] G. H. Fredrickson and H. C. Andersen, *Physical Review Letters* **53**, 1244 (1984).
- [67] K. Binder, in *Monte Carlo Methods in Statistical Physics* (Springer, 1986) pp. 1–45.
- [68] D. T. Gillespie, *The journal of physical chemistry* **81**, 2340 (1977).
- [69] I. H. Dinwoodie and S. L. Zabell, *The Annals of Probability* **20**, 1147 (1992).

- [70] R. S. Ellis, *Large deviations and statistical mechanics*, Vol. 17 (Springer, New York, 1985).
- [71] In an abuse of notation we have, for brevity, changed from writing the joint rate function in the form $\tilde{J}(q, a)$ in (15) to $\tilde{J}(\delta q, a)$ in (18) and subsequently.
- [72] Note that the reference model can be well behaved in this manner even when the *original* model exhibits anomalous fluctuations such that $J(a)$ is not quadratic about its minimum [39].
- [73] K. Duffy and A. P. Metcalfe, *Journal of applied probability* **42**, 267 (2005).
- [74] The choice of moving in the positive k_a direction is arbitrary. Moving in the negative k_a direction will yield an equivalent convergence criterion (similarly we can move in either the positive or negative $k_{\delta q}$ directions while holding $k_a = 0$).
- [75] The quantity $\hat{J}(a')$ on the right-side of (35) is determined from the Legendre transform (23) of $\hat{\theta}(k_{\delta q} = 0, k_a = \hat{k}'_a)$. The quantity \hat{k}'_a is calculated by finding the maximum of $\text{Err}[\hat{a}(k_a)]$, Eq. (33), while starting at the point $(k_{\delta q} = 0, k_a = 0)$ on the SCGF and scanning outwards by increasing k_a .
- [76] A. A. Budini, R. M. Turner, and J. P. Garrahan, *Journal of Statistical Mechanics: Theory and Experiment* **2014**, P03012 (2014).
- [77] B. Derrida, J. Lebowitz, and E. Speer, *Journal of statistical physics* **107**, 599 (2002).
- [78] A. B. Kolomeisky, G. M. Schütz, E. B. Kolomeisky, and J. P. Straley, *Journal of Physics A: Mathematical and General* **31**, 6911 (1998).
- [79] R. Blythe, M. Evans, F. Colaiori, and F. Essler, *Journal of Physics A: Mathematical and General* **33**, 2313 (2000).
- [80] J. P. Garrahan and D. Chandler, *Physical Review Letters* **89**, 035704 (2002).
- [81] M. C. Bañuls and J. P. Garrahan, arXiv preprint arXiv:1903.01570 (2019).
- [82] S. Whitelam, D. Jacobson, and I. Tamblyn, arXiv preprint arXiv:1909.00835 (2019).
- [83] The method works for Markov chains in discrete time, upon replacing the distributions (4) and (9) by $\delta(\Delta t - 1)$; see earlier versions of the method [39–41].
- [84] D. Frenkel and A. J. Ladd, *The Journal of Chemical Physics* **81**, 3188 (1984).
- [85] One advantage of bypassing the SCGF and calculating the rate function directly, as we do here, is the ability to reconstruct rate functions that are not strictly convex [28, 39, 89].
- [86] K. Spiliopoulos, *Applied Mathematics & Optimization* **67**, 123 (2013).
- [87] F. Bouchet, T. Grafke, T. Tangarife, and E. Vandeneijnden, *Journal of Statistical Physics* **162**, 793 (2016).
- [88] S. Whitelam, *Physical Review E* **97**, 062109 (2018).
- [89] D. Nickelsen and H. Touchette, arXiv preprint arXiv:1803.05708 (2018).