



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

Sampling rare fluctuations of discrete-time Markov chains

Stephen Whitelam

Phys. Rev. E **97**, 032122 — Published 19 March 2018

DOI: [10.1103/PhysRevE.97.032122](https://doi.org/10.1103/PhysRevE.97.032122)

Sampling rare fluctuations of discrete-time Markov chains

Stephen Whitelam*

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

We describe a simple method that can be used to sample the rare fluctuations of discrete-time Markov chains. The method is inspired by the ‘s-ensemble’ method of trajectory sampling but is different in several respects: for one, it computes probabilities directly, rather than via Legendre transform, and so can be used to calculate non-convex rate functions. We focus on the case of Markov chains with well-defined steady-state measures, and derive expressions for the large-deviation rate functions (and upper bounds on such functions) for dynamical quantities extensive in the length of the Markov chain. We illustrate the method using a series of simple examples, and use it to study the fluctuations of a lattice-based model of active matter that can undergo motility-induced phase separation.

I. INTRODUCTION

Quantifying the rare behavior of stochastic models is necessary to understand important physical processes such as phase separation and chemical reactions [1]. Computer simulation methods for quantifying rare behavior include umbrella sampling [2]; forward-flux sampling [3]; transition-path sampling [4]; and diffusion Monte Carlo methods [5–9], which often involve the use of an auxiliary or reference model whose typical behavior is in some sense equivalent to the rare behavior of the model of interest [5, 7, 8, 10–15]. Motivated by this latter approach, we show that the rare fluctuations of a discrete-time Markov chain generated by a stochastic model of interest can be quantified in a simple way using a stochastic *reference model* whose parameters are simply related to those of the original model. Partial information about the original model’s rare fluctuations – meaning a bound on the large-deviation rate function for time-extensive observables – can be obtained from typical trajectories of the reference model, while complete information about the original model’s rare fluctuations can be obtained by considering constrained fluctuations of the reference-model trajectory ensemble. In Ref. [16] we use this procedure to study the rare behavior of models of growth; here we focus on Markov chains with well-defined steady-state measures.

In what follows we describe the method, and use a series of simple examples to illustrate the method’s application to Markov chains with well-defined steady-state measures (Section II). We then apply the method to the lattice-based model of active matter introduced in Ref. [17], in order to quantify that model’s rare dynamic fluctuations (Section III). We conclude in Section IV.

II. SAMPLING THE RARE FLUCTUATIONS OF DISCRETE-TIME MARKOV CHAINS

Consider a discrete-time Markov chain of K steps. At a given instant the system is in microstate (or ‘state’) C , and with probability $p(C \rightarrow C')$ moves to state C' (which may be C). Probability conservation requires that $\sum_{C'} p(C \rightarrow C') = 1$. A trajectory \mathbf{x} is a sequence of microstates $\mathbf{x} = \{C_0, C_1, C_2, \dots, C_K\}$ generated by this dynamics, and its probability is

$$P[\mathbf{x}] = P_0(C_0) \prod_{k=0}^{K-1} p(C_k \rightarrow C_{k+1}), \quad (1)$$

where $P_0(C_0)$ is the probability of starting in state C_0 . The master equation for this discrete Markov chain is

$$\mathbf{P}(k+1) = \mathbf{W}\mathbf{P}(k), \quad (2)$$

where $\mathbf{P}(k)$ is the column vector whose elements are $P(i, k)$, the probability of residing in state i after k steps of the dynamics, and \mathbf{W} is the matrix whose $(j, i)^{\text{th}}$ element is $p(i \rightarrow j)$. For an M -state system, \mathbf{P} has M elements and \mathbf{W} has $M \times M$ elements.

Consider now an observable $A[\mathbf{x}]$, called the ‘activity’, extensive in the length of the trajectory [12]. Upon moving from C to C' we increment A by an amount $\alpha(C \rightarrow C')$, so that $A[\mathbf{x}] = \sum_{k=0}^{K-1} \alpha(C_k \rightarrow C_{k+1})$. Our aim is to compute the probability

$$\rho(a, K) \equiv \sum_{\mathbf{x}} P[\mathbf{x}] \delta(A[\mathbf{x}] - Ka) \quad (3)$$

that a trajectory of length K possesses a particular value $a \equiv A/K$ of the intensive counterpart of the observable A . For many models this probability adopts a large-deviation form $\rho(a, K) \sim e^{-KI(a)}$ for large values of K , where $I(a)$ is the large-deviation rate function [11]. Eq. (3) is the instruction to simulate the original model and count the trajectories that possess activity $A = Ka$. When a is close to a ‘typical’ value a_0 , for which $I(a_0) = 0$, such sampling is efficient. But when a is far from a_0 , so that it is rarely generated by direct simulation, a different strategy is required.

* swhitelam@lbl.gov

To motivate this strategy, note that the desired quantity $\rho(a, K)$ can be regarded as the normalization factor (or ‘partition function’) of the *microcanonical path ensemble* [18–21]

$$P_a[\mathbf{x}] = \frac{P[\mathbf{x}] \delta(A[\mathbf{x}] - Ka)}{\rho(a, K)}, \quad (4)$$

which defines an ensemble of trajectories conditioned upon the constraint $A[\mathbf{x}] = Ka$. (The term ‘microcanonical’ comes from analogy with the fixed-energy ensemble of equilibrium statistical mechanics; here it is activity that is fixed). Alternatively one can define a *canonical path ensemble* (or tilted ensemble or *s*-ensemble) [11] in which $A[\mathbf{x}]$ can fluctuate,

$$P_s[\mathbf{x}] = \frac{e^{-sA[\mathbf{x}]} P[\mathbf{x}]}{\sum_{\mathbf{x}} e^{-sA[\mathbf{x}]} P[\mathbf{x}]}. \quad (5)$$

The master equation in the *s*-ensemble is

$$\mathbf{P}_s(k+1) = \mathbf{W}_s \mathbf{P}_s(k), \quad (6)$$

where \mathbf{W}_s is the matrix whose $(j, i)^{\text{th}}$ element is

$$p_s(i \rightarrow j) = e^{-s\alpha(i \rightarrow j)} p(i \rightarrow j). \quad (7)$$

From the canonical path ensemble one can in certain cases recover the desired quantity $\rho(a, K)$ [20, 21]. The logarithm of the normalization factor in Eq. (5), $\lambda(s) \equiv K^{-1} \ln \sum_{\mathbf{x}} e^{-sA[\mathbf{x}]} P[\mathbf{x}]$, can be obtained by taking the logarithm of the principal eigenvalue of the generator \mathbf{W}_s of *s*-ensemble dynamics [11, 12]. Legendre transform of $\lambda(s)$ yields a rate function

$$I_s(a) = \max_s (-sa - \lambda(s)), \quad (8)$$

which is equal to the true rate function $I(a)$ when the latter is convex [11, 22, 23]. The *s*-ensemble has been used with great success, providing insight into the behavior of a large number of models [12], but it possesses two drawbacks. One is that \mathbf{W}_s is not a stochastic generator (its columns do not sum to unity, i.e. $\sum_j e^{-s\alpha(i \rightarrow j)} p(i \rightarrow j) \neq 1$), and so considerable effort and ingenuity is sometimes needed to identify and simulate the *s*-ensemble [5, 7, 8, 11–15]. More seriously, the calculation of probabilities is indirect, and fails when the rate function $I(a)$ of the original model is non-convex [11, 22, 23]. This is so in many physically interesting cases, where there exist phase transitions and coexistence of trajectories. Here (8) returns only the convex hull of the true rate function [11]: information about the trajectory ensemble is missing.

The method described here and in Ref. [16] draws inspiration from the *s*-ensemble method but calculates $\rho(a, K)$ directly, and so can recover non-convex rate functions. It is related to the method used in Ref. [24], generalized to Markov chains and applied to a probability-conserving model. Consider a stochastic reference model

whose transition probabilities $p_{\text{ref}}(i \rightarrow j)$ are normalized versions of those of the *s*-ensemble, $p_s(i \rightarrow j) / \sum_j p_s(i \rightarrow j)$ (the idea of using such a model is noted in Eq. (4.24) of [25]). The master equation of the reference model is then

$$\mathbf{P}_{\text{ref}}(k+1) = \mathbf{W}_{\text{ref}} \mathbf{P}_{\text{ref}}(k), \quad (9)$$

where element (j, i) of \mathbf{W}_{ref} is

$$p_{\text{ref}}(i \rightarrow j) = \frac{e^{-s\alpha(i \rightarrow j)} p(i \rightarrow j)}{\sum_j e^{-s\alpha(i \rightarrow j)} p(i \rightarrow j)}. \quad (10)$$

Note that the columns of \mathbf{W}_{ref} sum to unity – $\sum_j p_{\text{ref}}(i \rightarrow j) = 1$ – and so \mathbf{W}_{ref} is stochastic (probability-conserving). The idea of the scheme is then as follows. The reference model is not the *s*-ensemble, and is not intended to be. Simulations of the reference model (parameterized by *s*) produce values of *a* that for large *K* concentrate on a_s , where a_s minimizes the rate function $I_{\text{ref}}(a)$ of the reference model. Carrying out such simulations, and using the equations shown below, we calculate a piece $I(a_s)$ of the rate function of the original model. Repeating the process for a different value of *s* yields another piece of $I(a)$, and so on. We do not generate the microcanonical or canonical path ensembles. Instead, a simple modification of the original model (which admits all the processes of that model and no new ones) allows us to control which value of *a* is sampled, and from comparison of original and reference models we calculate the likelihood, Eq. (3), of generating those values of *a* using the original model.

To carry out this calculation note that Eq. (3) can be written

$$\rho(a, K) = \sum_{\mathbf{x}} P_{\text{ref}}[\mathbf{x}] w[\mathbf{x}] \delta(A[\mathbf{x}] - Ka), \quad (11)$$

where

$$P_{\text{ref}}[\mathbf{x}] = P_0(C_0) \prod_{k=0}^{K-1} p_{\text{ref}}(C_k \rightarrow C_{k+1}) \quad (12)$$

is the trajectory weight of the reference model, and $w[\mathbf{x}] \equiv P[\mathbf{x}] / P_{\text{ref}}[\mathbf{x}]$. It is convenient to write the transition probabilities of the original and reference models as $p(C \rightarrow C') = W(C \rightarrow C') / R(C)$ and $p_{\text{ref}}(C \rightarrow C') = W_{\text{ref}}(C \rightarrow C') / R_{\text{ref}}(C)$, where $W(C \rightarrow C')$ and

$$W_{\text{ref}}(C \rightarrow C') = e^{-s\alpha(C \rightarrow C')} W(C \rightarrow C') \quad (13)$$

are rates for the original and reference models, respectively, and $R(C) \equiv \sum_{C'} W(C \rightarrow C')$ and $R_{\text{ref}}(C) \equiv \sum_{C'} W_{\text{ref}}(C \rightarrow C')$ are the models’ escape rates from a given state.

We then have $w[\mathbf{x}] = \prod_{k=0}^{K-1} w_k$, where

$$w_k = e^{s\alpha(C_k \rightarrow C_{k+1})} \frac{R_{\text{ref}}(C_k)}{R(C_k)}. \quad (14)$$

It is also convenient to define the quantity

$$q[\mathbf{x}] \equiv \sum_{k=0}^{K-1} \ln \frac{R_{\text{ref}}(C_k)}{R(C_k)}, \quad (15)$$

using the pieces of $w[\mathbf{x}]$ not fixed by the delta-function path constraint.

The sum in Eq. (11) is the instruction to simulate the reference model, and take the arithmetic mean of values of w for trajectories that display activity a (the constraint imposed by the delta function). It is natural to consider the case $a = a_s$, where a_s is a value typical of the reference model (for which $I_{\text{ref}}(a_s) = 0$). If we generate \mathcal{N} trajectories of the reference model, and the trajectories labeled $i = 1, 2, \dots, \mathcal{M} \leq \mathcal{N}$ are typical in this way, then (11) reads

$$\rho(a_s, K) = \frac{\mathcal{M}}{\mathcal{N}} \cdot \frac{1}{\mathcal{M}} (w_1 + \dots + w_{\mathcal{M}}). \quad (16)$$

From (14) and (15) we have $w_\kappa = e^{sKa_s} e^{Kq_\kappa}$, where κ labels trajectories. Upon taking logarithms of (16) and assuming K large, so that $K^{-1} \ln(\mathcal{M}/\mathcal{N})$ is negligible, we get

$$I(a_s) = -sa_s - q_s - K^{-1} \ln \int dq P_s(q|a_s) e^{K\delta q}. \quad (17)$$

Here $I(a_s) = -K^{-1} \ln \rho(a_s, K)$ is the large-deviation rate function of the original model evaluated at a_s ; $q_s \equiv \int dq q P_s(q|a_s)$ is the mean value of $q[\mathbf{x}]$ for the ensemble of typical reference-model trajectories (those with $A[\mathbf{x}] = Ka_s$); $P_s(q|a_s)$ is the probability distribution of q for the ensemble of typical reference-model trajectories; and $\delta q \equiv q - q_s$.

By Jensen's inequality the first two terms in (17) provide a bound on the rate function $I_0(a_s) \geq I(a_s)$, where

$$I_0(a_s) = -sa_s - q_s. \quad (18)$$

The full rate function can be recovered by evaluation of the integral, which is demanding in general [24] (e.g. requiring a high-order cumulant expansion) but straightforward when $P_s(q|a_s)$ is Gaussian or close to it (this is not guaranteed in general, but is realized in several cases we have encountered [16]).

The computational cost $\mathcal{C}_{\text{ref}}^{(0)} K$ of generating a trajectory of length K of the reference model is in general greater than the cost $\mathcal{C}^{(0)} K$ of generating a trajectory of similar length using the original model, because to use

the reference model we must calculate the bias parameters $\alpha(C \rightarrow C')$. However, this cost is usually more than offset by our ability to direct the reference model to rarely-sampled values a_s of the dynamic parameter: the cost to sample the piece $I(a_s)$ of the rate function using the original model is $\mathcal{C} = \mathcal{C}^{(0)} K e^{KI(a_s)}$, while the cost using the reference model is $\mathcal{C}_{\text{ref}} = \mathcal{C}_{\text{ref}}^{(0)} K e^{KI_{\text{ref}}(a_s)} = \mathcal{C}_{\text{ref}}^{(0)} K$. Under most conditions we have encountered, the reference model is much more efficient than the original model. E.g. for the lattice model of Section III, $\alpha(C \rightarrow C')$ can be calculated via spatially local updates (updates involving only the neighborhood of the chosen lattice site), and the ratio $\mathcal{C}_{\text{ref}}^{(0)}/\mathcal{C}^{(0)}$ does not depend on the size of the system (it is of order 10). In this case the ratio $\mathcal{C}/\mathcal{C}_{\text{ref}} = (\mathcal{C}^{(0)}/\mathcal{C}_{\text{ref}}^{(0)}) e^{KI(a_s)}$ becomes large for any appreciable K (and nonzero $I(a_s)$), and so the reference-model method is more efficient than direct simulation. If for instance the computation of the bias $\alpha(C \rightarrow C')$ required updates of the whole lattice, and so $\mathcal{C}_{\text{ref}}^{(0)}$ grew with system size N , then there would exist an interval of K for which direct simulation may be more efficient.

In what follows we specialize the discussion to ergodic (irreducible, aperiodic) Markov chains possessing well-defined steady-state measures (probability of occupancy) $\pi(C)$ and $\pi_{\text{ref}}(C)$. These are obtained by requiring (2) and (9) to be stationary, and satisfy

$$\pi_{(\text{ref})}(C) = \sum_{C'} \pi_{(\text{ref})}(C') p_{(\text{ref})}(C' \rightarrow C). \quad (19)$$

The existence of such a measure is guaranteed when each microscopic process has a non-vanishing probability of reversal, which is true of the model of Ref. [17], but not true, for instance, of irreversible growth processes [16].

Eq. (17) can be written

$$I(a_s) = -K^{-1} \ln \langle \exp(\ln w_1 + \dots + \ln w_K) \rangle_{\text{ref}}, \quad (20)$$

where $\langle \cdot \rangle_{\text{ref}}$ denotes the average over the reference-model dynamics (for trajectories such that $a = a_s$). By Jensen's inequality we have Eq. (18), or

$$\begin{aligned} I_0(a_s) &= -K^{-1} \ln \exp(\langle \ln w_1 + \dots + \ln w_K \rangle_{\text{ref}}) \\ &= -\langle \ln w \rangle_{\text{ref}}. \end{aligned} \quad (21)$$

The average $\langle \ln w \rangle_{\text{ref}}$ can be calculated by keeping track of the terms $\ln w_i$ associated with each move in a reference-model simulation,

$$\theta(K) = K^{-1} \sum_{k=0}^{K-1} \ln w_k. \quad (22)$$

For large K , the series $\theta(K)$ will converge to θ_s , where

$$\theta_s = sa_s + q_s \quad (23)$$

$$= a \sum_C \pi_{\text{ref}}(C) \sum_{C'} p_{\text{ref}}(C \rightarrow C') \alpha(C \rightarrow C') + \sum_C \pi_{\text{ref}}(C) \ln \frac{R_{\text{ref}}(C)}{R(C)} \quad (24)$$

$$= \sum_C \pi_{\text{ref}}(C) \left\{ -s \frac{\partial}{\partial s} \ln R_{\text{ref}}(C) + \ln \frac{R_{\text{ref}}(C)}{R(C)} \right\}. \quad (25)$$

Note that for the choice $\alpha(C \rightarrow C') = B(C')$, made in Ref. [17], we have

$$a_s = \sum_C \pi_{\text{ref}}(C) \sum_{C'} p_{\text{ref}}(C \rightarrow C') B(C') \quad (26)$$

$$= \sum_{C'} B(C') \sum_C \pi_{\text{ref}}(C) p_{\text{ref}}(C \rightarrow C') \quad (27)$$

$$= \sum_C \pi_{\text{ref}}(C) B(C), \quad (28)$$

in which case

$$\theta_s = \sum_C \pi_{\text{ref}}(C) \left\{ sB(C) + \ln \frac{R_{\text{ref}}(C)}{R(C)} \right\}. \quad (29)$$

For simple models for which states C can be enumerated explicitly, the above expressions provide an analytic way of calculating the bound $I_0(a_s) = -\theta_s$. The forms (24) and (25) emphasize that only the typical occupancies $\pi_{\text{ref}}(C)$ need be calculated; the form (25) emphasizes in addition that the prefactors of these occupancies can be derived from the object $\ln \sum_{C'} e^{-s\alpha(C \rightarrow C')} W(C \rightarrow C')$.

The rate function calculated using the reference-model method is not constrained to be convex: Eq. (17), which is a direct evaluation of a probability, follows from simple arguments that do not depend on the form of the rate function in question. Convexity of the s -ensemble rate function, $I_s(a)$, is guaranteed by the structure of Eq. (8), which imposes a particular relationship between the two terms in that expression. Upon maximization, the quantity a in the first term of (8) is given by the derivative of the second term, $-\partial\lambda(s)/\partial s$ (where this derivative exists). The same relationship does not exist between the terms of Eq. (17), or between the terms of its bound, Eq. (18), $I_0(a_s) = -sa_s - q_s$. To see this, note from Eq. (25) that

$$a_s = - \sum_C \pi_{\text{ref}}(C) \partial_s \ln R_{\text{ref}}(C), \quad (30)$$

and

$$q_s = \sum_C \pi_{\text{ref}}(C) \ln[R_{\text{ref}}(C)/R(C)]. \quad (31)$$

Thus

$$-\partial_s q_s = a_s - \sum_C (\partial_s \pi_{\text{ref}}(C)) \ln[R_{\text{ref}}(C)/R(C)], \quad (32)$$

and so a_s and q_s are not related in the way they would be if they were the elements of a Legendre transform




FIG. 1. The coin-toss model represented as a Markov process, with transition probabilities as indicated.

(except in the special case of constant exit-rate ratio $R_{\text{ref}}(C)/R(C)$).

Thus the rate-function bound (18), which is simple to calculate, can be non-convex and so can give information about coexistence of trajectories (see e.g. Fig. 4 of Ref. [16]).

A. Example 1

To illustrate the reference-model method we start with a simple example in which the s -ensemble and reference-model method are equivalent. Consider the unbiased coin-toss model, represented in Fig. 1 as a two-state Markov chain. The microstates 0 and 1 correspond to a tail or a head, respectively, and the probability of generating either is 1/2. The master equation in the constant- K ensemble, Eq. (2), has generator (matrix of transition probabilities)

$$\mathbf{W} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}. \quad (33)$$

We define a dynamic observable A that increases by +1 upon any move into microstate 1, and by -1 upon any move into microstate 0. The observable $a = A/K$ is then the total number of heads minus the total number of tails, divided by the number of coin tosses, and its rate function is [26]

$$I_{\text{coin}}(a) = \frac{1-a}{2} \ln(1-a) + \frac{1+a}{2} \ln(1+a). \quad (34)$$

The s -ensemble can be used to derive this result [11]. The s -ensemble of the coin-toss model has generator

$$\mathbf{W}_s = \begin{pmatrix} \frac{e^s}{2} & \frac{e^s}{2} \\ \frac{e^{-s}}{2} & \frac{e^{-s}}{2} \end{pmatrix}, \quad (35)$$

whose principal eigenvalue is $\cosh s$. Then $\lambda(s) = \ln \cosh s$, and Eq. (8) yields the rate function

$$\begin{aligned} I_s(a) &= \max_s (-sa - \ln \cosh s) \\ &= a \tanh^{-1} a - \ln \cosh \tanh^{-1} a \\ &= \frac{1-a}{2} \ln(1-a) + \frac{1+a}{2} \ln(1+a), \end{aligned} \quad (36)$$

as expected.

Consider now the reference-model dynamics, which is generated by

$$\mathbf{W}_{\text{ref}} = \begin{pmatrix} \frac{e^s}{2 \cosh s} & \frac{e^s}{2 \cosh s} \\ \frac{e^{-s}}{2 \cosh s} & \frac{e^{-s}}{2 \cosh s} \end{pmatrix}. \quad (37)$$

This stochastic process can be simulated using standard methods [27], in order to calculate Eq. (17) numerically [16]. For this problem the evaluation can also be done analytically. The steady-state measure of the reference model is $\pi(0) = e^s/(2 \cosh s)$ and $\pi(1) = 1 - \pi(0)$, and the typical activity is $a_s = \sum_C \pi(C) \sum_{C'} p_{\text{ref}}(C \rightarrow C') \alpha(C \rightarrow C') = -\tanh s$ (thus a_s and s have the relationship that a and s possess in Eq. (36)). The exit-rate ratio is always $q_s = \ln(R_{\text{ref}}(C)/R(C)) = \cosh s$, and so Eq. (17) is

$$\begin{aligned} I(a_s) &= a_s \tanh^{-1} a_s - \ln \cosh \tanh^{-1} a_s \\ &= \frac{1-a_s}{2} \ln(1-a_s) + \frac{1+a_s}{2} \ln(1+a_s), \end{aligned} \quad (38)$$

giving us one point on the rate-function curve $I_{\text{coin}}(a)$. Repeating the method for a range of values of s (which here amounts simply to replacing $a_s \rightarrow a$) us to reconstruct the full rate-function curve.

B. Example 2

A second example illustrates a case in which the s -ensemble and reference-model ensemble are different (but both can be used to obtain $\rho(a, K)$). Consider the 4-state model shown in Fig. 2, with each rate set to unity (indicated are the transition probabilities, all $1/2$). Define the activity A as the number of configuration changes from state 1 to 2 or state 2 to 1 in the course of a simulation of K steps. By symmetry of the network the mean number of such configuration changes will be $A = K/4$, but some trajectories will exhibit more changes, and some less. We can use the s -ensemble and the reference-model method to calculate the logarithmic probability distribution $I(a) = -K^{-1} \ln \rho(a, K)$ associated with the quantity $a = A/K$.

The generator of the original model is

$$\mathbf{W} = \begin{pmatrix} 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \end{pmatrix}, \quad (39)$$

and that of the s -ensemble is

$$\mathbf{W}_s = \begin{pmatrix} 0 & \frac{e^{-s}}{2} & 0 & \frac{1}{2} \\ \frac{e^{-s}}{2} & 0 & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \end{pmatrix}. \quad (40)$$

The principal eigenvalue of \mathbf{W}_s is $e^{\lambda(s)} = \frac{1}{4} (e^{-s} \sqrt{-2e^s + 5e^{2s} + 1} + e^{-s} + 1)$, from which, using Eq. (8), $I(a)$ can be calculated numerically. The result is the green line in Fig. 2.

The reference-model generator is

$$\mathbf{W}_{\text{ref}} = \begin{pmatrix} 0 & \frac{e^{-s}}{1+e^{-s}} & 0 & \frac{1}{2} \\ \frac{e^{-s}}{1+e^{-s}} & 0 & \frac{1}{2} & 0 \\ 0 & \frac{1}{1+e^{-s}} & 0 & \frac{1}{2} \\ \frac{1}{1+e^{-s}} & 0 & \frac{1}{2} & 0 \end{pmatrix}, \quad (41)$$

which can be simulated using standard stochastic methods [27]. For a range of values of s (from -3 to 3 , in intervals of 0.25) we carried out a single simulation of the reference model of $K = 5 \times 10^8$ configuration changes. We computed the rate-function bound (18) by a) keeping track of explicit configuration changes, as in Eq. (22), and b) by evaluating Eq. (24). To evaluate the latter, note that the steady-state measure is

$$\pi_{\text{ref}}(1) = \pi_{\text{ref}}(2) = \frac{1 + e^{-s}}{2(3 + e^{-s})}. \quad (42)$$

The typical activity of the reference model is then $a_s = (\pi_{\text{ref}}(1) + \pi_{\text{ref}}(2)) \times e^{-s} / (1 + e^{-s})$, from which we obtain $s = \ln[(1/a_s - 1)/3]$. The relationship between a_s and s is different to the relationship $a = -\partial_s \lambda(s)$ obtained from the s -ensemble: here the two methods produce different typical activity for given s .

From Eq. (24) we have

$$I_0(a) = -\theta_s(a) = -a \ln \frac{1-a}{3a} - \frac{1+2a}{3} \ln \frac{1+2a}{2-2a}. \quad (43)$$

As shown in Fig. 2, both ways of calculating $I_0(a)$, analytically using Eq. (43) and numerically using (22), the blue and red lines in the figure, give the same result.

The full rate function can be estimated by evaluating numerically the integral in Eq. (17). To do so we took each trajectory of 5×10^8 steps and split it into 5×10^3 pieces of length $K' = 10^5$. Of that set we retained only those pieces with extensive order parameter $A' = K'a_s$ (it is necessary to explicitly enforce the delta-function constraint of Eq. (11) for short trajectories). The resulting distribution of q , $P(q|a_s)$, is Gaussian with variance σ_s^2 . In this case Eq. (17) reads

$$I(a_s) = -sa_s - q_s - K\sigma_s^2/2. \quad (44)$$

Evaluation of Eq. (44) gives the dotted black line in Fig. 2, consistent with the result calculated via the s -ensemble.




FIG. 2. For the 4-state model shown we compute the large-deviation rate function $I(a)$ associated with the order parameter $a = A/K$, where A is incremented by 1 for every transition along the blue arrows in the course of a trajectory of K steps. The green line is the result obtained using the s -ensemble. The blue and red lines on the right are upper bounds (18) on the rate function computed using the reference-model method. We keep track of explicit configuration changes during numerical simulations (red line, Eq. (22)) or evaluate analytically Eq. (24) = Eq. (43) (blue line). The black dashed line is obtained by evaluation of Eq. (44), which accounts for fluctuations of the weight for typical reference-model trajectories.

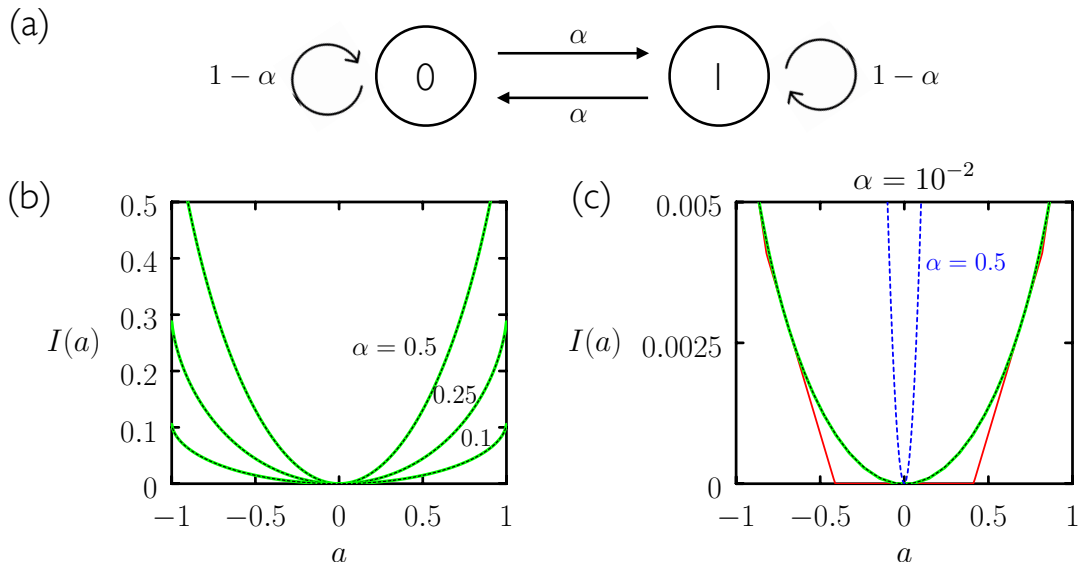


FIG. 3. (a) Minimal model of a one-dimensional active particle [11]. (b) For the values of α shown, rate functions obtained from the s -ensemble (green) and reference-model method (black dashed lines) agree. (c) For small values of α the rate function is very broad (red, green, and black lines, $\alpha = 10^{-2}$; the blue line shows the case $\alpha = 1/2$ for comparison). Here the rate function can be obtained from the s -ensemble via Legendre transform (green line) only if a very fine numerical grid of s is used; if not, the rate function displays spurious flat portions (red line). The reference-model method (black line) evaluates $I(a)$ directly, and does not require closely-spaced values of s .

C. Example 3

A third example illustrates some numerical difficulties associated with obtaining a rate function via Legendre transform, even when the rate function is convex. Consider the two-state Markov chain of example IV.4 of [11] (see also [28]), shown in Fig. 3(a). We shall construct a dynamic observable A by associating a value of $+1$ with

any move into state 1, and -1 with any move into state 0. In this case we can consider the model to be a minimal representation of a one-dimensional active particle (see e.g. [29, 30]), with 0 being a left-pointing, left-moving state, and 1 being a right-pointing, right-moving state. At each step the particle flips direction with probability α (and takes a step in the new direction), or else retains its current orientation and takes a step in the corresponding

direction. $a = A/K$ is then the position of the particle after K steps.

For $\alpha = 1/2$ we have the random walker (or unbiased coin-toss) of Section II A. For $\alpha = 0$ the two states cannot interconvert (the Markov chain is reducible), and there exist two rate functions (which vanish for $a = 1$ or $a = -1$, and are infinite otherwise). For $0 < \alpha < 1/2$ we have an ‘active walker’ that moves persistently right or left between flips. In physical terms, fluctuations of A/K for this model are related to fluctuations of magnetization M/N in the one-dimensional Ising model in zero magnetic field, where N is system size. Flips between runs $\cdots \text{LLLLRRRRLLLL} \cdots$ occur independently with probability α , much as domain walls between Ising spins $\cdots - - - + + + - - - \cdots$ occur independently with probability $(1 + e^{\beta J})^{-1}$, where J is the Ising coupling. For nonzero α (finite J) we expect domain lengths to be exponentially distributed with mean α^{-1} or $1 + e^{\beta J}$; for there to be no bias between L and R or between $-$ and $+$; and for the mean values of A/K and M/N in the limits of large K and large N to be zero. Thus the rate function $I(a)$ should be convex, symmetric, and have a (unique) minimum at zero.

We first calculate the rate function using the s -ensemble, following Ref. [11]. The generator

$$\mathbf{W} = \begin{pmatrix} 1 - \alpha & \alpha \\ \alpha & 1 - \alpha \end{pmatrix} \quad (45)$$

becomes, in the s -ensemble,

$$\mathbf{W}_s = \begin{pmatrix} (1 - \alpha)e^s & \alpha e^s \\ \alpha e^{-s} & (1 - \alpha)e^{-s} \end{pmatrix}. \quad (46)$$

From this we extract $\lambda(s)$ in the usual way [11], and calculate $I(a)$ via Legendre transform. Some results for three values of α are shown in Fig. 3(b) (green lines).

We next calculate $I(a)$ using the reference-model method. Its generator is

$$\mathbf{W}_{\text{ref}} = \begin{pmatrix} \frac{(1 - \alpha)e^s}{(1 - \alpha)e^s + \alpha e^{-s}} & \frac{\alpha e^s}{(1 - \alpha)e^{-s} + \alpha e^s} \\ \frac{\alpha e^{-s}}{(1 - \alpha)e^s + \alpha e^{-s}} & \frac{(1 - \alpha)e^{-s}}{(1 - \alpha)e^{-s} + \alpha e^s} \end{pmatrix}. \quad (47)$$

Calculation of $I(a)$ can be done analytically, because the Jensen bound (18) is exact for this model. This is so because each value of activity A is associated with a unique path weight: every move into state 1 (and only state 1) generates activity $+1$, and every move out of the state generates a piece $\ln(R_{\text{ref}}(1)/R(1))$ of the path weight. Similarly for state 0. By conservation of probability, every move into a state must be balanced by a move out of the state, and so the path weight takes a unique value for each value of A . Solving as before for the stationary measure $\pi_{\text{ref}}(C)$, we evaluate the terms in Eq. (24) and write down the rate function in the parametric form $(a_s(s), I(s))$; we get

$$a_s(s) = \frac{(\alpha - 1) \sinh(2s)}{\alpha - (\alpha - 1) \cosh(2s)} \quad (48)$$

and

$$\begin{aligned} I(s) = & \frac{(\alpha - 1)(e^{4s} - 1)s + (\alpha(e^{2s} - 1) + 1) \ln((2\alpha - 1) \sinh s + \cosh s)}{\alpha(e^{2s} - 1)^2 - e^{4s} - 1} \\ & + \frac{e^{3s}(-2\alpha \sinh s + \sinh s + \cosh s) \ln(-2\alpha \sinh s + \sinh s + \cosh s)}{\alpha(e^{2s} - 1)^2 - e^{4s} - 1}. \end{aligned} \quad (49)$$

In Fig. 3(b) we show $I(a)$ obtained from the reference-model method for three values of α (black dotted lines). These results agree with those obtained using the s -ensemble.

For smaller values of α the Legendre transform encounters numerical problems. In Fig. 3(c) we consider the case $\alpha = 10^{-2}$ (the blue line shows the case $\alpha = 1/2$ for comparison). The rate function, calculated via the reference-model method, is shown in black. The same result can be obtained using the s -ensemble via numerical Legendre transform, but only if many closely-spaced values of s are used to calculate Eq. (8): the associated scaled cumulant-generating function is sharply kinked, and its numerical differentiation requires fine intervals of s . In Fig. 1(c) the green line was calculated by numeri-

cal Legendre transform using intervals of s of 5×10^{-4} , and this agrees with the reference-model result. The red line was calculated by Legendre transform using coarser intervals of spacing 5×10^{-3} , and the result has spurious flat portions. It is clear that this result is unphysical: the mean value of a of the original model is 0, and so the true rate function must have a unique minimum at zero. The smaller the value of α , the finer the mesh required by the Legendre transform.

In this case the rate function is ‘simple’, i.e. it is convex and quadratic (for α small we can expand (48) and (49) about $s = 0$ to obtain $I(a) \approx \alpha a^2/2$) but its calculation via Legendre transform requires fine sampling of values of s . This is so for geometrical reasons. In the s -ensemble the calculated quantity $\lambda(s)$ is the Legendre

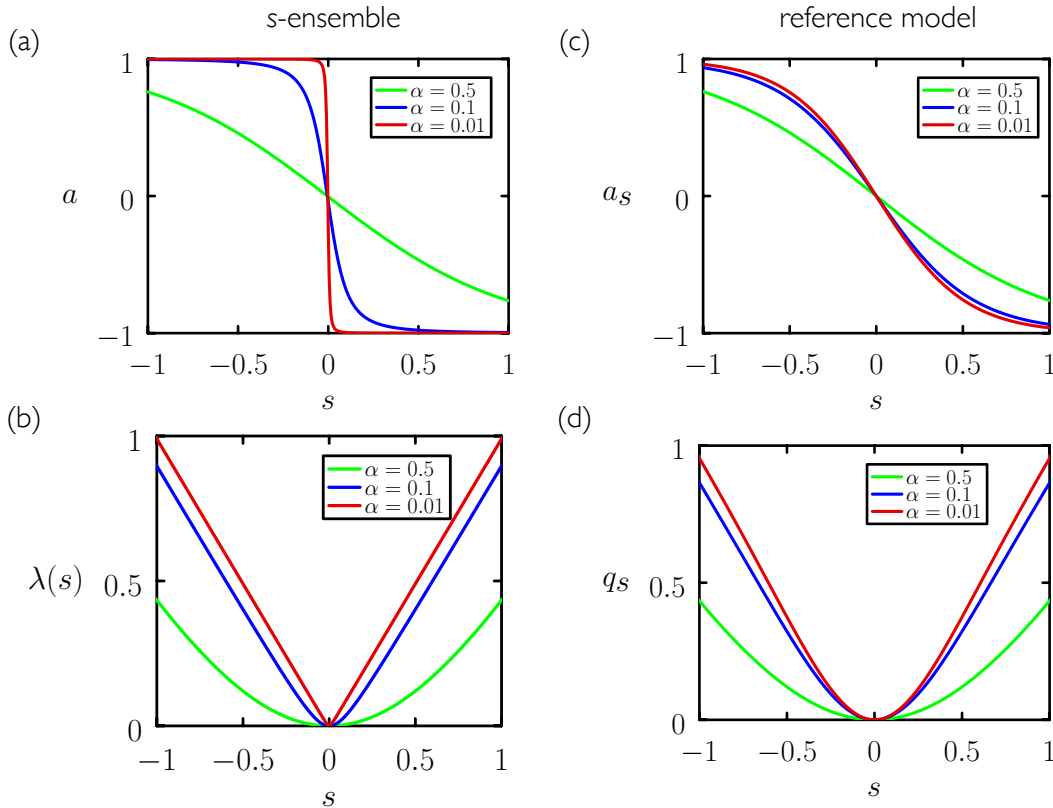


FIG. 4. The activity $a = -\partial_s \lambda(s)$ (panel (a)) and scaled cumulant-generating function $\lambda(s)$ (b) in the s -ensemble are constrained to vary sharply whenever the rate function $I(a)$ of the original model (shown in Fig. 3) has small gradient. In this case this is so near $a = 0$ as α becomes small. As a result, fine sampling near $s = 0$ is required to allow recovery of $I(a)$ by Legendre transform of $\lambda(s)$; see Fig. 3(b). By contrast, in the reference-model method the quantities a_s (c) and q_s (d) are not constrained to vary sharply when $I(a)$ has small gradient. Their combination $-sa_s - q_s$ yields the piece $I(a_s)$ of the rate function $I(a)$.

transform of the rate function $I(a)$. As a result (see Fig. 4 of [11]), $s(a)$ is given locally by (minus) the slope of the rate function $I(a)$. Thus a varies sharply with s whenever the rate function has small curvature. In Fig. 4(a) we show $a = -\partial_s \lambda(s)$ in the s -ensemble, for the model of Fig. 3; a changes sharply near $s = 0$ as α becomes small (here $a \approx -s/\alpha$ near $s = 0$). The corresponding scaled cumulant-generating function (b) becomes sharply kinked, and fine sampling near $s = 0$ is needed to convert $\lambda(s)$ into $I(a)$.

By contrast, the key quantities a_s and q_s of the reference-model method, panels (c) and (d) of Fig. 4, are not constrained to change sharply where original model's rate function $I(a)$ is slowly varying. The reference-model rate functions shown in Fig. 3 are evaluated directly, by taking the combination $-sa_s - q_s$. Each point on the function can be evaluated independently, and the grid of s can be as coarse or as fine as desired. The evaluation works when the rate function has very small curvature or is non-convex (see e.g. Fig. 4 of [16]). In the latter case the rate function cannot be recovered by Legendre transform [11].

In physical terms the results of this section relate to the probability distribution $P(a)$ of the position a

of an active particle, with rotation rate α , in discrete space. This distribution adopts a large-deviation form $P(a) \sim e^{-KI(a)}$ in the limit of a large number of steps K , with $I(a)$ given by Equations (48) and (49), shown in Fig. 3. For small α , fluctuations of a become large in the sense of having a large variance, but are to leading order Gaussian: $I(a) \approx \alpha a^2/2$, for a, α small. This distribution can be calculated using the s -ensemble, whose key quantities $\lambda(s)$ and $a(s)$ (see Fig. 4(a,b)) show a sharp jump when the rate function has small gradient. This is so for the geometrical reasons shown in Fig. 4 of [11]. The distribution $P(a)$ can also be calculated using the reference-model method, whose key quantities (see Fig. 4(c,d)) are not required to change sharply. Indeed, here the original model is similar to the 1D Ising model at nonzero temperature, for the reasons discussed above, and the set of reference models is similar to the 1D Ising model in a field $h \propto s$. Neither shows a sharp transition.

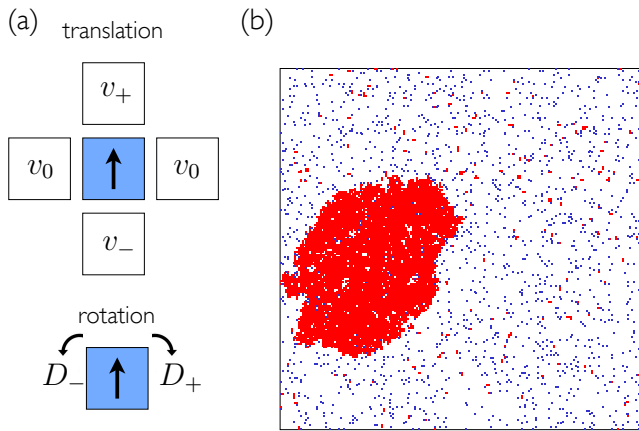


FIG. 5. (a) Rates of motion for a model of lattice-based active particles [17]. (b) A typical phase-separated configuration of collections of volume-excluding particles of this kind, produced using a continuous-time Monte Carlo algorithm [27], starting from well-mixed conditions. The packing fraction is $\phi = 1/5$; the rates are $v_+ = 16$, $v_- = v_0 = 1$, $D_+ = D_- = 1/10$. The lattice size is 200^2 , and periodic boundaries are applied in each direction. Particles are colored red if they point toward a nearest-neighbor particle, and blue otherwise.

III. SAMPLING NONEQUILIBRIUM TRAJECTORIES OF A LATTICE MODEL OF ACTIVE MATTER.

For models with larger state space, matrix diagonalization becomes impractical, and other methods are required to determine the s -ensemble [5, 7, 8, 11, 13–15]. The reference-model method is implemented in the same way, however, as we now illustrate using the lattice model of active particles introduced in Ref. [17].

Active particles propel themselves by consuming energy, and move by diffusion and by drift in a direction that fluctuates [31–46]. Active particles form clusters and undergo phase separation even in the absence of interparticle attractions [31, 37, 39, 44, 47, 48]. This phenomenon, called motility-induced phase separation, can be reproduced on the computer using simple model particles [30, 39, 47, 48] (motility-induced phase separation also occurs in the presence of velocity alignment [36, 49–51]). Associated with this phenomenon are various kinds of large fluctuations [34, 47, 52].

With an eye to studying such fluctuations in detail we introduced in Ref. [17] the following lattice model of active matter. Consider particles that move in isolation with the rates shown in Fig. 5(a). Particles possess both positions and orientations, and move on a two-dimensional square lattice. An isolated particle moves forward with rate v_+ , backward with rate v_- , to either side with rate v_0 , and rotates $\pi/2$ with rate D_+ or D_- (here we set $D_+ = D_- \equiv D_{\text{rot}}$). An on-lattice particle of this nature moves in a manner similar to an off-lattice Brownian particle [17]. As expected from the behavior

of off-lattice particles [30, 39, 47, 48], on-lattice active particles undergo motility-induced phase separation: see Fig. 5(b) and Fig. 6. We used N hard particles on a periodically-replicated lattice of size L^2 (for a packing fraction $\phi = N/L^2$), starting from disordered configurations, and propagated dynamics using a continuous-time Monte Carlo algorithm [27]. At each step of the simulation the system moves from the current configuration C to a new configuration C' with probability

$$p(C \rightarrow C') = \frac{W(C \rightarrow C')}{\sum_{C'} W(C \rightarrow C')}, \quad (50)$$

where the rates W for each possible process are taken from Fig. 5. Typical trajectories of this model lead to phase separation at a density-dependent value of the Péclet number [17, 39].

To quantify the rare behavior of this model we chose a dynamic order parameter a motivated by studies of glasses. There, authors sometimes choose to count the number of events that occur in a particular time, with the mean time taken to leave configuration C being $1/R(C)$. This choice allows the identification of phase transitions out of equilibrium [12, 53, 54]. Here we choose to measure the values of the escape rates $R(C)$ of states explored by a fixed number of configuration changes. With this choice the bias $\alpha(C \rightarrow C') = B(C')$, where $B(C') = R(C')$ is the escape rate from the state to which the model is moving, and the transition probabilities of the reference model are

$$p_{\text{ref}}(C \rightarrow C') = \frac{e^{-s[B(C') - B(C)]} W(C \rightarrow C')}{\sum_{C'} e^{-s[B(C') - B(C)]} W(C \rightarrow C')}. \quad (51)$$

The factors $e^{sB(C)}$ in numerator and denominator have been introduced so that numbers appearing in the arguments of exponentials are not too large.

The coding overhead for the reference model is only slightly greater than that of the original model. One needs to keep track of the escape rate $R(C')$ of each configuration C' accessible from the current configuration (rather than just the escape rate of the current configuration), but this can be done efficiently, in the course of the simulation, by noting which particles can touch the moving particle, before and after its move, one move into the future.

The order parameter against which dynamics is conditioned is $a = A/K = K^{-1} \sum_{k=0}^{K-1} B(C_{k+1})$, the mean relaxation rate of configurations comprising the trajectory, which we report in the form of a scaled ‘activity’, $\nu \equiv a/(\Sigma N)$. Here N is the number of particles and $\Sigma \equiv v_+ + 2v_0 + v_- + 2D_{\text{rot}}$ is the sum of rates of processes of isolated particles. We used Eq. (22) to calculate the rate-function bound, which in this case reads

$$I_0(a_s) = -sK^{-1} \sum_{k=0}^{K-1} [B(C_{k+1}) - B(C_k)] - K^{-1} \sum_{k=0}^{K-1} \ln \frac{\sum_{C'} e^{-s[B(C') - B(C_k)]} W(C_k \rightarrow C')}{\sum_{C'} W(C_k \rightarrow C')}. \quad (52)$$




FIG. 6. Active lattice-based particles undergo phase separation at a density-dependent value of their motility. Rows are labeled by the packing fraction, ϕ , and columns are labeled by the rate of forward motion, v_+ . The other rates are as Fig. 5, giving a Péclet number $Pe = 5(v_+ - 1)$ [17]. The lattice size is 200^2 .

The first term on the right-hand side of (52) becomes negligible for large K . We evaluated this expression for trajectories of about $K = 10^8$ steps (the bound looks no different trajectories of $K \sim 10^7$ steps are used).

Results derived from these simulations are shown in Fig. 7, for packing fraction $\phi = 1/5$. Here the rate-function bound (which is not convex) shows that the typical activity of the model (where $I_0(\nu) = 0$) shifts from large ν to small ν as the drift rate v_+ of the active particles increases, because typical trajectories exhibit clustering (see Fig. 6) and (hence) reduced dynamism. In addition, trajectories show an active-to-less active crossover as a function of the bias parameter s . For model parameters far from the ‘typical’ nonequilibrium phase transition (which occurs near $v_+ \approx 7$ for $\phi = 1/5$) this crossover is seen in the ensemble of rare trajectories (for which $s \neq 0$), while for model parameters near the ‘typical’ transition the crossover occurs in the ensemble of typical trajectories, for which $s \approx 0$. This behavior resembles transitions observed in models of glasses [53] or proteins [55] using the s -ensemble.

IV. CONCLUSIONS

We have described a method to sample the rare fluctuations of discrete-time Markov chains, and have applied it to the lattice-based model of active matter introduced in [17]. The method is inspired by the s -ensemble method but is distinct: we use a probability-conserving reference model to calculate directly the likelihood that a model displays a particular value of a dynamic order parameter a . As a result, the method can be used to calculate

non-convex rate functions, so giving information about trajectory coexistence near nonequilibrium phase transitions. The method is straightforward to implement for models for which continuous-time Monte Carlo dynamics can be carried out. Take the rates of the original model $W(C \rightarrow C')$, change them to $e^{-s\alpha(C \rightarrow C')}W(C \rightarrow C')$, run simulations and evaluate the quantities needed for equation (17) (repeat for a range of values of s). In general, as noted in Ref. [24], high-precision evaluation of integrals of the form appearing in Eq. (17) is demanding. For several models we have studied, however, calculation of the variance of q appears to give a good approximation of the rate function $I(a)$ (because fluctuations of q are often approximately Gaussian), while calculation of the mean of q – which is straightforward – gives via Eq. (18) a physically meaningful, and potentially non-convex, rate-function bound. Thus the method provides a simple way of obtaining non-trivial information about a model and its phase transitions away from equilibrium, and so provides a complement to existing trajectory-sampling methods [5, 6, 10–13, 53, 56–59].

We end with a note about time. The reference-model method operates within the constant-event-number ensemble (the discrete-time Markov chain), because placing no restriction upon jump times makes the sampling procedure easier (see Appendix B of [16]). However, the procedure used to move through configuration space is the same as that used in a continuous-time Monte Carlo scheme (except that no random number is drawn to compute the time of escape from a given state), and all temporal information can be recovered after the fact. The escape time t_k from configuration C_k is an exponentially distributed random variable with mean $1/R(C_k)$, and

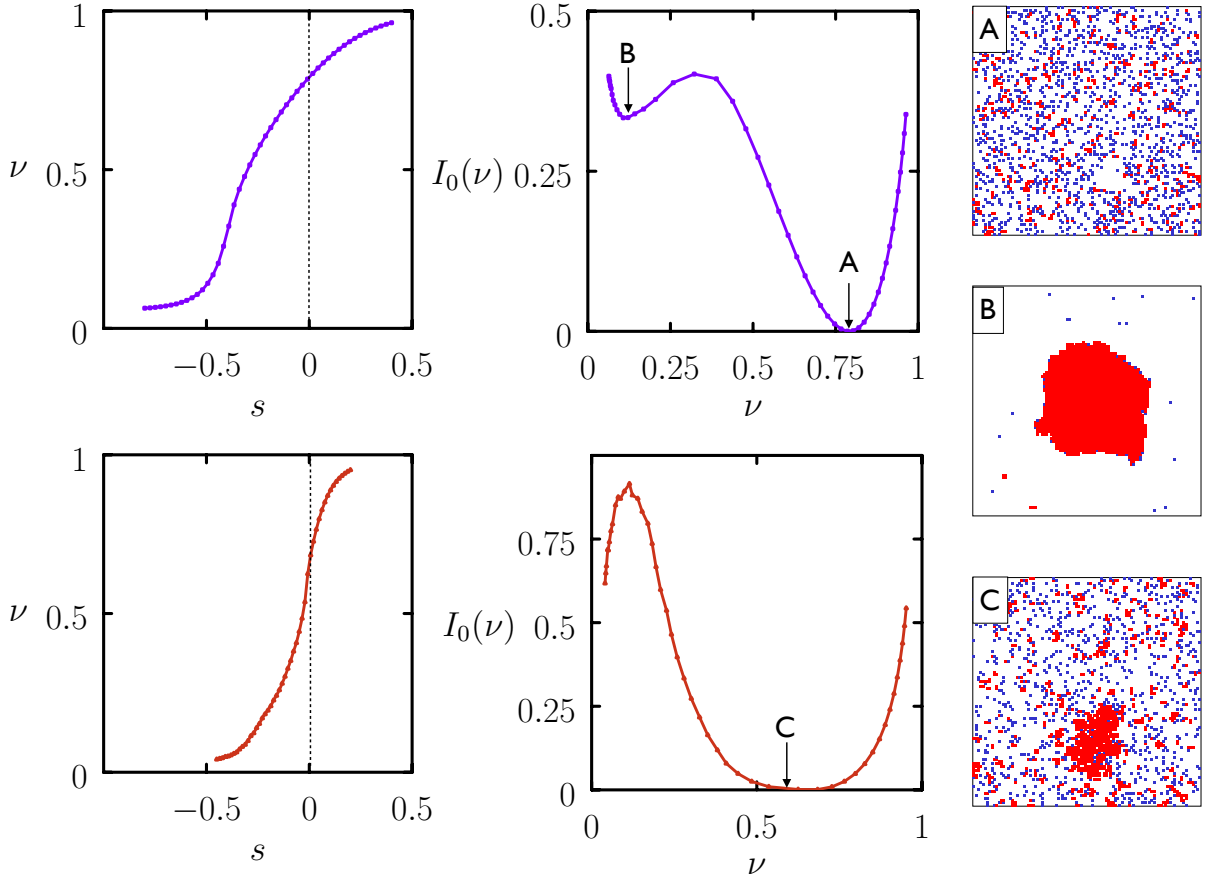


FIG. 7. Activity ν as a function of s (left graphs) and rate-function bound $I_0(\nu)$ as a function of ν (right graphs) from the reference-model method applied to the active lattice model. Packing fraction is $\phi = 1/5$ and drift rate $v_+ = 2$ (blue lines, top) or $v_+ = 7$ (red lines, bottom). Lattice size is 100^2 . Trajectories prepared at $s = 0$ are typical of the *original* model, while trajectories prepared at $s \neq 0$ are rare in the original model; how rare can be understood from the right-hand graphs. For $v_+ = 2$ we are far from the ‘typical’ phase transition (see Fig. 6), and the crossover from active to inactive trajectories occurs for $s \neq 0$. For $v_+ = 7$, close to the ‘typical’ transition, the crossover occurs near $s = 0$. Shown right are typical and rare configurations of the model.

so the total time of a trajectory, $T = \sum_{k=0}^{K-1} t_k$, is a stochastic variable with a known distribution $P(T)$ [60]. This distribution reduces to the Erlang distribution if all $R(C_k)$ are identical, and to the hypoexponential distribution if all $R(C_k)$ are distinct [61]. Thus the complete distribution of trajectory times, for fixed number of events and specified activity a , can be recovered from the constant-event-number ensemble without additional

simulation.

ACKNOWLEDGMENTS

I thank Hugo Touchette, Juan Garrahan, and Katherine Klymko for discussions. This work was performed 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.

[1] D. Frenkel and B. Smit, *Understanding molecular simulation: from algorithms to applications*, Vol. 1 (Academic press, 2001).

[2] G. M. Torrie and J. P. Valleau, *Journal of Computational Physics* **23**, 187 (1977).

[3] R. J. Allen, C. Valeriani, and P. R. ten Wolde, *Journal*

- of physics: Condensed matter **21**, 463102 (2009).
- [4] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, *Annual review of physical chemistry* **53**, 291 (2002).
 - [5] C. Giardinà, J. Kurchan, V. Lecomte, and J. Tailleur, *Journal of statistical physics* **145**, 787 (2011).
 - [6] C. Giardinà, J. Kurchan, and L. Peliti, *Physical Review Letters* **96**, 120603 (2006).
 - [7] T. Nemoto, F. Bouchet, R. L. Jack, and V. Lecomte, *Physical Review E* **93**, 062123 (2016).
 - [8] V. Lecomte and J. Tailleur, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P03004 (2007).
 - [9] T. Nemoto, in *Phenomenological Structure for the Large Deviation Principle in Time-Series Statistics* (Springer, 2016) pp. 17–39.
 - [10] J. A. Bucklew, *Large deviation techniques in decision, simulation, and estimation* (Wiley New York, 1990).
 - [11] H. Touchette, *Physics Reports* **478**, 1 (2009).
 - [12] 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).
 - [13] C. Maes and K. Netočný, *EPL (Europhysics Letters)* **82**, 30003 (2008).
 - [14] R. L. Jack and P. Sollich, *The European Physical Journal Special Topics* **224**, 2351 (2015).
 - [15] T. Nemoto and S.-i. Sasa, *Physical Review Letters* **112**, 090602 (2014).
 - [16] K. Klymko, P. L. Geissler, J. P. Garrahan, and S. Whitlam, *arXiv preprint arXiv:1707.00767* (2017).
 - [17] S. Whitlam, K. Klymko, and D. Mandal, *arXiv preprint arXiv:1709.03951* (2017).
 - [18] R. Evans, *Physical Review Letters* **92**, 150601 (2004).
 - [19] R. Evans, *Journal of Physics A: Mathematical and General* **38**, 293 (2004).
 - [20] R. Chetrite and H. Touchette, *Physical Review Letters* **111**, 120601 (2013).
 - [21] R. Chetrite and H. Touchette, *Annales Henri Poincaré, Ann. Henri Poincaré*, 1 (2014).
 - [22] I. H. Dinwoodie, S. L. Zabell, *et al.*, *The Annals of Probability* **20**, 1147 (1992).
 - [23] I. H. Dinwoodie, *Annals of probability* **21**, 216 (1993).
 - [24] C. M. Rohwer, F. Angeletti, and H. Touchette, *Physical Review E* **92**, 052104 (2015).
 - [25] H. Touchette, *arXiv preprint arXiv:1106.4146* (2011).
 - [26] J. T. Lewis and R. Russell, *Dublin Institute for Advanced Studies*, 1 (1997).
 - [27] D. Gillespie, *Journal of Computational Physics* **22**, 403 (2005).
 - [28] Y. Oono, *Progress of Theoretical Physics Supplement* **99**, 165 (1989).
 - [29] P. Pietzonka, K. Kleinbeck, and U. Seifert, *New Journal of Physics* **18**, 052001 (2016).
 - [30] A. Thompson, J. Tailleur, M. Cates, and R. Blythe, *Journal of Statistical Mechanics: Theory and Experiment* **2011**, P02029 (2011).
 - [31] J. Tailleur and M. Cates, *Physical Review Letters* **100**, 218103 (2008).
 - [32] S. Weitz, A. Deutsch, and F. Peruani, *Physical Review E* **92**, 012322 (2015).
 - [33] A. Baskaran and M. C. Marchetti, *Proceedings of the National Academy of Sciences* **106**, 15567 (2009).
 - [34] F. Peruani, J. Starrau, V. Jakovljevic, L. Søgaard-Andersen, A. Deutsch, and M. Bär, *Physical Review Letters* **108**, 098102 (2012).
 - [35] J. Barré, R. Chétrite, M. Muratori, and F. Peruani, *Journal of Statistical Physics* **158**, 589 (2015).
 - [36] F. Peruani, T. Klaus, A. Deutsch, and A. Voss-Boehme, *Physical Review Letters* **106**, 128101 (2011).
 - [37] S. Ramaswamy, *The Annual Review of Condensed Matter Physics* **1**, 323 (2010).
 - [38] P. Romanczuk, M. Bar, W. Ebeling, B. Lindner, and L. Schimansky-Geier, *European Physical Journal Special Topics* **202**, 1 (2012).
 - [39] G. S. Redner, M. F. Hagan, and A. Baskaran, *Physical Review Letters* **110**, 055701 (2013).
 - [40] M. C. Marchetti, J. Joanny, S. Ramaswamy, T. Liverpool, J. Prost, M. Rao, and R. A. Simha, *Reviews of Modern Physics* **85**, 1143 (2013).
 - [41] I. Buttinoni, J. Bialké, F. Kümmel, H. Löwen, C. Bechinger, and T. Speck, *Physical Review Letters* **110**, 238301 (2013).
 - [42] J. Stenhammar, A. Tiribocchi, R. J. Allen, D. Marenduzzo, and M. E. Cates, *Physical Review Letters* **111**, 145702 (2013).
 - [43] J. Yeomans, D. Pushkin, and H. Shum, *European Physical Journal Special Topics* **223**, 1771 (2014).
 - [44] M. E. Cates and J. Tailleur, *Annu. Rev. Condens. Matter Phys.* **6**, 219 (2015).
 - [45] A. M. Menzel, *Physics Reports* **554**, 1 (2015).
 - [46] C. Bechinger, R. D. Leonardo, H. Lowen, C. Reichhardt, G. Volpe, and G. Volpe, *Reviews of Modern Physics* **88**, 045006 (2016).
 - [47] Y. Fily and M. C. Marchetti, *Physical Review Letters* **108**, 235702 (2012).
 - [48] T. Speck, J. Bialké, A. M. Menzel, and H. Löwen, *Physical Review Letters* **112**, 218304 (2014).
 - [49] A. P. Solon and J. Tailleur, *Physical Review E* **92**, 042119 (2015).
 - [50] A. Solon and J. Tailleur, *Physical Review Letters* **111**, 078101 (2013).
 - [51] K. R. Pilkieicz and J. D. Eaves, *Phys. Rev. E* **89**, 012718 (2014).
 - [52] F. Peruani, L. Schimansky-Geier, and M. Bär, *The European Physical Journal Special Topics* **191**, 173 (2010).
 - [53] L. O. Hedges, R. L. Jack, J. P. Garrahan, and D. Chandler, *Science* **323**, 1309 (2009).
 - [54] A. A. Budini, R. M. Turner, and J. P. Garrahan, *Journal of Statistical Mechanics: Theory and Experiment* **2014**, P03012 (2014).
 - [55] A. S. Mey, P. L. Geissler, and J. P. Garrahan, *Physical Review E* **89**, 032109 (2014).
 - [56] D. Ruelle, *Thermodynamic formalism: the mathematical structure of equilibrium statistical mechanics* (Cambridge University Press, 2004).
 - [57] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, *Physical Review Letters* **98**, 195702 (2007).
 - [58] V. Lecomte, C. Appert-Rolland, and F. van Wijland, *J. Stat. Phys.* **127**, 51 (2007).
 - [59] R. Chetrite and H. Touchette, *Journal of Statistical Mechanics: Theory and Experiment* **2015**, P12001 (2015).
 - [60] H. Jasiulewicz and W. Kordecki, *Demonstratio Mathematica. Warsaw Technical University Institute of Mathematics* **36**, 231 (2003).
 - [61] P. Buchholz, J. Kriege, and I. Felko, *Input modeling with phase-type distributions and Markov models: theory and applications* (Springer, 2014).