Importance sampling large deviations in nonequilibrium steady states. I

Ushnish Ray, Garnet Kin-Lic Chan, and David T. Limmer

Citation: *The Journal of Chemical Physics* **148**, 124120 (2018); doi: 10.1063/1.5003151

View online: <https://doi.org/10.1063/1.5003151>

View Table of Contents: <http://aip.scitation.org/toc/jcp/148/12>

Published by the American Institute of Physics

Importance sampling large deviations in nonequilibrium steady states. I

Ushnish Ray,^{1,a)} Garnet Kin-Lic Chan,^{1,b)} and David T. Limmer^{2,3,4,c)}

¹*Division of Chemistry and Chemical Engineering, California Institute of Technology, Pasadena, California 91125, USA*

²*Department of Chemistry, University of California, Berkeley, California 94609, USA*

³*Kavli Energy NanoScience Institute, Berkeley, California 94609, USA*

⁴*Materials Science Division, Lawrence Berkeley National Laboratory, Berkeley, California 94609, USA*

(Received 3 September 2017; accepted 5 March 2018; published online 29 March 2018)

Large deviation functions contain information on the stability and response of systems driven into nonequilibrium steady states and in such a way are similar to free energies for systems at equilibrium. As with equilibrium free energies, evaluating large deviation functions numerically for all but the simplest systems is difficult because by construction they depend on exponentially rare events. In this first paper of a series, we evaluate different trajectory-based sampling methods capable of computing large deviation functions of time integrated observables within nonequilibrium steady states. We illustrate some convergence criteria and best practices using a number of different models, including a biased Brownian walker, a driven lattice gas, and a model of self-assembly. We show how two popular methods for sampling trajectory ensembles, transition path sampling and diffusion Monte Carlo, suffer from exponentially diverging correlations in trajectory space as a function of the bias parameter when estimating large deviation functions. Improving the efficiencies of these algorithms requires introducing guiding functions for the trajectories. *Published by AIP Publishing.*
<https://doi.org/10.1063/1.5003151>

I. INTRODUCTION

Free energy calculations play a central role in molecular simulations of chemical, biological, and material systems. As the governing function for stability, free energies are a principle quantity of statistical physics. With the exception of special cases, the evaluation of free energies requires numerical computation. The development and ease of implementation of methods such as umbrella sampling, metadynamics, and Wang-Landau sampling^{1–3} has made the computation of equilibrium free energies textbook material.⁴ With these and related techniques, free energy calculations of realistic models of matter are routinely used to compute phase diagrams, to infer thermodynamic response, and to estimate rates of rare events.^{5–7} However, such calculations are traditionally limited to instances where systems are at thermodynamic equilibrium, as most algorithms are typically formulated with an implicit reliance on an underlying Boltzmann distribution of configurations. Large deviation theory provides a theoretical framework for extending free energy calculations to systems away from equilibrium, with large deviation functions acting as the analog of free energies for such systems.⁸ Though not identical to a thermodynamic potential, large deviation functions nevertheless encode stability and often response of nonequilibrium systems. While methods exist to evaluate large deviation functions numerically,^{9–13} few studies on their performance^{14,15} or extensions have been reported.¹⁶ Here we evaluate

the convergence of two methods, transition path sampling (TPS)¹⁷ and diffusion Monte Carlo (DMC).¹⁸ We find that each samples the same distribution of trajectories and can in principle converge large deviation functions. However, both methods suffer from sampling deficiencies that make applications to high-dimensional systems difficult. By illustrating these properties, we provide a basis for what is necessary to systematically extend them for use in molecular models of nonequilibrium steady states.

Large deviation theory underlies much of the recent progress in nonequilibrium statistical mechanics. Within its context, strikingly simple and general properties for systems driven away from thermal equilibrium have been derived. These include relationships between far-from-equilibrium averages and equilibrium free energies, such as the Jarzynski equality and Crooks fluctuation theorem, and symmetries within nonequilibrium ensembles, such as the Lebowitz-Spohn and Gallovti-Cohen symmetries.^{19–22} Large deviation functions are cumulant generating functions for the fluctuations of quantities in arbitrary physical systems. Many physical systems of fundamental and applied interest operate away from thermal equilibrium, either because boundary conditions generate a flux of material or energy, or because the individual constituents take in energy to move or exert force. Therefore computing large deviation functions as a means of characterizing the likelihood of fluctuations about a nonequilibrium state has the potential to broaden our understanding of physical processes ranging from crystal growth and molecular self-assembly, to nanoscale fluid flows and instabilities, to the organization of living systems and other active matter.^{23–29}

^{a)}Electronic mail: uray@caltech.edu

^{b)}Electronic mail: garnetc@caltech.edu

^{c)}Electronic mail: dlimmer@berkeley.edu

Large deviation functions are averages taken over a trajectory ensemble. As such, to compute large deviation functions for high-dimensional systems, methods for sampling trajectory ensembles must be constructed and employed. Beginning with transition path sampling,³⁰ followed by its generalizations, such as transition interface sampling³¹ and the finite temperature string method,³² such tools have been developed. As will be elaborated on below, such methods sample trajectory space by a sequential random update, resulting in a Monte Carlo walk that uniformly samples trajectories. Similarly, methods that advance an ensemble of trajectories in parallel, such as forward flux sampling³³ and the cloning algorithm,¹⁸ have been established. These methods also sample a trajectory ensemble though they do so by directly representing the ensemble by a collection of trajectories that are simultaneously updated in time. Such methods also sample a trajectory ensemble uniformly and, as with sequential update Monte Carlo methods, admit the addition of importance sampling to compute large deviation functions.

Overwhelmingly, these methods have been applied to sample rare trajectories evolved with detailed balance dynamics.³⁴ Notably, such studies have found much success in the study of the glass transition^{35–37} and protein folding.^{29,38} There have been some efforts to extend trajectory based importance sampling to driven systems, but these have been largely confined to the calculation of simpler observables, rather than large deviation functions.^{39–41} The few applications of path sampling methods to compute large deviations have been limited either to processes occurring with moderately high probability within the nonequilibrium steady state or to low-dimensional model systems.^{42–44} This is because the statistical efficiency of the standard algorithms becomes exponentially small with the exponentially rare fluctuations needed to be sampled. This has made the computational effort to converge large deviation functions intractable for many realistic systems. As will be illustrated below, the origin of the low sampling efficiency is the small overlap between the distribution of trajectories generated with a model's dynamics and the distribution whose dynamical fluctuations most contribute to the large deviation function at a given bias. This difficulty manifests itself differently in various methods, with either exponentially small acceptance probability or exponentially large variance in the weights. The generality points to a need to develop robust importance sampling schemes on top of these trajectory based Monte Carlo procedures.

In the following, we outline the basic theory of large deviations of physical processes and how they are estimated stochastically with path sampling. This serves to clarify how disparate methods like transition path sampling and diffusion Monte Carlo sample the same trajectory spaces in essentially the same way. We then explore the convergence of standard implementations of different methods by studying systems of increasing complexity: a biased Brownian particle on a periodic potential, a 1d driven lattice gas with volume exclusion, and a 2d model of self-assembly. We use these models to illuminate convergence criteria and the origin of statistical inefficiencies. Throughout we focus on the commonalities between

two canonical Monte Carlo methods and aim to draw general conclusions about their efficiencies and how they can be extended to complex systems. These lessons serve to motivate the formulation of a hierarchy of methods to mitigate the low sampling efficiencies by introducing an auxiliary dynamics. The formalism follows the technique of using a trial wavefunction in diffusion Monte Carlo and is discussed in Paper II of this work.⁴⁵

II. THEORY AND MONTE CARLO ALGORITHMS

For compactness and ease of notation, we focus on large deviation functions for processes described by time-independent, Markovian stochastic dynamics. While extensions of large deviation theory to deterministic, non-Markovian dynamics, and time-dependent processes exist, these generalizations will not be considered here. Similarly, we will only consider ensembles of trajectories with a fixed observation time though generalizations to a fixed number of events and fluctuating observation times are straightforward.⁴⁶ Much of the theory of large deviations summarized here is expanded upon in Ref. 8.

In general, Markovian processes obey a continuity equation of the form

$$\partial_t \rho(\mathcal{C}_t, t) = \mathcal{W} \rho(\mathcal{C}_t, t), \quad (1)$$

where $\rho(\mathcal{C}_t, t)$ is the probability of observing a configuration of the system, $\mathcal{C}$, at a time t and $\mathcal{W}$ is the linear operator in that configurational space that propagates trajectories. If $\mathcal{C}$ spans a discrete space, then $\mathcal{W}$ is a transition rate matrix and Eq. (1) is a master equation. If $\mathcal{C}$ spans a continuous space, then $\mathcal{W}$ is a Liouvillian and Eq. (1) is a generalized Fokker-Planck equation. In either case, provided $\mathcal{W}$ is irreducible, Eq. (1) generates a unique steady state in the long time limit, which in general produces non-vanishing currents and whose configurations do not follow Boltzmann statistics or some generalization. Only in instances where $\mathcal{W}$ obeys detailed balance is the steady state an equilibrium distribution in which all currents vanish.⁴⁷

For general nonequilibrium steady states, there exist two generic classes of time extensive observables: *configuration type* observables that are sums of functions of individual configurations,

$$\hat{\mathcal{O}}_a = \sum_{t=1}^{t_N} o_a(\mathcal{C}_t), \quad (2)$$

where o_a is an arbitrary function of a single configuration, and *current type* observables that are sums of functions of changes in configurations,

$$\hat{\mathcal{O}}_b = \sum_{t=1}^{t_N} o_b(\mathcal{C}_{t-1} \rightarrow \mathcal{C}_t), \quad (3)$$

where o_b is an arbitrary function of configurations at adjacent times. Within a nonequilibrium steady state, either observable type has an associated probability distribution that can be constructed as a marginal distribution over an ensemble of trajectories. For example, for an observable that is a combination of configurational and current types, $\mathcal{O} = \mathcal{O}_a + \mathcal{O}_b$,

$$p[\mathcal{O}] = \left\langle \delta \left[\mathcal{O} - \sum_{t=0}^{t_N} o_a(C_t) + \sum_{t=1}^{t_N} o_b(C_{t-1} \rightarrow C_t) \right] \right\rangle$$

$$= \sum_{\mathcal{C}(t_N)} P[\mathcal{C}(t_N)] \delta \left[\mathcal{O} - \sum_{t=0}^{t_N} o_a(C_t) + \sum_{t=1}^{t_N} o_b(C_{t-1} \rightarrow C_t) \right], \quad (4)$$

where $\langle \dots \rangle$ denotes the path ensemble average within the steady state distribution generated by the system dynamics and $\delta(x)$ is Dirac's delta function. A member of the trajectory ensemble is denoted $\mathcal{C}(t_N) = \{C_0, C_1, \dots, C_{t_N}\}$, which is a vector of all of the configurations that a system has visited over an observation time, t_N . Generically the probability for observing a given trajectory, $\mathcal{C}(t_N)$, is given by

$$P[\mathcal{C}(t_N)] = \rho(C_0) \prod_{t=1}^{t_N} u(C_{t-1} \rightarrow C_t), \quad (5)$$

where $\rho(C_0)$ represents the distribution of initial conditions, the suppression of an explicit time argument denotes a stationary distribution, and $u(\dots)$ are the transition probabilities. For a given nonequilibrium steady state, $\rho(C_0)$ is the stationary distribution and the transition probabilities are encoded in the propagator, $\mathcal{W}$.

Observables that are correlated over a finite amount of time admit an asymptotic, time intensive form, for their distribution function,

$$t_N^{-1} \ln p[\mathcal{O}] \sim -\phi(\mathcal{O}/t_N), \quad (6)$$

where $\phi(\mathcal{O}/t_N)$ is the rate function or generalized entropy and $\sim$ denotes the long time limit that is an equality up to sub-extensive terms $\propto 1/t_N$. In this limit, a generating function can be computed by taking the Laplace transform of Eq. (4),

$$\langle e^{-\lambda \mathcal{O}} \rangle = \sum_{\mathcal{C}(t_N)} P[\mathcal{C}(t_N)] e^{-\lambda \mathcal{O}} \sim e^{\psi(\lambda)t_N}, \quad (7)$$

where $\psi(\lambda)$ is known as the large deviation function or generalized free energy and λ is a counting field conjugate to the observable. Like an equilibrium free energy, the large deviation function is a cumulant generating function, and by taking derivatives with respect to λ ,

$$\partial^n \psi(\lambda) / \partial \lambda^n |_{\lambda=0} = (-1)^n \frac{1}{t_N} \langle (\delta \mathcal{O})^n \rangle, \quad (8)$$

the n th scaled cumulant of $\mathcal{O}$ is generated. Unlike in equilibrium, in general λ is not related to a physical field, like temperature. When the asymptotic limits exists and the rate function is a smooth convex function, then the rate function and large deviation function are related through a Legendre transform,

$$\psi(\lambda) = \inf_{\mathcal{O}} [\phi(\mathcal{O}/t_N) + \lambda \mathcal{O}/t_N], \quad (9)$$

where $\inf$ is the largest lower bound taken over all $\mathcal{O}$. This structure is reminiscent of thermodynamics, where Legendre transforms change ensembles between conjugate variables. Like in equilibrium thermodynamics, when the large deviation function is non-analytic, or nonconvex, this transform loses information about the fluctuations of the system, as it will return the convex hull of $\phi(\mathcal{O}/t_N)$. Away from these points, when the rate function is convex, the long time limit ensures

an equivalence of ensembles⁴⁸ between those conditioned on a value of $\mathcal{O}$ and those with fixed field λ , resulting in an average value of $\mathcal{O}$ as in Eq. (8). This is analogous to equilibrium in the thermodynamic limit away from points of phase transitions.⁴⁹

As an exponential average, estimating large deviation functions through straightforward simulations is difficult for large λ . Stochastic algorithms have been designed to solve this problem by means of Monte Carlo sampling. In equilibrium free energy calculations, difficulties associated with exponentially rare fluctuations are overcome by introducing a sampling bias in the form of additional terms in the Hamiltonian. Provided the additional terms, the system can be simulated or sampled and the bias can be corrected for easily because the distribution function is the known Boltzmann form. This basic procedure is the basis of numerous methods such as umbrella sampling, metadynamics, and Wang-Landau sampling.¹⁻³

However, equilibrium importance sampling methods do not translate straightforwardly to nonequilibrium steady states, as the distribution of configurations is not known generically. Added biases in the form of terms in the Hamiltonian, or propagator, are not easily reweighted. However, provided a means of sampling trajectories uniformly from their steady state distribution, added biases in the evolution of the system can be constructed to push the system into rare regions and their effects exactly corrected in order to return unbiased estimates of those rare fluctuations. To importance sample large deviation functions in nonequilibrium steady states, the trajectory ensemble is biased, or tilted, to the form

$$P_\lambda[\mathcal{C}(t_N)] = P[\mathcal{C}(t_N)] e^{-\lambda \mathcal{O}[\mathcal{C}(t_N)] - \psi(\lambda)t_N}, \quad (10)$$

where the large deviation function $\psi(\lambda)$ is the normalization constant computable as in Eq. (9). The new distribution of trajectories, $P_\lambda[\mathcal{C}(t_N)]$, represent those that contribute most to the large deviation function at a given value of λ . Ensemble averages for arbitrary observables, $\mathcal{O}'$, within the unbiased distribution and the tilted one, denoted $\langle \dots \rangle_\lambda$, are related by

$$\langle \mathcal{O}'[\mathcal{C}(t_N)] \rangle_\lambda = \frac{\langle \mathcal{O}'[\mathcal{C}(t_N)] e^{-\lambda \mathcal{O}[\mathcal{C}(t_N)]} \rangle}{\langle e^{-\lambda \mathcal{O}[\mathcal{C}(t_N)]} \rangle}, \quad (11)$$

where the denominator is $\exp[\psi(\lambda)t_N]$. In any stochastic sampling method, the ensemble average is estimated by a finite sum of trajectories. Differences between methods dictate how the ensemble of trajectories is constructed, and their efficiencies are determined by how well correlations between trajectories can be engineered and how easily biases or other constraints can be added and removed. In this manner, the importance sampling is not that different from equilibrium free energy calculations, with the exception that the simplifying forms to correct for the added bias provided by the Boltzmann distribution are replaced by implicit forms that must be computed along each trajectory. Below we discuss two Monte Carlo methods for sampling trajectories from nonequilibrium steady states with the correct weights, to which additional importance sampling can be added in order to efficiently evaluate large deviation and rate functions.

A. Transition path sampling

The distribution defined by Eq. (5) can be sampled using a Monte Carlo algorithm, analogous to path integral quantum Monte Carlo. Chandler and co-workers first identified this analogy and developed transition path sampling (TPS) to importance sample dynamical events occurring in generic classical systems.³⁰ Since its inception, TPS has become an important tool for investigating rare events, typically of the form of barrier crossings in systems evolving around an equilibrium.¹⁷ In recent years, it has been used to compute large deviation functions for systems evolving with detailed balance preserving dynamics. Most notably, this has been done for model glass-formers^{35,37,50} and glassy polymers.^{29,38} Only a very few studies have employed it in the computation of large deviation functions for driven systems, or nonequilibrium steady states.^{40,43} Nonequilibrium systems constrain the types of Monte Carlo moves that can be easily accommodated, due to incomplete knowledge of the steady state distribution. Below we review the basics of TPS and how it can be used to compute large deviation functions.

TPS generates an unbiased ensemble of trajectories by a sequential random update to an existing trajectory, using a variety of different update moves and accepting or rejecting each move according to a Metropolis rule. Usually, the update rules are collective and result in the simultaneous change of many configurations along the trajectory. To derive the Metropolis acceptance rules, the targeted stationary distribution must be known. In order to sample exponentially rare fluctuations of an observable, $\mathcal{O}$, and obtain a reasonable estimate of the large deviation function, trajectories are weighted in proportion to the path ensemble defined in Eq. (10). Given this distribution, the acceptance criterion for updating the MC walk between two trajectories $\mathcal{C}_0(t_N)$ and $\mathcal{C}_n(t_N)$ can be determined by the steady state criterion

$$P_\lambda[\mathcal{C}_0]\Gamma[\mathcal{C}_0 \rightarrow \mathcal{C}_n]\text{Ac}[\mathcal{C}_0 \rightarrow \mathcal{C}_n] \\ = P_\lambda[\mathcal{C}_n]\Gamma[\mathcal{C}_n \rightarrow \mathcal{C}_0]\text{Ac}[\mathcal{C}_n \rightarrow \mathcal{C}_0], \quad (12)$$

where $\text{Ac}[\dots]$ is the acceptance probability and $\Gamma[\dots]$ is the generation probability and we have suppressed the arguments of the trajectories, assuming that they are all t_N in length. Rearranging, we can put the acceptance probability in a Metropolis form

$$\text{Ac}[\mathcal{C}_0 \rightarrow \mathcal{C}_n] = \text{Min}\left(1, e^{-\lambda(\mathcal{O}[\mathcal{C}_n] - \mathcal{O}[\mathcal{C}_0]) - \Omega}\right), \quad (13)$$

where the first term proportional to λ is the bias away from typical events and

$$\Omega = -\ln \frac{P[\mathcal{C}_n]\Gamma[\mathcal{C}_n \rightarrow \mathcal{C}_0]}{P[\mathcal{C}_0]\Gamma[\mathcal{C}_0 \rightarrow \mathcal{C}_n]} \quad (14)$$

depends on how new trajectories are generated. If they are generated using physical dynamics, independent of λ , then Ω is a measure of the entropy production to transform between the two trajectories.⁵¹ The specific form of the entropy production depends on the details of the Monte Carlo move.

For a so-called forward shooting move, a point is picked at random along the trajectory, $\mathcal{C}_{t^*}$, and the system is reintegrated.

FIG. 1. Illustration of trajectory update moves used in TPS: (a) forward shooting, (b) backward shooting, and (c) forward shifting.

This is illustrated in Fig. 1(a). The acceptance can be computed by expanding Eq. (14),

$$e^{-\Omega} = \frac{\prod_{t=t^*+1}^{t_N} u(\mathcal{C}_{t-1}^n \rightarrow \mathcal{C}_t^n) (g[\mathcal{C}_{t^*}^n \rightarrow \mathcal{C}_{t^*}^o] \prod_{t=t^*+1}^{t_N} u[\mathcal{C}_{t-1}^o \rightarrow \mathcal{C}_t^o])}{\prod_{t=t^*+1}^{t_N} u[\mathcal{C}_{t-1}^o \rightarrow \mathcal{C}_t^o] (g[\mathcal{C}_{t^*}^o \rightarrow \mathcal{C}_{t^*}^n] \prod_{t=t^*+1}^{t_N} u[\mathcal{C}_{t-1}^n \rightarrow \mathcal{C}_t^n])} \\ = \frac{g[\mathcal{C}_{t^*}^n \rightarrow \mathcal{C}_{t^*}^o]}{g[\mathcal{C}_{t^*}^o \rightarrow \mathcal{C}_{t^*}^n]} = 1, \quad (15)$$

where transition probabilities up to t^* trivially cancel in the numerator and denominator because those pieces of the trajectory are the same. The last equality follows from assuming that the shooting point is changed symmetrically with probability $g[\mathcal{C}_{t^*}^n \rightarrow \mathcal{C}_{t^*}^o]$, independent of the configuration, for example, by drawing new random noise for the stochastic integration. With such a procedure, the move is accepted in proportion to the ratio of $\exp(-\lambda\mathcal{O})$ between the new and old trajectory.

In order to have an ergodic Markov Chain, the entire trajectory from the initial condition must be sampled. Unlike in equilibrium, for dynamics that do not obey detailed balance this is problematic, as in general the acceptance probability for backward moves cannot be manipulated to yield unity even for $\lambda = 0$. For a backward shooting move, illustrated in Fig. 1(b), the acceptance criterion is given by

$$e^{-\Omega} = \frac{(\rho(\mathcal{C}_0^n) \prod_{t=1}^{t^*} u[\mathcal{C}_{t-1}^n \rightarrow \mathcal{C}_t^n]) (\prod_{t=t^*+1}^{t_N} \bar{u}[\tilde{\mathcal{C}}_{t-1}^o \rightarrow \tilde{\mathcal{C}}_t^o])}{(\rho(\mathcal{C}_0^o) \prod_{t=1}^{t^*} u[\mathcal{C}_{t-1}^o \rightarrow \mathcal{C}_t^o]) (\prod_{t=t^*+1}^{t_N} \bar{u}[\tilde{\mathcal{C}}_{t-1}^n \rightarrow \tilde{\mathcal{C}}_t^n])}, \quad (16)$$

where $\bar{u}[\dots]$ is the time-reverse transition rate and $\tilde{\mathcal{C}}_i$ is the corresponding time reversed state of the system. From stochastic thermodynamics,^{20,51} the time reverse rates are related to the forward rates by the entropy produced, Ξ ,

$$\frac{u[\mathcal{C}_{t-1}^n \rightarrow \mathcal{C}_t^n]}{\bar{u}[\tilde{\mathcal{C}}_{t-1}^n \rightarrow \tilde{\mathcal{C}}_t^n]} = e^{\Xi^n(t-1,t)}. \quad (17)$$

Substituting this relation into the acceptance probability, it is clear that the acceptance rate is exponentially small and contains two terms: a time extensive entropy production and a boundary term with information on the initial conditions,

$$e^{-\Omega} = e^{-[W^n(t_0, t^*) - W^o(t_0, t^*)]}, \quad (18)$$

where

$$W^{o,n}(t_0, t^*) = \sum_{t=1}^{t^*} \Xi^{o,n}(t-1, t) - \ln \rho(C_0^{o,n}). \quad (19)$$

This ratio is problematic because the relative weights for the initial distribution are often not known. Even in cases where the initial distribution is known, or can be neglected as in a long time limit, for a nonequilibrium steady state, the time extensive part means that backward shooting moves are exponentially unlikely to be accepted, with a rate proportional to the length of the move and the number of driven degrees of freedom in the configuration space.

Alternatively, the initial condition and early time segments of the trajectory can be updated using forward shifting moves, illustrated in Fig. 1(c). In such a move, a configuration from the middle of the trajectory is randomly selected, C_{t^*} , and moved to the beginning of a new proposed trajectory. In order to conserve the total length of the trajectory, an additional segment is generated from the last configuration of the old trajectory. The acceptance criterion for such a move can be written as

$$e^{-\Omega} = \frac{\rho(C_{t^*}^n) \prod_{t=1}^{t^*} \bar{u}[\tilde{C}_t^n \rightarrow \tilde{C}_{t-1}^n]}{\rho(C_0^n) \prod_{t=1}^{t^*} u[C_{t-1}^o \rightarrow C_t^o]} = 1, \quad (20)$$

where the final equality assumes that the nonequilibrium driving is time-reversal symmetric and that the system is evolving within a stationary steady state. This cancelation occurs because under these assumptions the proposal and path weights are generated with the same dynamics. This is analogous to previous work whereby new initial conditions are generated by short trajectories, followed by reintegration of the full path.³⁹ Because the dynamics we consider are stochastic and irreducible, forward shooting and forward shifting are sufficient to ensure an ergodic Markov chain, as over repeated TPS moves, the entire trajectory is replaced. This is distinct from other forward only propagating methods like forward flux sampling where the initial conditions are not updated, which can result in biased sampling if sufficiently large initial distributions are not used.^{34,52} However, trajectories decorrelate more slowly without backward shooting as many shooting and shifting moves are needed to generate completely new trajectories. The acceptance criterion for backward shifting can be analogously derived and presents similar problems as backward shooting.

The moves outlined above, together with the acceptance criterion of Eq. (13), allow TPS to be used to sample trajectory ensembles of arbitrary nonequilibrium systems in proportion to their correct statistical weight. Expectation values of the form of Eq. (11) can be estimated straightforwardly by

$$\langle \mathcal{O}' \rangle_\lambda = \frac{1}{N_{\text{traj}}} \sum_{i=1}^{N_{\text{traj}}} \mathcal{O}'[\mathcal{C}_i(t_N) | \lambda], \quad (21)$$

where N_{traj} is the number of trajectories harvested by TPS. This is because by construction they sample $P_\lambda[\mathcal{C}(t_N)]$ without a bias. Such expectation values can be related to the unbiased

steady state distribution using Eq. (11). In order to estimate the large deviation or rate function, simulations can then be run with multiple values of λ , and provided the distributions collected at various values of λ overlap, $\psi(\lambda)$ can be computed from the relation

$$\phi_\lambda(\mathcal{O}) = \phi(\mathcal{O}) + \lambda \mathcal{O}/t_N + \psi(\lambda), \quad (22)$$

where $\phi_\lambda(\mathcal{O})$ is the marginal distribution, or unnormalized rate function, generated from Eq. (10). For multiple distributions, calculations with many values of λ can be related to the standard histogram reweighting techniques, such as the weighted histogram analysis method or the multistate Bennet acceptance ratio (MBAR), which provide optimal estimates of $\psi(\lambda)$, and unbiased $\phi(\mathcal{O})$ computed far into the tails of its distribution.^{53,54} Histogram reweighting techniques can only determine $\psi(\lambda)$ up to a global constant but because the unbiased dynamics are normalized, $\psi(0) = 0$, so the full function of $\psi(\lambda)$ can be uniquely determined. These procedures are a straightforward generalization of traditional umbrella sampling calculations in equilibrium, only now with an ensemble of trajectories.

B. Diffusion Monte Carlo

An alternative to sampling trajectory space with sequential random updates is to sample trajectories in parallel. This is the strategy proposed in diffusion Monte Carlo (DMC) and has been adopted for the study of nonequilibrium classical stochastic dynamics by algorithms known as forward flux sampling and the so-called cloning algorithm.^{12,16,55} In both methods, an ensemble of initial conditions is generated using some prior distribution. Each initial condition, referred to here as a walker, is propagated for some short time, creating an ensemble of short trajectories. Depending on the specific algorithm and observable of interest, importance sampling is incorporated by assigning weights to each walker that accumulate over the short trajectories and upon iterations of parallel updates to the full trajectory ensemble. Generically, walkers are weighted in proportion to the biased ensemble defined in Eq. (10). The cloning algorithm offsets the exponential weights with a population dynamics, whereby walkers with low weight are deleted and those with high weight are replicated or “cloned.” The cloning algorithm in particular has been successful in computing large deviation functions for model lattice systems^{56,57} and some small atomistic systems.⁵⁸ Mostly it has been used to study models of transport in 1 dimension.^{18,59,60} Below we review the basics of DMC and how it can be used to compute large deviation functions.

The form of $P_\lambda[\mathcal{C}(t)]$ suggests an alternative way in which trajectories may be generated. Expanding P_λ , we get

$$P_\lambda[\mathcal{C}(t_N)] = \frac{\rho(C_0) \prod_{t=1}^{t_N} u(C_{t-1} \rightarrow C_t) e^{-\lambda o(C_{t-1} \rightarrow C_t)}}{\sum_{\mathcal{C}(t_N)} \rho(C_0) \prod_{t=1}^{t_N} u(C_{t-1} \rightarrow C_t) e^{-\lambda o(C_{t-1} \rightarrow C_t)}}, \quad (23)$$

where we have written $\mathcal{O}$ as a current type observable, $\mathcal{O}[\mathcal{C}(t_N)] = \sum_{t=1}^{t_N} o(C_{t-1} \rightarrow C_t)$, otherwise it could be an additive function of single configurations. The argument inside the product is the transition probability times a bias factor.

This combination of probabilities does not represent a physical dynamics, as it is unnormalized. However, it can be interpreted as a population dynamics as was done by Giardinà, Kurchan, and Peliti and is done in the standard DMC, where the nonconservative part proportional to the bias is represented by adding and removing walkers.¹¹

In particular, in the cloning algorithm, trajectories are propagated in two steps. First N_w walkers are propagated according to the normalized dynamics specified by $u[C_{t-1} \rightarrow C_t]$ for a time, Δt . Over this time, the bias is accumulated according to

$$w_i(t, \Delta t) = \exp(-\lambda O[C_i(t - \Delta t), C_i(t)]), \quad (24)$$

where $O[C_{t-\Delta t}, C_t] = \sum_{t'=t-\Delta t}^t o[C_{t'} \rightarrow C_{t'+1}]$ is the collected current, which due to the multiplicative structure of the Markov chain is simply summed in the exponential. After the trajectory integration, $n_i(t)$ identical copies of the i th trajectory are generated in proportion to $w_i(t, \Delta t)$. This is known as a branching step. In the standard quantum-DMC, Δt is chosen to be the time step corresponding to trotterization of the Hamiltonian.⁶¹ Additionally, instead of using the bias to branch, it is advantageous to incorporate it into the integration step and use the normalization as the branching weight. However, in many applications of non-equilibrium stationary states, calculating the normalization may not be straightforward.

There are different ways of performing the branching step that either holds the total number of walkers constant or not. In a simple formulation with a non-constant walker population, the number of replicated walkers is determined stochastically by

$$n_i(t) = \lfloor w_i(t, \Delta t) + \xi \rfloor \quad (25)$$

where ξ is a uniform random number between 0 and 1 and $\lfloor \dots \rfloor$ is the floor function. Alternatively, the floor function may be avoided by using a partial weight for the walker, which decreases the statistical noise. The process of integration followed by branching is continued until $t = t_N$. In this case, the large deviation function can be directly evaluated from the walker population as

$$\psi(\lambda) = \frac{1}{t_N} \ln \prod_{t=1}^{t_N/\Delta t} \frac{N_w(t)}{N_w(t-1)}, \quad (26)$$

which follows from the equivalence between the normalization constant of the biased distribution and the large deviation function. This relation clarifies that the number of walkers needed to resolve $\psi(\lambda)$ is exponentially large in λ , and the time extensive observable, $\mathcal{O}$. Because of this exponential dependence, this method is not practical and requires that the population is controlled in some way to mitigate walker explosion or total walker deletion.

One way to do population control is to formulate the DMC algorithm with a constant walker population. In this case, as before each walker is initially integrated over a time Δt . At this end of the short trajectory, $n_i(t)$ copies of the i th walker are generated stochastically by

$$n_i(t) = \left\lfloor N_w \frac{w_i(t, \Delta t)}{\sum_{j=1}^{N_w} w_j(t, \Delta t)} + \xi \right\rfloor, \quad (27)$$

where ξ is a uniform random number between 0 and 1. Generally this process will result in a different population of walkers, and so deficiencies or excesses in the walker population are compensated by cloning or deleting walkers uniformly. In this case, the large deviation function can be evaluated at each time as

$$\psi^t(\lambda) = \ln \frac{1}{N_w} \sum_{i=1}^{N_w} w_i(t), \quad (28)$$

which is an exponential average over the bias factors of each walker. This local estimate can be improved by averaging over the observation time,

$$\psi(\lambda) = \frac{1}{t_N} \sum_{t=1}^{t_N/\Delta t} \psi^t(\lambda), \quad (29)$$

which upon repeated cycles of integration and population dynamics yields a refined estimate of $\psi(\lambda)$. Since the population dynamics represents the path probability, averages can be computed with walker weights equal to

$$\langle \mathcal{O}' \rangle_\lambda = \frac{1}{N_w} \sum_{i=1}^{N_w} \mathcal{O}'[\mathcal{C}_i(t_N)], \quad (30)$$

where the trajectory $\mathcal{C}(t)$ has been generated according to $P_\lambda[\mathcal{C}(t)]$ in Eq. (23). This estimator is only valid if the number of walkers is held constant. We note here that instead of using a constant population of walkers, it is possible to use a non-constant walker population with a chemical potential as is done in the standard quantum DMC. This can lead to better statistics since in this case the uniform redistribution of excesses or deficit of walkers is avoided.⁶²

The cloning algorithm utilizes the multiplicative structure of the Markov dynamics to tune the relative times between population steps, Δt , in order to effect adjustments in the size of the biases used to add or remove trajectories. The choice of Δt is critical as it scales the overlap between the exponential bias and the generating dynamics. Population sizes can be decreased for smaller Δt since the distribution from the generating dynamics is wider for smaller Δt and the exponential decays more slowly since $\mathcal{O}$ is time-extensive. However, the wider the distribution, the larger the variance in $\psi^t(\lambda)$, so local estimates of the trajectory distribution and its bias are poorer. This tension results in a sensitivity to Δt that requires it to be chosen such that it is large enough to affect good local estimates of the large deviation function, but small enough that the overlap between the bias factor and the bare dynamics is large. However, it cannot be too large since population control is only done after Δt , and so the systematic error from finite population will grow with Δt as will the stochastic error from finite walker population.

The time-series that result from DMC can be considered in two ways: in terms of continuous or discontinuous trajectories as illustrated in Fig. 2. The continuous trajectories can be constructed by tracing the configurations backwards from $\mathcal{C}(t_N)$ keeping those branches that have been cloned. In this case, trajectories sample exactly $P_\lambda[\mathcal{C}(t)]$ and expectation values can be computed just as in TPS. This means that histogram reweighing procedures outlined in Sec. II A can be used to compute $\phi(\mathcal{O})$, instead of the typical method of

FIG. 2. Illustration of the diffusion Monte Carlo algorithm. The red and black lines represent different walker trajectories. (a) Illustrates the multiplicative weight of trajectories constructed during forward propagation. (b) Non-constant walker population method, whereby branching and deletion is done in proportion to w_i . (c) Constant walker population method where branching and deletion result in continuous (solid lines) and discontinuous (dashed lines) trajectories.

computing $\phi(\mathcal{O})$ from an inverse of the Legendre transform. The discontinuous paths correspond to trajectories constructed by tracking the configurations of a walker as they are propagated forward in time. In this case, each cloning event introduces an abrupt change in configuration space since a walker's configuration may be replaced by configurations from other walkers that carry larger weights. The ensemble described by discontinuous trajectories is not equal to the trajectories of the biased ensemble generally. For sufficiently large number of walkers and long enough propagation time, the two ensembles are strongly overlapping.

C. Comparisons between TPS and DMC

Both TPS and DMC are Monte Carlo procedures to sample trajectory space with the correct path weight. While their implementations and executions are different, both evaluate expectation values stochastically and do so by updating the system of interest using its unmodified propagator. In both methods, the computational cost scales as the number of trajectories generated times the observation time. Systematic errors arising from finite length trajectories or finite walker populations can be alleviated in both, as will be shown below. The resulting Monte Carlo walk in trajectory space is explored differently in the different algorithms. However, provided each can relax to a stationary distribution, estimates of comparable statistical accuracy require the same number of uncorrelated trajectories and thus the same computational effort.

In TPS, the sequential update means that complete trajectories of length t_N are constructed and so global constraints such as temporal boundary conditions imposed on the trajectory ensembles are naturally accommodated. Indeed, transition path sampling was conceived of as a means of sampling equilibrium trajectory ensembles with hard constraints on the initial and final configurations, as needed in the study of barrier crossing events. However, because large deviation functions result from a long time limit, individual long trajectories must

be constructed. Maintaining correlations for such long trajectories with the standard shooting and shifting moves is difficult especially when the ensemble sampled by the natural system dynamics does not have significant overlap with the ensemble that contributes to the large deviation function at finite λ . As will be discussed below, this can result in significant statistical errors.

In contrast to TPS, DMC constructs paths in parallel using an ensemble of walkers and implicitly benefits from modern computer architectures. However, the propagation in DMC is uni-directional with path construction relying heavily on the distribution of paths in a local interval of time, Δt . Accurate sampling of the local distribution therefore becomes exceedingly important and due to the exponential scaling of the bias becomes hard to accomplish with limited walker populations. Formally the number of walkers needed to accurately converge the calculation scales exponentially with λ . This complexity results in individual trajectories carrying exponentially larger weights than others, which means they are likely to be replicated more often. This can produce an ensemble of walkers that becomes highly correlated over long times and finite walker populations. Modulating these walker correlations is the key to controlling statistical errors in DMC.

III. CONVERGENCE CRITERIA AND STATISTICAL UNCERTAINTY

In order to understand the utility of these algorithms, and uncover their limitations, we have studied three different models. The first is forced Brownian motion that evolves in continuous space and time, but whose large deviation function can be computed exactly. The second is a driven lattice gas that evolves in continuous time but discrete space, whose two-body interactions result in correlated many-body dynamics. The third is a minimal model of nonequilibrium self-assembly that evolves in discrete space and time but exhibits a nontrivial phase diagram including critical fluctuations. These three systems allow us to ground a discussion of the systematic and statistical errors that occur in calculating large deviation functions.

A. Long time or large population limit

Finite time and finite walker population¹⁴ are the largest source of systematic error in the calculation of large deviation functions. Other sources of error that exist independent of importance sampling, such as those resulting from the discretization of continuous time dynamics are well understood and controllable.⁶ The definitions for the limiting forms of the probability distribution or its cumulant generating function, Eqs. (6) and (9), are exponential equalities. Finite time effects result from sub-exponential terms whose origins are from finite time correlations or from temporal boundary conditions. In TPS, the long time limit is approached by increasing t_N directly. Deviations from the asymptotic values of ϕ or ψ come directly from the temporal boundary conditions that are free to fluctuate independently and will bias initial and final parts of the trajectory toward values typical of the system at $\lambda = 0$. In DMC, provided trajectories have been integrated into the steady state, the trajectory space is rotated into a

space of walkers, and the effective long time limit can be approached by increasing walker population at fixed integration time. Because of the assumed ergodicity of the nonequilibrium steady state, averages computed from an ensemble of trajectories are expected to be equivalent to those computed from an infinitely long single trajectory. In DMC, provided the fixed integration time is long compared to the correlation times of the system dynamics, ϕ or ψ can be converged solely by increasing the walker population. Alternatively, one can take a hybrid approach, where a finite number of walkers can be compensated by longer propagation times.

To explore the systematic errors arising in the calculations of large deviation functions, we consider a paradigmatic model for driven systems: Brownian motion in a one-dimensional periodic potential with an external force.^{63,64} The system is sketched in Fig. 3(a). This model has been studied extensively in the context of the fluctuation theorem and is exactly solvable numerically by diagonalization of the Fokker-Planck operator through basis set expansion techniques.⁶⁵ As such, it provides an ideal reference to understand errors resulting from finite trajectory length or finite walker population.

The overdamped dynamics of a Brownian particle moving in a closed ring geometry with circumference L and external periodic potential $V(x) = v_o \cos(2\pi x/L)$ obeys a Langevin equation

$$\dot{x} = -\nabla V(x) + f + \xi, \quad (31)$$

where the dot denotes the derivative with respect to t . The particle is driven out of equilibrium through the non-gradient force f , resulting in a non-vanishing current around the ring. The

random force ξ describes the effective interactions between the particle and environment fixed at temperature T with statistics

$$\langle \xi(t) \rangle = 0 \quad \langle \xi(t)\xi(t') \rangle = 2\mu T \delta(t - t'), \quad (32)$$

where the mobility of the particle has been set to unity. In the results we present below, we set $L = T = 1$, $v_o = 2$ and consider the effects of changing a non-dimensionalized external force, $f^* = f/2\pi v_o \geq 0$.

The dynamics generated by Eq. (32) are well characterized by two limiting behaviors. In the limit that f^* is small, a typical trajectory is localized near a potential minimum with rare excursions to adjacent minima. Under such conditions, the average particle flux around the ring is small. In the limit that f^* is large, a typical trajectory exhibits frequent hops between adjacent minima, biased in the direction of the non-gradient force. Under these conditions, the average particle flux around the ring is large. Examples of both types of trajectories are shown in Fig. 3(a). A parameter that characterizes the dynamical behavior of the system is the entropy production

$$\Sigma(t) = \int_0^t f dx(t') = f \Delta x(t) \quad (33)$$

which is time extensive and in this case simply proportional to the particle flux around the ring. We denote its time intensive counterpart, $\sigma(t_N) = \Sigma(t_N)/t_N$. In order to evaluate the finite time effects for TPS and DMC, we compute the rate function for entropy production, $\phi(\sigma)$, and its associated large deviation function, $\psi(\lambda)$, for this model.

To simulate the model of driven Brownian motion, we use a second-order stochastic Runge-Kutta integrator with time step 0.001.⁶⁶ For calculations using TPS, we simulated trajectories of length $t_N = 2000$, using forward shooting and shifting moves and correlated noises.⁵¹ Trajectories are harvested using the acceptance criterion given in Eq. (13), with a bias of $\exp[-\lambda t_N \sigma(t)]$, for λ values ranging between -1 and 1 in increments of 0.025 . Shooting and shifting points were optimized to result in a 30% acceptance ratio and around 10^6 trajectories were harvested for each value of λ . Rate functions were computed using MBAR and error bars were estimated using a bootstrapping analysis. For calculations using DMC, we set $\Delta t = 0.05$ and $t_N = 20$, and a constant walker population of $N_w = 10\,000$. Branching probabilities were computed from Eq. (27). Values of λ between -1 and 1 were used to compute $\psi(\lambda)$ using Eq. (29), and a numerical Legendre transform was used to determine $\phi(\sigma)$.

Figure 3(b) compares the rate functions for three different cases, characterized by an increasing force, f^* . For finite f^* , $\phi(\sigma)$ exhibits a minimum at a finite value for the entropy production. With increasing f^* , the mean value of σ increases and the curves broaden. In addition, there is an asymmetry which becomes more notable with increasing f^* . This latter feature arises from breaking microscopic reversibility and is codified in the fluctuation theorem,

$$\phi(\sigma) - \phi(-\sigma) = \sigma. \quad (34)$$

As derived for diffusion processes by Kurchan,⁶⁷ the fluctuation theorem relates the probability of generating entropy to

FIG. 3. Rare entropy production fluctuations of a driven Brownian walker. (a) Typical trajectories for instances of large $f^* = 1.5$ (red) and small $f^* = 0.5$ (blue). Inset is a schematic of the model. (b) Rate functions for the entropy production for $f^* = 0.8$ (triangles), 1.0 (squares), and 1.5 (circles) computed with TPS (red) and DMC (blue). The black curves are numerically exact results.

the probability of destroying it. Equation (34) implies that as the entropy production gets larger, the likelihood of seeing a fluctuation that reduces entropy becomes exponentially rarer. Figure 3(b) compares $\phi(\sigma)$ computed with TPS and DMC to numerically exact results. For all three cases studied, both methods are able to recover the exact result quantitatively. From the fluctuation theorem, Eq. (34), it is clear that trajectories that generate negative entropy are rarer, and as a result, the error bars for both TPS and DMC are larger for $\sigma < 0$. The ability to reproduce the exact result itself is not surprising for this simple model. There are no approximations introduced in the algorithms, so in the limit of well sampled ensembles, the converged expectation values should agree with the exact result within their statistical error bar. It is then only a question of how quickly the long time or large walker population limit is reached so that finite time effects are converged to a given accuracy.

For the model of driven Brownian motion, we find that the long time limit for the rate function is approached as

$$\phi(\sigma) - \tilde{\phi}(\sigma, t_N) \propto 1/t_N + \mathcal{O}(1/t_N^2), \quad (35)$$

where $\tilde{\phi}(\sigma, t_N)$ is the finite time estimate of the asymptotic $\phi(\sigma)$. Therefore to leading order, finite time corrections scale as $1/t_N$ with a proportionality constant that depends on the correlation time for fluctuations in σ . Data taken from TPS calculations with varying t_N are shown in Fig. 4(a) for a number of different λ values. In this model, time correlations arise due to barrier crossing events that result in large changes in σ relative to the typical small amplitude changes from fluctuations of the particle at the bottom of the potential well. For small f^* or $\lambda > 0$, these barrier crossing events occur infrequently,

and as a consequence, the times along a trajectory where they occur are uncorrelated. By contrast, for large f^* or $\lambda < 0$, barrier crossings occur often, and accordingly σ fluctuations exhibit longer correlation times. Because of this, the prefactor accompanying the $1/t_N$ correction to the rate function becomes larger for increasing λ , or increasing f^* , as shown in Fig. 4(a).

Provided Δt is large enough to produce the steady state, the systematic errors in DMC depend on the number of walkers. Formally, for branching rates obeying Eq. (27), the number of walkers required to offset the bias is exponential in $\lambda \sigma \Delta t$. For a constant walker number algorithm, Nemoto *et al.* have shown⁶⁸ that the leading order error in the rate function scales as

$$\psi(\lambda) - \tilde{\psi}(\lambda, N_w) \propto 1/N_w + \mathcal{O}(1/N_w^2), \quad (36)$$

where $\psi(\lambda)$ is the Legendre transform of $\phi(\sigma)$ and $\tilde{\psi}(\lambda, N_w)$, its finite population estimate. This is indeed what we find for the driven Brownian motion model. As shown in Fig. 4(b), for various values of λ , we find that systematic errors in the estimate of the large deviation function are inversely proportional to the walker population. This is due to the exponential error in representing the steady state distribution made by DMC with a finite number of walkers. As an expectation value of an exponential quantity, $\psi(\lambda)$ consequently depends linearly on this population error. The coefficient determining the proportionality between the stochastic error and the number of walkers is related to the number of correlated walkers, which increase with increasing λ .

B. Correlations in trajectory space

In TPS and DMC with branching, expectation values are estimated by summing over trajectories generated with the correct path weights so that each entry into the sum is weighted equally. Because in practice this sum is finite, the existence of statistical errors determines the accuracy with which expectation values, like large deviation functions, can be computed. For a given number of trajectories, the statistical efficiency is determined by the number of trajectories that are uncorrelated. In TPS, the sequential update results in correlations between trajectories that originate from a common trajectory and share some configurations in common. The relevant measure of these correlations is the time, in number of TPS moves, over which trajectories remain correlated. In DMC, the parallel branching rules replace lower weight walkers with those of higher weight, resulting in a given walker history that can be traced back to some number of common ancestors. Correlations within the trajectory ensemble are measured by the number of walkers that originate from the same source. In both cases, for traditional implementations of these algorithms, we find that correlations rapidly increase for increasingly rare fluctuations. In practice, this renders the convergence of large deviation functions untenable for all but the smallest system size, or λ .

To explore the origin of statistical uncertainties in the estimation of large deviation functions, we consider a family of models known as simple exclusion processes (SEPs). These models are well studied examples of one dimensional transport and are considered paradigmatic models of nonequi-

FIG. 4. Systematic errors in large deviation function of the entropy production for the driven Brownian walker with $f^* = 0.8$, $v_o = 2$. (a) Finite time error computed with TPS for $\lambda = -0.25$ (blue), 0.0 (red), and 0.25 (black). Solid lines are best fits of the form A/t_N . (b) Finite walker number error computed with DMC for $\lambda = -0.5$ (blue), -0.2 (red), and 0.2 (black). Solid lines are best fits of the form B/N_w .

librium many-body physics.⁶⁹ The simple exclusion process evolves on a lattice with L sites, and in the following, we consider the open case whereby particles are injected or removed at site 0 and $L - 1$. At a given time, t , the configuration of the system is defined by a set of occupation numbers, $n_i = \{0, 1\}$, e.g., $\mathcal{C}(t) = \{0, 1, \dots, 1, 1\}$. The probability of observing a given configuration at a specified time, $\rho(\mathcal{C}_t)$, can be found from the solution of the continuity equation in Eq. (1), with a Markovian rate matrix, $\mathcal{W}$, whose elements have the form

$$\mathcal{W}(\mathcal{C}, \mathcal{C}') = -r(\mathcal{C}) + \sum_{\mathcal{C} \neq \mathcal{C}'} w(\mathcal{C}, \mathcal{C}'), \quad (37)$$

where $r(\mathcal{C}) = \sum_{\mathcal{C}'} w(\mathcal{C}, \mathcal{C}')$ is the exit rate and $w(\mathcal{C}, \mathcal{C}')$ has the elements $w = w_L + \sum_i w_i + w_R$, where

$$w_L = \begin{bmatrix} -\alpha & \gamma \\ \alpha & -\gamma \end{bmatrix}, \quad w_R = \begin{bmatrix} -\delta & \beta \\ \delta & -\beta \end{bmatrix},$$

$$w_i = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & -q & p & 0 \\ 0 & q & -p & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}.$$

The matrices w_L and w_R act on the first and last sites, respectively, in a basis of single particle occupancy, i.e., $n_i = \{0, 1\}$. Particles are inserted at the boundaries with rates α and γ and removed with rates β and δ . The matrices w_i act on all L sites and are expressed in a multi-particle basis, i.e., $n_i n_{i+1} = \{00, 01, 10, 11\}$. Particles move to the right with rate p and to the left with rate q , subject to the constraint of single site occupancy. This hard core constraint results in correlations between particles moving on the lattice. These rates and the hard core exclusion are illustrated in Fig. 5(a).

For most values of $\alpha, \gamma, \delta, \beta, q$, and p , an average current of particles flows through the system. We denote the particle current by $Q(t_N)$, which is equal to the number of hops to the right minus the number of hops to the left over a time t_N ,

$$Q(t_N) = \sum_{t=1}^{t_N} \sum_{i=0}^{L-1} \delta_{i+1}(t) \delta_i(t-1) - \delta_i(t) \delta_{i+1}(t-1), \quad (38)$$

where δ_i is the Kronecker delta function for site i and the sum runs over the lattice and t_N . Finite steady state currents can be generated either by boundary driven or bulk driven processes. For $p = q$, boundary driven currents occur when α, β , and γ, δ are such that they establish mean densities on the left and right ends of the lattice that are unequal. For $q \neq p$, bulk driven currents are produced and remain finite in the infinite lattice size limit, unlike boundary driven currents which decay following Fick's law as $1/L$. For $q \neq p$ and different values of the boundary rates, the model exhibits a nontrivial phase diagram, with different density profiles and mean currents.⁷⁰ A representative trajectory in the maximum current phase⁷⁰ is shown in Fig. 5(a). As a relatively low dimensional many-body system, SEP models represent an ideal testing ground for studying statistical errors in the computation of large deviation functions.

We have computed large deviation functions for parameters that include both bulk and boundary driven transport. For

FIG. 5. Large deviation functions for the particle current in the SEP models. (a) Schematic of the SEP model, 1d lattice gas with nearest neighbor hopping rates and double occupancy exclusion. A typical trajectory shown in space-time, where the dark lattice sites denote an occupied position. (b) Large deviation function for $L = 8$ boundary (circles) and bulk (squares) driven SEP models, computed with TPS (red), DMC (blue), and exact diagonalization (solid line).

the case of boundary driven transport, we study a model with $\alpha = \delta = 0.9$, $\beta = \gamma = 0.1$, and $q = p = 0.5$. For the case of bulk driven transport, we study a model with $\alpha = \delta = \beta = \gamma = 0.5$, and $p = 0.75$ and $q = 0.25$. For both cases, we consider fairly small lattices, $L = 8$. The dynamics are propagated by sampling the trotterized rate matrix using kinetic Monte Carlo,⁷¹ with a time step of 0.05 that is commensurate with single particle moves. For calculations using TPS, we simulated trajectories of length $t_N = 20$, using forward shooting and shifting moves. Trajectories are harvested using the acceptance criterion given in Eq. (13), with the bias $\exp[-\lambda Q(t_N)]$. Shooting and shifting points were optimized to result in a 30% acceptance ratio and around 10^5 – 10^6 trajectories were harvested for each value of λ . Large deviation functions were solved self-consistently with MBAR as in Eq. (22). For calculations using DMC, we used integration times, $\Delta t = 0.5$, $t_N = 100$, and a constant walker population of $N_w = 2000$. Branching probabilities were computed from Eq. (27), and $\psi(\lambda)$ was estimated from Eq. (29).

Figure 5(b) show the large deviation functions computed for fluctuations of Q . As before, there is quantitative agreement between both TPS and DMC. Figure 5(b) compares the results for both bulk and boundary driven cases obtained by sampling with TPS and DMC and exact numerical results obtained by computing the maximum eigenvalue of the many-body rate matrix using exact diagonalization. Across the range of λ 's, the sampled values agree with those numerically exact results. In all cases, the simple convex shape of the large deviation functions is a consequence of the fluctuation theorem, the one derived by Lebowitz and Spohn for Markov jump processes. For a 1d system such as SEP, the total entropy

production along a trajectory is linearly proportional to the current.⁷² The proportionality constant is known as the affinity, $A = \ln(\alpha\delta p^{L-1}/\beta\gamma q^{L-1})/L + 1$, and results in a slightly different symmetry for current fluctuations,

$$\psi(\lambda) = \psi(A - \lambda), \quad (39)$$

where for an equilibrium system, $A = 0$, and the fluctuation theorem is again a statement of microscopic reversibility.²¹ For SEP models, the value of A depends on whether the current is generated by bulk or boundary driving, and as a consequence, the curves in Fig. 5(b) are symmetric about different values.

As before, the main systematic errors in converging $\psi(\lambda)$ scale as $1/t_N$ for TPS and $1/N_w$ for DMC. For the systems considered here with $L = 8$, for $|\lambda| \leq 1$ the large deviation functions can be converged. However with increasing λ , the statistical effort grows significantly. For the TPS algorithm, the origin of the statistical inefficiencies can be understood by computing the average acceptance probability for moves made at different times along the trajectory. Shown in Fig. 6(a) is the average acceptance probability for the boundary driven SEP model, computed for $\lambda = 0.1$. We find that the acceptance probability is large at the ends of the trajectory and vanishing in its middle. This confinement of acceptance illustrates that

FIG. 6. Statistical efficiencies of large deviation function calculations for the SEP model. (a) The main figure shows the average acceptance ratio computed with TPS for $t_N = 50$ and $\lambda = 0.1$. The red line is computed for forward shifting moves and the blue line is computed for forward shooting moves. The inset is the correlation time in a number of TPS moves as a function of λ . (b) The fraction of independent walkers computed with DMC for $\Delta t = 0.5$ and $\lambda = 0.1$. The inset is the number of correlated walks as a function of λ for the same system.

at this value of λ only small perturbations to a trajectory can be accepted. The decay of acceptance probability away from the ends becomes increasingly abrupt for larger λ .

The efficiency of a Monte Carlo algorithm depends on how many independent samples can be generated with a given computational effort. Diffusion in trajectory space will be small if only small changes to the trajectory are made at each Monte Carlo step. In order to quantify how correlated trajectories are during the sequential updating of TPS, we compute a correlation function of current fluctuations. For trajectories separated by τ TPS moves, we define the trajectory space function as

$$I(\tau) = \langle \delta Q[\mathcal{C}_0] \delta Q[\mathcal{C}_\tau] \rangle, \quad (40)$$

where δQ is trajectory's current minus its ensemble average, $Q - \langle Q \rangle$. This correlation function is monotonically decreasing, so it is uniquely characterized by its decay time which we define by

$$\tau_c = \int_0^\infty d\tau I(\tau)/I(0). \quad (41)$$

The correlation time τ_c is a time measured in TPS Monte Carlo moves. The inset of Fig. 6(a) shows τ_c as a function of λ . The correlation time has been computed using moves that result in an acceptance probability of 30% so that we are effectively probing the characteristic decay of acceptance probability off of the ends trajectory. Because this results in only small changes to a trajectory to be accepted, subsequent trajectories stay correlated over more moves for increasing λ . We find that over a small range of $\lambda = 0.01$ – 0.2 , this increase is abrupt, growing over 4 orders of magnitude. To have accurate estimates, many statistically independent samples must be generated by a Monte Carlo procedure, but this rapidly decreasing statistical efficiency prohibits convergence for $\lambda > 1$. As the bias being offset by the acceptance criteria is an exponential of λ times an extensive order parameter, we expect that the correlation time is an exponentially sensitive function of λ . Because the order parameter grows with the system size or trajectory length, the range of possible updates to the trajectory becomes small for large systems or t_N , making studying all but the smallest fluctuations for large systems cumbersome, if not numerically impossible.

A similar analysis can be done to understand the statistical efficiencies of DMC. In the cloning algorithm, rather than shifting the bias into acceptance criteria as TPS does, it funnels probability into a population dynamics. Correlations between trajectories result from walkers that have common ancestors. Recently, Nemoto *et al.* introduced a useful measure of this correlation, as the fraction of independent clones, $F_i(t)$, which measures the probability that a descendant of a walker at a time t survives up to the final time t_N .¹⁶ This quantity is shown in Fig. 6(b) along a $\Delta t = 0.5$ step for the bulk driven SEP model with $N_w = 2000$. Just as for the acceptance ratio in TPS, $F_i(t)$ is strongly peaked at t_N and decays quickly from the boundary. Analogously, $F_i(t)$ reports on the number of independent walkers and the decay time from t_N sets a rate over which walkers become correlated.

We can quantify the statistical efficiency of DMC by computing the total number of correlated walkers, N_c , over the

repeated iterations of propagation and population dynamics. In terms of $F_i(t)$,

$$N_c = N_w(1 - F_i(0)), \quad (42)$$

where $F_i(t)$ is evaluated in the limit of the time going to 0 where we generically find that it plateaus. This estimate is the number of clones that share a common ancestor somewhere along the simulation. This quantity is plotted in the inset of Fig. 6(b) as a function of λ . We find that the number of correlated walkers is strongly increasing as a function of λ . This means that as we probe rarer fluctuations, fewer walkers carry significant weight. In fact, to represent the large deviation function an exponentially large number of walkers as a function of λ must be used. As the statistical errors are inversely proportional to the number of independent walkers,⁶⁸ the estimates for large λ become exponentially worse.

For both algorithms, the origin of sampling difficulties is in the exponentially decreasing overlap between the ensemble generated by the bare dynamics, $P[\mathcal{C}(t_N)]$, and the ensemble that contributes most to the large deviation function at the given λ 's through the acceptance criteria in TPS or the branching weights in DMC, $P_\lambda[\mathcal{C}(t_N)]$. Both algorithms propagate trajectories without any information from the exponential basis that distinguishes these distributions and attempt to offset all of the exponential weight onto a subsequent importance sampling step. In the case of TPS, all of the weight is put on the acceptance criteria for making moves in trajectory space. This results in an acceptance rate that becomes exponentially small, and subsequently trajectories that become correlated over exponentially large number of moves. In the case of DMC, all of the weight is put into the branching step. This results in a walker population that must be exponentially large in order to overcome the exponentially different walker weights. Because the targeted ensemble is weighted by a factor that is an exponential of a time and system size extensive quantity, these deficiencies severely limit the current practical utility of these standard algorithms.

C. Alternate biasing form

The final system we consider is the two dimensional, multicomponent, restricted solid-on-solid (MRSOS) model.⁷³ This model was recently studied in the context of nonequilibrium crystal growth, where it was observed that by tuning the growth rate the system can undergo an order-disorder transition with a concomitant critical point and regions of bistability.⁷⁴ As a two-dimensional model that exhibits critical fluctuations, it presents an ideal testing ground for exploring sampling difficulties that result from large typical fluctuations and strong finite size effects. Also, because it has the potential for bistability, over short times the rate function may exhibit non-convex parts, analogous to equilibrium free energies for finite size systems at phase coexistence. Indeed few importance sampling techniques have been able to importance sample these fluctuations efficiently. A notable exception is recent work by Klymko *et al.*, who derived an effective dynamics capable of sampling rare trajectories for this model.⁷⁵

The model is defined by an energy function,

$$E = -J \sum_{\langle i,j \rangle} s_i s_j - \mu \sum_i |s_i|, \quad (43)$$

where $s_i = \{\pm 1, 0\}$, $J > 0$ is the standard ferroelectric coupling, μ is the chemical potential for adding a spin of either sign, and $\langle i, j \rangle$ denotes a sum over only nearest neighbors. The model is evolved with a Metropolis Monte Carlo dynamics, with uniform addition and deletion probabilities. The spin additions or deletions are constrained to only a thin layer at the interface of the growing lattice. Specifically, we impose the solid-on-solid restriction⁷⁶ that does not allow for the addition or removal of a spin with $s = \pm 1$ to a site if the difference in the new height and the height of adjacent sites with $s = \pm 1$ exceeds 1. This essentially freezes the interior of the lattice, allowing for spin flips and deletions only at the interface. Starting from an initial condition with $s_i = 0$ for all i , and a flat interface at the origin, for $\mu > -0.5$, on average the lattice will grow through the sequential additions of spins, creating a net mass current. Periodic boundary conditions are imposed in the direction parallel to the interface. In the direction perpendicular to the interface, mixed boundaries are imposed with a fixed boundary at the origin and an open boundary in the direction of net growth.

The energy function in Eq. (43) is just the Ising Hamiltonian and at equilibrium the structure should belong to a two dimensional system independent of growth mechanism. Consequently, in a quasi-static growth limit, we expect the system to follow the physics of the two dimensional Ising model with a phase transition occurring at an inverse temperature $J/k_B T = \ln(1 + \sqrt{2})/2 \sim 0.44$. At $J/k_B T = 0.75$, the system is deep within the ordered state so that the expectation is that the structure should be demixed, with large domains of spins of one sign. Remarkably the system undergoes a dynamical phase transition dependent on the growth rate as determined by the chemical potential, from a demixed state in the quasi-static limit to a mixed state at high growth rates.⁷⁴ Example configurations on either side of the transition are shown in Fig. 7. For fixed J , and small $\mu = -0.5$, typical configurations are large domains of mostly spin up or down with equal probability. At $\mu = 0.0$, domains of spin up and spin down of many sizes span the system. For large $\mu = 0.5$, the fast growth results in completely mixed lattice of spin up and spin down, with no long ranged correlations. Because of the kinetic constraint, these nonequilibrium patterns cannot easily anneal.

Depending on which side of the dynamical transition the system is on, the different spatial correlations will be reflected in the distribution of magnetization-per-particle,

$$M(t_N) = \sum_{i=1}^{N(t_N)} s_i, \quad (44)$$

where $N(t_N)$ is the number of added spins $s_i = \pm 1$. At long times, the system relaxes to a nonequilibrium steady state with fixed growth rate $N(t) \propto t_N$. In order to compute the distribution function for $m = M(t_N)/N(t_N)$ in the long time limit, we could bias ensembles of trajectories in TPS or DMC with a field conjugate to m as was done in Secs. III A and III B. However, because of the potential bistability of $p(m)$, such a linear bias would not help sampling near $m = 0$ when $\langle |m| \rangle$ is

FIG. 7. Rare fluctuations of the net magnetization for the MRSOS model computed with TPS (red) and DMC (blue) for $\mu =$ (a) -0.5 , (b) 0.0 , and (c) 0.5 . The top panels show representative configurations of the lattice computed with these parameters.

finite. Instead, we choose to bias the system with a factor of the form

$$p_{\kappa, M^*}(m) = p(m)e^{-\kappa(M-M^*)^2 - t_N \psi(\kappa, M^*)}, \quad (45)$$

where the biased distribution is now a function of two parameters, κ and M^* , as is the normalization constant $\psi(\kappa, M^*)$. In the limit that κ is large and positive, trajectories will be localized near M^* , allowing us to sample arbitrary magnetizations, independent of the geometry of the underlying distribution. A consequence of such a bias is that $\psi(\kappa, M^*)$ is no longer a cumulant generating function, but inverting Eq. (45),

$$\phi(m) = \phi_{\kappa, M^*}(m) - \kappa(M - M^*)^2/t_N - \psi(\kappa, M^*) \quad (46)$$

we still can relate the biased rate function to the unbiased one, where $\phi(m) = -\ln p(m)/t_N$. Such a bias changes the acceptance criteria in TPS to

$$\text{Acc}[\mathcal{C}_0 \rightarrow \mathcal{C}_n] = \text{Min} \left(1, e^{-\Omega} e^{-\eta[\mathcal{C}_n, \kappa, M^*] + \eta[\mathcal{C}_0, \kappa, M^*]} \right), \quad (47)$$

where Ω is the path entropy that depends on the type of MC move and

$$\eta[\mathcal{C}, \kappa, M^*] = \kappa(M[\mathcal{C}] - M^*)^2, \quad (48)$$

where $M[\mathcal{C}]$ is the time integrated value of the magnetization for the trajectory $\mathcal{C}$.

Due to the quadratic bias, the branching probability for DMC has to be altered so that the long time limit adheres to the required functional form of the weight. With a linear bias, the weight is simply multiplicative, so the branching rate can be easily computed. In the non-linear case, this is no longer possible and we must resort to the telescopic form, which requires additional book keeping. For harmonic bias, the DMC branching probability is given by

$$k_i(t) = \frac{\exp[-\kappa(\Delta_i[C_0, C_t] - \Delta_i[C_0, C_{t-\Delta t}])]}{\sum_j^{N_w} \exp[-\kappa(\Delta_j[C_0, C_t] - \Delta_j[C_0, C_{t-\Delta t}])]}, \quad (49)$$

where

$$\Delta_i[C_0, C_t] = (M_i[C_0 \rightarrow C_t] - M^*)^2 \quad (50)$$

is the harmonic bias up to time t . In both TPS and DMC, provided κ and M^* are chosen such that different ensembles of trajectories overlap, we can use histogram reweighting together with the relation in Eq. (46) to reconstruct $\phi(m)$ deep into the tails of the distribution.

Shown in Fig. 7 are TPS and DMC results for the MRSOS model using harmonic biases. Calculations were

accomplished for $\mu = -0.5, 0.0, 0.5$ and $J = 0.75$ using system sizes that were approximately $L = 40$ and $t_N = 200$. The TPS calculations were run with $\kappa = 2$ and 40 windows, with M^*/t_N chosen to cover $m = -1, 1$, and trajectories of 200 MC steps were sampled with shifting and shooting moves. In all, 10^5 trajectories were collected for each μ . The DMC calculations were run using $N_w = 10\,000$ and $\Delta t = 1$. For all cases a spring constant, $\kappa = 2$, was used and 40 values of M^* were spread uniformly across the range of $m = -1, 1$.

For all cases studied, the DMC and TPS results are statistically identical. For large $\mu = 0.5$ as expected, the system exhibits a single minimum, centered at $m = 0$, with gaussian fluctuations about the average. At intermediate driving, $\mu = 0$, the system exhibits one broad minimum, with fluctuations that go like $\phi(m) \propto m^6$ for small m . The concomitant susceptibility computed under such conditions is divergent in the large system, long time limit, as expected for a critical transition. For small $\mu = -0.5$, we find the system exhibits two minima, each stable basin centered at finite values of m that are equal and opposite as required by symmetry.⁷⁷ If the non-convexity of $\phi(m)$ is related to a surface tension in spacetime, we would expect that in the limit of large t_N and large L the rate function becomes flat about $m = 0$. We indeed find this to be the case, and further find that depending on the ratio of t_N and L , interfaces can preferentially form in the smaller of the spatial or temporal directions, yielding a second metastable minimum about $m = 0$, consistent with recent calculations.⁷⁵

While generalizing the calculations beyond linear biases allows TPS and DMC to sample non-convex rate functions and study a complex model of self-assembly, the system sizes and times are still restrictive. Indeed without additional importance sampling, performing finite size scaling to determine the critical exponents for the transition is prohibitively expensive.

IV. CONCLUSION

We have shown how large deviation functions within nonequilibrium steady states can be computed with trajectory based importance sampling. In particular, we have studied the systematic and statistical errors associated with two Monte Carlo algorithms, TPS and DMC, to estimate large deviation functions. We found that though they are implemented in different ways, the two algorithms sample the same ensemble of trajectories and suffer from related difficulties.

The unifying feature of these two algorithms is that they augment the propagation of the bare system dynamics with *a posteriori* importance sampling. In the case of TPS, this is done with a Metropolis acceptance step after proposing a new trajectory from an existing one. In the case of DMC, the bare dynamics are augmented with a population dynamics, whereby walkers are cloned or deleted after short trajectories are generated. In both cases, the proposal step of the Monte Carlo chain is done without any guidance from the rare events that contribute to the large deviation function. As a result, for exponentially rarer fluctuations, the Monte Carlo chain requires exponentially more samples of the targeted stationary distribution as the overlap between it and the proposed distribution becomes exponentially small. This means that in practice low dimensional systems, or small fluctuations, can be systematically converged, but rare fluctuations of high dimensional systems are impractical to compute. We have illustrated this with a number of simple models where exact numerical calculations can be done, on and off lattice. By making simple adjustments to the algorithm and using the standard histogram reweighting tools, we could push the algorithm to study the MRSOS model. This model exhibits a complex phase diagram, including a critical point. However, systematic studies on the finite size scaling properties are outside the reach of these current algorithms.

These results highlight a potential avenue for further development, where guiding functions are incorporated into the proposal step, in analogy to what is done in quantum diffusion Monte Carlo.⁶¹ To do this in the context of nonequilibrium steady states requires developing auxiliary dynamics that can push the trajectories to the rare fluctuations that contribute to the large deviation function. In principle, for Markovian dynamics, an optimal auxiliary process always exists and can be constructed by a canonical transformation known as the Doob transform.⁴⁸ However, to construct such a process exactly requires finding the large deviation function. In Paper II of this work, we construct approximations to the Doob transform and use those as guiding functions for Monte Carlo sampling⁴⁵ and show that orders of magnitude in statistical efficiency can be gained by relatively simple approximate guiding functions.

ACKNOWLEDGMENTS

D.T.L. was supported by UC Berkeley College of Chemistry. U.R. was supported by the Simons Collaboration on the Many-Electron Problem and the California Institute of Technology. G.K.-L.C. is a Simons Investigator in Theoretical Physics and was supported by the California Institute of Technology.

¹G. Torrie and J. Valleau, "Nonphysical sampling distributions in Monte Carlo free-energy estimation: Umbrella sampling," *J. Comput. Phys.* **23**, 187–199 (1977).

²A. Laio and M. Parrinello, "Escaping free-energy minima," *Proc. Natl. Acad. Sci. U. S. A.* **99**, 12562–12566 (2002).

³F. Wang and D. P. Landau, "Efficient, multiple-range random walk algorithm to calculate the density of states," *Phys. Rev. Lett.* **86**, 2050–2053 (2001).

⁴C. Chipot and A. Pohorille, *Free Energy Calculations* (Springer, 2007).

⁵M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Oxford University Press, 1989).

⁶D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications*, Computational Science Series (Elsevier Science, 2001).

⁷D. P. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University Press, 2014).

⁸H. Touchette, "The large deviation approach to statistical mechanics," *Phys. Rep.* **478**, 1–69 (2009).

⁹M. Merolle, J. P. Garrahan, and D. Chandler, "Space-time thermodynamics of the glass transition," *Proc. Natl. Acad. Sci. U. S. A.* **102**, 10837–10840 (2005).

¹⁰C. Valeriani, R. J. Allen, M. J. Morelli, D. Frenkel, and P. Rein ten Wolde, "Computing stationary distributions in equilibrium and nonequilibrium systems with forward flux sampling," *J. Chem. Phys.* **127**, 114109 (2007).

¹¹C. Giardinà, J. Kurchan, and L. Peliti, "Direct evaluation of large-deviation functions," *Phys. Rev. Lett.* **96**, 120603 (2006).

¹²M. Tchernookov and A. R. Dinner, "A list-based algorithm for evaluation of large deviation functions," *J. Stat. Mech.: Theory Exp.* **2010**, P02006.

¹³T. Nemoto and S.-i. Sasa, "Computation of large deviation statistics via iterative measurement-and-feedback procedure," *Phys. Rev. Lett.* **112**, 090602 (2014).

¹⁴T. Nemoto, E. G. Hidalgo, and V. Lecomte, "Finite-time and finite-size scalings in the evaluation of large-deviation functions: Analytical study using a birth-death process," *Phys. Rev. E* **95**, 012102 (2017).

¹⁵E. G. Hidalgo, T. Nemoto, and V. Lecomte, "Finite-time and -size scalings in the evaluation of large deviation functions: Numerical approach in continuous time," *Phys. Rev. E* **95**, 062134 (2017).

¹⁶T. Nemoto, F. Bouchet, R. L. Jack, and V. Lecomte, "Population-dynamics method with a multicanonical feedback control," *Phys. Rev. E* **93**, 062123 (2016).

¹⁷P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, "Transition path sampling: Throwing ropes over rough mountain passes, in the dark," *Annu. Rev. Phys. Chem.* **53**, 291–318 (2002).

¹⁸C. Giardinà, J. Kurchan, V. Lecomte, and J. Tailleur, "Simulating rare events in dynamical processes," *J. Stat. Phys.* **145**, 787–811 (2011).

¹⁹C. Jarzynski, "Nonequilibrium equality for free energy differences," *Phys. Rev. Lett.* **78**, 2690–2693 (1997).

²⁰G. E. Crooks, "Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences," *Phys. Rev. E* **60**, 2721–2726 (1999).

²¹J. L. Lebowitz and H. Spohn, "A Gallavotti–Cohen-type symmetry in the large deviation functional for stochastic dynamics," *J. Stat. Phys.* **95**, 333–365 (1999).

²²G. Gallavotti and E. G. D. Cohen, "Dynamical ensembles in nonequilibrium statistical mechanics," *Phys. Rev. Lett.* **74**, 2694–2697 (1995).

²³S. Auer and D. Frenkel, "Prediction of absolute crystal-nucleation rate in hard-sphere colloids," *Nature* **409**, 1020 (2000).

²⁴S. Auer and D. Frenkel, "Crystallization of weakly charged colloidal spheres: A numerical study," *J. Phys.: Condens. Matter* **14**, 7667 (2002).

²⁵W. Ren, E. Vanden-Eijnden, P. Maragakis, and W. E, "Transition pathways in complex systems: Application of the finite-temperature string method to the alanine dipeptide," *J. Chem. Phys.* **123**, 134109 (2005).

²⁶M. Berhanu, R. Monchaux, S. Fauve, N. Mordant, F. Pétrélis, A. Chiffau, F. Daviaud, B. Dubrulle, L. Marié, F. Ravelet, M. Bourgoin, P. Odier, J.-F. Pinton, and R. Volk, "Magnetic field reversals in an experimental turbulent dynamo," *Europhys. Lett.* **77**, 59001 (2007).

²⁷F. Bouchet and E. Simonnet, "Random changes of flow topology in two-dimensional and geophysical turbulence," *Phys. Rev. Lett.* **102**, 094504 (2009).

²⁸D. R. Easterling, G. A. Meehl, C. Parmesan, S. A. Changnon, T. R. Karl, and L. O. Mearns, "Climate extremes: Observations, modeling, and impacts," *Science* **289**, 2068–2074 (2000).

²⁹J. K. Weber, R. L. Jack, and V. S. Pande, "Emergence of glass-like behavior in Markov state models of protein folding dynamics," *J. Am. Chem. Soc.* **135**, 5501–5504 (2013).

³⁰C. Dellago, P. G. Bolhuis, and D. Chandler, "Efficient transition path sampling: Application to Lennard-Jones cluster rearrangements," *J. Chem. Phys.* **108**, 9236–9245 (1998).

³¹T. S. Van Erp and P. G. Bolhuis, "Elaborating transition interface sampling methods," *J. Comput. Phys.* **205**, 157–181 (2005).

³²E. Weinan, W. Ren, and E. Vanden-Eijnden, "Transition pathways in complex systems: Reaction coordinates, isocommittor surfaces, and transition tubes," *Chem. Phys. Lett.* **413**, 242–247 (2005).

³³R. J. Allen, C. Valeriani, and P. R. ten Wolde, "Forward flux sampling for rare event simulations," *J. Phys.: Condens. Matter* **21**, 463102 (2009).

- ³⁴P. Bolhuis and C. Dellago, "Practical and conceptual path sampling issues," *Eur. Phys. J.: Spec. Top.* **224**, 2409–2427 (2015).
- ³⁵L. O. Hedges, R. L. Jack, J. P. Garrahan, and D. Chandler, "Dynamic order-disorder in atomistic models of structural glass formers," *Science* **323**, 1309–1313 (2009).
- ³⁶T. Speck and D. Chandler, "Constrained dynamics of localized excitations causes a non-equilibrium phase transition in an atomistic model of glass formers," *J. Chem. Phys.* **136**, 184509 (2012).
- ³⁷D. T. Limmer and D. Chandler, "Theory of amorphous ices," *Proc. Natl. Acad. Sci. U. S. A.* **111**, 9413–9418 (2014).
- ³⁸A. S. Mey, P. L. Geissler, and J. P. Garrahan, "Rare-event trajectory ensemble analysis reveals metastable dynamical phases in lattice proteins," *Phys. Rev. E* **89**, 032109 (2014).
- ³⁹G. E. Crooks and D. Chandler, "Efficient transition path sampling for nonequilibrium stochastic dynamics," *Phys. Rev. E* **64**, 026109 (2001).
- ⁴⁰T. R. Gingrich, G. M. Rotskoff, G. E. Crooks, and P. L. Geissler, "Near-optimal protocols in complex nonequilibrium transformations," *Proc. Natl. Acad. Sci. U. S. A.* **113**, 10263 (2016).
- ⁴¹R. J. Allen, P. B. Warren, and P. R. Ten Wolde, "Sampling rare switching events in biochemical networks," *Phys. Rev. Lett.* **94**, 018104 (2005).
- ⁴²P. I. Hurtado and P. L. Garrido, "Large fluctuations of the macroscopic current in diffusive systems: A numerical test of the additivity principle," *Phys. Rev. E* **81**, 041102 (2010).
- ⁴³T. Speck and J. P. Garrahan, "Space-time phase transitions in driven kinetically constrained lattice models," *Eur. Phys. J. B* **79**, 1–6 (2011).
- ⁴⁴P. I. Hurtado, C. P. Espigares, J. J. del Pozo, and P. L. Garrido, "Thermodynamics of currents in nonequilibrium diffusive systems: Theory and simulation," *J. Stat. Phys.* **154**, 214–264 (2014).
- ⁴⁵U. Ray, G. K.-L. Chan, and D. T. Limmer, "Importance sampling large deviations in nonequilibrium steady states: Part II," *J. Chem. Phys.* (unpublished).
- ⁴⁶A. A. Budini, R. M. Turner, and J. P. Garrahan, "Fluctuating observation time ensembles in the thermodynamics of trajectories," *J. Stat. Mech.: Theory Exp.* **2014**, P03012.
- ⁴⁷N. Van Kampen, *Stochastic Processes in Physics and Chemistry*, North-Holland Personal Library (Elsevier Science, 1992).
- ⁴⁸R. Chetrite and H. Touchette, "Nonequilibrium Markov processes conditioned on large deviations," *Ann. Henri Poincaré* **16**, 2005–2057 (2015).
- ⁴⁹D. Chandler, *Introduction to Modern Statistical Mechanics* (Oxford University Press, 1987), p. 288.
- ⁵⁰T. Speck, A. Malins, and C. P. Royall, "First-order phase transition in a model glass former: Coupling of local structure and dynamics," *Phys. Rev. Lett.* **109**, 195703 (2012).
- ⁵¹T. R. Gingrich and P. L. Geissler, "Preserving correlations between trajectories for efficient path sampling," *J. Chem. Phys.* **142**, 234104 (2015).
- ⁵²T. S. Van Erp, "Dynamical rare event simulation techniques for equilibrium and nonequilibrium systems," *Adv. Chem. Phys.* **151**, 27 (2012).
- ⁵³S. Kumar, J. M. Rosenberg, D. Bouzida, R. H. Swendsen, and P. A. Kollman, "The weighted histogram analysis method for free-energy calculations on biomolecules. I. The method," *J. Comput. Chem.* **13**, 1011–1021 (1992).
- ⁵⁴M. R. Shirts and J. D. Chodera, "Statistically optimal analysis of samples from multiple equilibrium states," *J. Chem. Phys.* **129**, 124105 (2008).
- ⁵⁵P. Grassberger, "Go with the winners: A general Monte Carlo strategy," *Comput. Phys. Commun.* **147**, 64–70 (2002).
- ⁵⁶J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, "First-order dynamical phase transition in models of glasses: An approach based on ensembles of histories," *J. Phys. A: Math. Theor.* **42**, 075007 (2009).
- ⁵⁷T. Bodineau, V. Lecomte, and C. Toninelli, "Finite size scaling of the dynamical free-energy in a kinetically constrained model," *J. Stat. Phys.* **147**, 1–17 (2012).
- ⁵⁸E. Pitard, V. Lecomte, and F. Van Wijland, "Dynamic transition in an atomic glass former: A molecular-dynamics evidence," *Europhys. Lett.* **96**, 56002 (2011).
- ⁵⁹P. I. Hurtado and P. L. Garrido, "Spontaneous symmetry breaking at the fluctuating level," *Phys. Rev. Lett.* **107**, 180601 (2011).
- ⁶⁰T. Mitsudo and S. Takesue, "Numerical estimation of the current large deviation function in the asymmetric simple exclusion process with open boundary conditions," *J. Phys. Soc. Jpn.* **80**, 114001 (2011).
- ⁶¹P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester, "Fixed-node quantum Monte Carlo for molecules," *J. Chem. Phys.* **77**, 5593–5603 (1982).
- ⁶²C. J. Umrigar, M. P. Nightingale, and K. J. Runge, "A diffusion Monte Carlo algorithm with very small time-step errors," *J. Chem. Phys.* **99**, 2865–2890 (1993).
- ⁶³P. Reimann, C. Van den Broeck, H. Linke, P. Hänggi, J. Rubi, and A. Pérez-Madrid, "Giant acceleration of free diffusion by use of tilted periodic potentials," *Phys. Rev. Lett.* **87**, 010602 (2001).
- ⁶⁴J. Mehl, T. Speck, and U. Seifert, "Large deviation function for entropy production in driven one-dimensional systems," *Phys. Rev. E* **78**, 011123 (2008).
- ⁶⁵P. Tsobgni Nyawo and H. Touchette, "Large deviations of the current for driven periodic diffusions," *Phys. Rev. E* **94**, 032101 (2016).
- ⁶⁶A. Braňka and D. M. Heyes, "Algorithms for Brownian dynamics computer simulations: Multivariable case," *Phys. Rev. E* **60**, 2381 (1999).
- ⁶⁷J. Kurchan, "Fluctuation theorem for stochastic dynamics," *J. Phys. A: Math. Gen.* **31**, 3719 (1998).
- ⁶⁸T. Nemoto, E. G. Hidalgo, and V. Lecomte, "Finite-time and finite-size scalings in the evaluation of large-deviation functions: Analytical study using a birth-death process," *Phys. Rev. E* **95**, 012102 (2017).
- ⁶⁹G. M. Schütz, *Integrable Stochastic Many-Body Systems* (Forschungszentrum Jülich, Zentralbibliothek, 1998).
- ⁷⁰A. B. Kolomeisky, G. M. Schütz, E. B. Kolomeisky, and J. P. Straley, "Phase diagram of one-dimensional driven lattice gases with open boundaries," *J. Phys. A: Math. Gen.* **31**, 6911 (1998).
- ⁷¹A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, "A new algorithm for Monte Carlo simulation of Ising spin systems," *J. Comput. Phys.* **17**, 10–18 (1975).
- ⁷²M. Esposito and C. Van den Broeck, "Three faces of the second law. I. Master equation formulation," *Phys. Rev. E* **82**, 011143 (2010).
- ⁷³R. Jullien and R. Botet, "Scaling properties of the surface of the Eden model in $d = 2, 3, 4$," *J. Phys. A: Math. Gen.* **18**, 2279 (1985).
- ⁷⁴S. Whitelam, L. O. Hedges, and J. D. Schmit, "Self-assembly at a nonequilibrium critical point," *Phys. Rev. Lett.* **112**, 155504 (2014).
- ⁷⁵K. Klymko, P. L. Geissler, J. P. Garrahan, and W. Stephen, "Rare behavior of growth processes via umbrella sampling of trajectories," e-print [arXiv:1707.00767](https://arxiv.org/abs/1707.00767) (2017).
- ⁷⁶J. M. Kim and J. Kosterlitz, "Growth in a restricted solid-on-solid model," *Phys. Rev. Lett.* **62**, 2289 (1989).
- ⁷⁷M. Nguyen and S. Vaikuntanathan, "Design principles for nonequilibrium self-assembly," *Proc. Natl. Acad. Sci. U. S. A.* **113**, 14231–14236 (2016).