Population-dynamics method with a multicanonical feedback control

Takahiro Nemoto,[1,2] Freddy Bouchet,[2] Robert L. Jack,[3] and Vivien Lecomte[1]
[1]*Laboratoire de Probabilités et Modèles Aléatoires, Sorbonne Paris Cité, UMR 7599 CNRS, Université Paris Diderot, 75013 Paris, France*
[2]*Laboratoire de Physique, ENS de Lyon, Université de Lyon, CNRS, 46 allée d'Italie, 69364 Lyon, France*
[3]*Department of Physics, University of Bath, Bath BA2 7AY, United Kingdom*
(Received 25 January 2016; revised manuscript received 5 April 2016; published 15 June 2016)

We discuss the Giardinà-Kurchan-Peliti population dynamics method for evaluating large deviations of time-averaged quantities in Markov processes [Phys. Rev. Lett. **96**, 120603 (2006)]. This method exhibits systematic errors which can be large in some circumstances, particularly for systems with weak noise, with many degrees of freedom, or close to dynamical phase transitions. We show how these errors can be mitigated by introducing control forces within the algorithm. These forces are determined by an iteration-and-feedback scheme, inspired by multicanonical methods in equilibrium sampling. We demonstrate substantially improved results in a simple model, and we discuss potential applications to more complex systems.

DOI: 10.1103/PhysRevE.93.062123

I. INTRODUCTION

In many physical systems, interesting and important behavior is associated with rare events; examples include crystal nucleation, slow transitions in biomolecules [1–3], rare transitions in turbulent flows [4,5], and extreme events in climate dynamics [6]. Many computational methods for sampling these events have been proposed and exploited [1,3,5,7–18]. One family of methods is based around population dynamics [19–24], in which several copies of a system evolve in parallel: the copies which exhibit the rare behavior of interest are copied (or *cloned*) while other copies are discarded. The result is that *typical* copies within the population dynamics reproduce the desired rare events in the original system. One such method has recently been employed to characterize a particular class of rare events [7,8], in which time-averaged physical quantities exhibit *large deviations* [25,26] from their typical values in the large time limit. Studies of such events have revealed new and unexpected features in glass formers [27], biomolecules [28–30], nonequilibrium transport [31,32], and integrable systems [8]. In this article, we identify a pitfall that limits the computational efficiency of the population dynamics method, and we show that the method can be modified so as to avoid this problem. The issue at stake is the number of copies of the system that must be considered in order to obtain accurate results; if very many copies are required, then the method is difficult to apply, especially if even a single system is complex or contains many degrees of freedom. In some relevant cases then the standard population dynamics method requires an exponentially large population to be effective [33]. However, the method that we propose here, which is an improved version of the population dynamics, inspired by multicanonical methods in equilibrium systems [13,14] (or adaptive importance sampling [15–18]), can still be effective in these cases.

The intuitive description of the problem that we identify is the following. The population dynamics is characterized by two different distributions, which describe the state of the system at some fixed final time and its state at intermediate times. We show that in situations where the two distributions have a small overlap, the population dynamics is affected by a serious sampling problem, in which statistical estimators of the quantities of interest become dominated by just a few samples. One relevant case is that of systems with weak noise, for which the two distributions become more and more concentrated around their most likely values, so that they quite generally have zero overlap: this leads to an unavoidable failure of the population dynamics. In this article we describe how to modify the population dynamics so as to maintain the two distributions close to each other, thus solving the sampling problem. We argue that this new method will provide a step change in the complexity of the systems for which large deviation computations can be performed.

The structure of the paper is as follows: we introduce our model and the population dynamics algorithm in Sec. II. We discuss sampling problems associated with this algorithm in Sec. III. In Sec. IV we introduce our main idea, which is to combine a controlling force with the population dynamics algorithm, in order to resolve the sampling issues. In Sec. V we numerically demonstrate this method in a simple Brownian particle model. Finally, in Sec. VI we describe the potential for future applications and extensions of this work.

II. MODEL AND METHODS

A. Rare event problem

The rare events that we consider can take place in a variety of models. To illustrate the method, consider a particle moving in d dimensions, whose position $x \in \mathbf{R}^d$ obeys a Langevin equation

$$\dot{x}_t = F(x_t) + B(x_t)\xi_t, \tag{1}$$

where ξ is a d-dimensional Gaussian white noise of unit variance, $F(x) \in \mathbf{R}^d$ a deterministic force, and the matrix $B(x)$ specifies the action of the noise on the particle. [34]. We use the Itô convention [35] for stochastic calculus throughout this paper, although one can also work with the Stratonovich convention by using a transformation formula to relate one convention to the other [36].

We restrict to ergodic systems, and we focus on rare events in which a time-averaged quantity $\Lambda(\tau)$ takes some nontypical value. Here τ is the long time period over which the average

is taken, and

$$\Lambda(\tau) = \Lambda_{\rm d}(\tau) + \Lambda_{\rm c}(\tau) \tag{2}$$

consists of a "scalar" contribution

$$\Lambda_{\rm d}(\tau) = \frac{1}{\tau}\int_0^\tau \lambda_{\rm d}(x_t)\,dt \tag{3}$$

and a "vector" one

$$\Lambda_{\rm c}(\tau) = \frac{1}{\tau}\int_0^\tau \lambda_{\rm c}(x_t)\cdot dx_t, \tag{4}$$

where $\lambda_{\rm d,c}$ are arbitrary functions of the particle position x. The first contribution $\Lambda_{\rm d}(\tau)$ is a time average of a quantity $\lambda_{\rm d}$ that depends only on the position x (i.e., a time average of a static function such as a particle density or an energy density), whereas the second contribution $\Lambda_{\rm c}(\tau)$ includes transitions of x as seen from the form $\lambda_{\rm c}(x_t)\cdot dx_t$ [i.e., $\Lambda_{\rm c}(\tau)$ is an average of a dynamic function such as a particle current or an energy current [37]]. See also the explanation around Eq. (34) in Ref. [38] for a pedagogical introduction of $\Lambda(\tau)$. This class of observable includes many physically and mathematically interesting examples, and fluctuations of these quantities have been intensively studied recently, where examples are entropy production [39,40], dynamical activity [27,41], and particle fluxes [42].

In the limit of large τ, ergodicity of the system means that the observable $\Lambda(\tau)$ is almost surely equal to its typical value $\overline{\Lambda}$. Our aims are (i) to estimate the (small) probability of deviations from this value and (ii) to generate the rare trajectories that lead to these deviations. This is an important problem because these nontypical trajectories can exhibit interesting and unusual structures, including misfolded proteins [29,30], stable glass states [27], and traveling waves in models of particle transport [31].

To achieve these aims, the standard theoretical route [25,39] is to introduce a *biasing field* h, which controls deviations of $\Lambda(\tau)$ from its typical value. Specifically, we consider an ensemble of paths $X = (x_t)_{t=0}^\tau$ with (unnormalized) probability density

$$P_h[X] = \pi_0(x_0)\exp\left[-\int_0^\tau \mathcal{L}(x_t,\dot{x}_t)\,dt + h\tau\Lambda(\tau)\right], \tag{5}$$

where

$$\mathcal{L}(x,\dot{x}) = \tfrac{1}{2}[\dot{x} - F(x)]\cdot\kappa(x)^{-1}[\dot{x} - F(x)] \tag{6}$$

is a Lagrangian density that describes the (unbiased) model (1); $\pi_0(x)$ is the initial condition for the trajectories, which can be arbitrary and which we take to be the stationary probability distribution of the unbiased model in the numerical examples. Also, $\kappa = BB^T$ where the notation B^T indicates a matrix transpose [43].

Normalised averages with respect to P_h are denoted by $\langle\cdot\rangle_h$, and we use these averages to characterize the rare trajectories associated with deviations of $\Lambda(\tau)$ from $\overline{\Lambda}$, for the model in Eq. (1). We define the scaled cumulant generating function (CGF):

$$G(h) = \lim_{\tau\to\infty}\tau^{-1}\log\langle {\rm e}^{\tau h\Lambda(\tau)}\rangle_0. \tag{7}$$

In the limit of large τ, the probability distribution of $\Lambda(\tau)$ satisfies a large deviation principle, and can be obtained by a Legendre transformation of $G(h)$ [for which we assume that the large deviation function of $\Lambda(\tau)$ is convex [25,26]]. In the same limit, for a given deviation Λ from $\overline{\Lambda}$, there exists a bias $h^\star(\Lambda)$ for which $\langle\cdot\rangle_{h^\star(\Lambda)}$ is equivalent to a conditional average over trajectories with $\Lambda(\tau) = \Lambda$ [44]. Biased averages with respect to the biased distribution P_h, which are numerically evaluated through the population dynamics, thus enable us to characterize the trajectories of the original dynamics for which time-averaged physical quantities exhibit large deviations from their typical values in the large time limit.

B. Population dynamics method

There are several computational methods that allow evaluation of averages with respect to P_h [7,11,12,45]. In the population dynamics method [7], one considers N_c copies (or clones) of the system. These clones evolve independently as a function of the time t, except that (for $h > 0$) clones with small $\Lambda(t)$ are periodically removed (eliminated) from the system, while clones with large $\Lambda(t)$ are duplicated (cloned), to maintain a constant population. The algorithm is illustrated in Fig. 1 and described fully in Appendix A 1. This method biases the dynamics towards the rare events of interest. For sufficiently large N_c (and large enough τ), the method provides accurate estimates of $G(h)$, and it generates sample paths consistent with the biased distribution P_h.

C. Numerical example

To show the operation of the population dynamics method, we introduce a simple model of diffusion in a quartic potential; that is, $F(x) = -x^3$ and $B(x) = \sqrt{2\epsilon}$, where ϵ is the noise strength (or temperature). We take $\lambda_{\rm c} = 0$ and $\lambda_{\rm d} = x(x+1)$. For $h < 0$ the distribution P_h is concentrated on trajectories with small values of $\lambda_{\rm d}$, which tend to localize near $x \approx -\frac{1}{2}$. Here we focus on the case $h > 0$, which leads to unusually large values of $\lambda_{\rm d}$. Those can be realized either for $x > 0$ or $x < 0$, but at large τ this rare event is almost always realized by trajectories that have $x > 0$ (as illustrated in Fig. 1). This simple problem can be solved exactly in the zero-noise limit (see Appendix D).

The operation of the population dynamics method is illustrated in Fig. 1. Figure 1(a) shows four copies of the system that evolve in time, except that some trajectory segments stop and others branch, as the cloning operates. Figure 1(b) shows four representative trajectories (sample paths) for the distribution $P_h[X]$, which have been reconstructed from Fig. 1(a), by tracing backwards in time from the clones that survived up to the final time τ.

III. SAMPLING ERRORS WITHIN POPULATION DYNAMICS

A. Distributions $p_{\rm end}$ and $p_{\rm ave}$

The accuracy of the population dynamics is limited by the number of clones N_c used in its numerical implementation, as we now explain. Consider the distribution

$$p_{\rm ave}(x) = \lim_{\tau\to\infty}\left\langle \tau^{-1}\int_0^\tau \delta(x_t - x)\,dt\right\rangle_h, \tag{8}$$

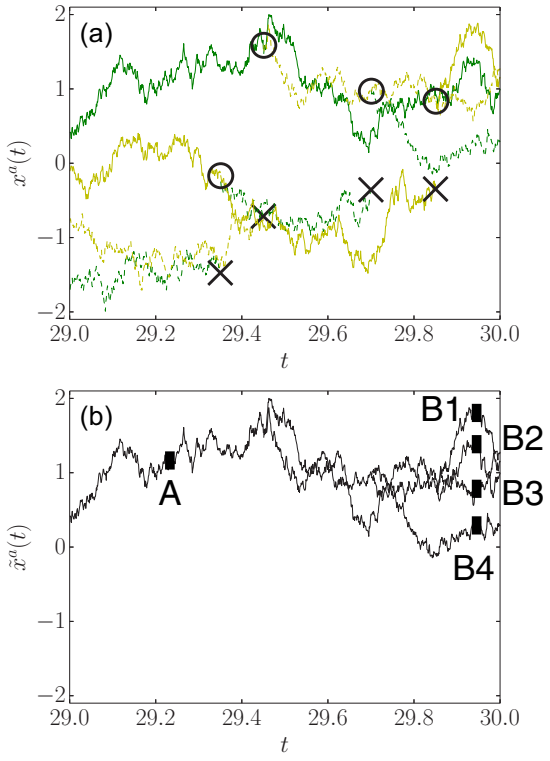

FIG. 1. (a) Trajectories $x^a(t)$ generated by population dynamics at fixed total population $N_c = 4$ for the model system described in Sec. II C ($\epsilon = 1, h = 1$). The different colors or line types, the green (dark gray) solid line, green dashed line, yellow (light gray) solid line, and yellow dashed line, represent different copies. At certain times, some copies of the system are removed ($\times$ symbols) and others are duplicated ($\circ$ symbols). The time interval ΔT for the cloning procedure is set to be 0.05, and the time step for solving the Langevin equation is 0.001 (see Appendix A 1 for the details of the algorithm). (b) Representative sample paths $\tilde{x}^a(t)$ for the distribution $P_h[X]$, derived from those in (a) by keeping only trajectories surviving up to final time $\tau = 30$. For each cloning event, we also copy the history of the trajectory, which replaces the history of the eliminated trajectory. This means that the trajectories $(\tilde{x}^a(t))_{a=1}^{N_c}$ overlap, especially for early times. For example, in panel (b), the point A appears in the past of the four points B1, . . . , B4. For any point $x^a(t)$ (such as A, B1, B2, . . .), we define the multiplicity $m_a(t,\tau)$ as the number of trajectories that include this point and survive until the final time τ. So for point A, the multiplicity is $m_a(t,\tau) = 4$, but for B1, . . . , B4 then $m_a(t,\tau) = 1$. [For all points in the trajectories that die before τ, which are not drawn in panel (b), $m_a(t,\tau) = 0$.]

which indicates the fraction of time spent at position x, within the biased ensemble. We also define

$$p_{\rm end}(x) = \lim_{\tau\to\infty} \langle \delta(x_\tau - x)\rangle_h, \tag{9}$$

which indicates the fraction of trajectories for which the particle's final position is x. For the stationary state of the dynamics (1), which corresponds to $h = 0$, time-translational invariance ensures that $p_{\rm ave} = p_{\rm end}$. However, this is not the case for biased ensembles where $h \neq 0$, as illustrated in Refs. [7,46] and in Fig. 2.

The population dynamics method provides estimates of both $p_{\rm ave}$ and $p_{\rm end}$. Let the position of clone a at time t be $x^a(t)$, with $a = 1, \ldots, N_c$. Recalling Fig. 1(a), note that the functions $x^a(t)$ are not continuous in time and do not represent sample paths for the distribution $P_h[X]$. However, from the definition of the population dynamics algorithm (as explained in Appendix A 1), the distribution of $x^a(t)$ can be used to estimate $p_{\rm end}(x)$, as

$$p_{\rm end}(x) \simeq \frac{1}{\tau N_c} \int_0^\tau \sum_{a=1}^{N_c} \delta[x - x^a(t)]\, dt. \tag{10}$$

In order to construct sample paths, which we denote by $\tilde{x}^a(t)$, we trace backwards in time from the clones that survive up to τ, as shown in Fig. 1(b). There are still N_c paths $\tilde{x}^a$, but these overlap, particularly at early times. Since these trajectories correspond to $P_h[X]$, the distribution of $\tilde{x}^a$ gives an estimate of $p_{\rm ave}(x)$, as:

$$p_{\rm ave}(x) \simeq \frac{1}{\tau N_c} \int_0^\tau \sum_{a=1}^{N_c} \delta[x - \tilde{x}^a(t)]\, dt. \tag{11}$$

The approximate equalities in the relations (10) and (11) become exact in the limit $N_c \to \infty$ and $\tau \to \infty$, in which the population dynamics gives exact results.

We show numerical examples of these functions in Fig. 2, for a particle moving in a quartic potential, as introduced in Sec. II C. We estimate $p_{\rm ave}$ and $p_{\rm end}$ from (10) and (11), and show them in Fig. 2. In the same figure, we also plot the numerically exact distributions, obtained by numerical solution of a modified Fokker-Planck equation (see Ref. [25] and Appendix B 2). The population dynamics converges to the exact result as N_c is increased. Also shown in Fig. 2 are results using the control-with-feedback method that we introduce in this paper; these results will be discussed in later sections.

B. Multiplicity

The population dynamics method gives accurate results in the limit of large N_c. The central idea is that in a large population, short-lived rare fluctuations will occur. Based on these short-lived fluctuations, we duplicate some of the clones: repeated application of this procedure generates the *long-lived* fluctuations that are relevant for large deviation theory. For this to be effective, the population on which the cloning operates must be large enough to capture the relevant short-lived fluctuations; that is, the cloning part of the algorithm can allocate extra statistical weight to configurations that are already present in the population, but new configurations are generated only by the natural (unbiased) dynamics of the system.

Assuming that N_c is large enough for efficient operation of the algorithm, the configurations that are associated with long-lived dynamical fluctuations are distributed as $p_{\rm ave}$, but the cloning step operates on a population distributed as $p_{\rm end}$. From the argument above, it is clear that if *typical* samples from $p_{\rm ave}$ are rare with respect to $p_{\rm end}$, then a large population is required in order to obtain accurate results. To quantify this, it is useful to define the multiplicity $m_a(t,\tau)$ of clone a at time t as the number of its descendants that survive until the final time τ (see Fig. 1). Rewriting (8) as $p_{\rm ave}(x) \simeq \frac{1}{\tau N_c} \int_0^\tau \sum_{a=1}^{N_c} m_a(t,\tau)\delta[x - x^a(t)]\, dt$ and comparing with (9), one sees that for a clone with position $x = x^a(t)$, the *expected*

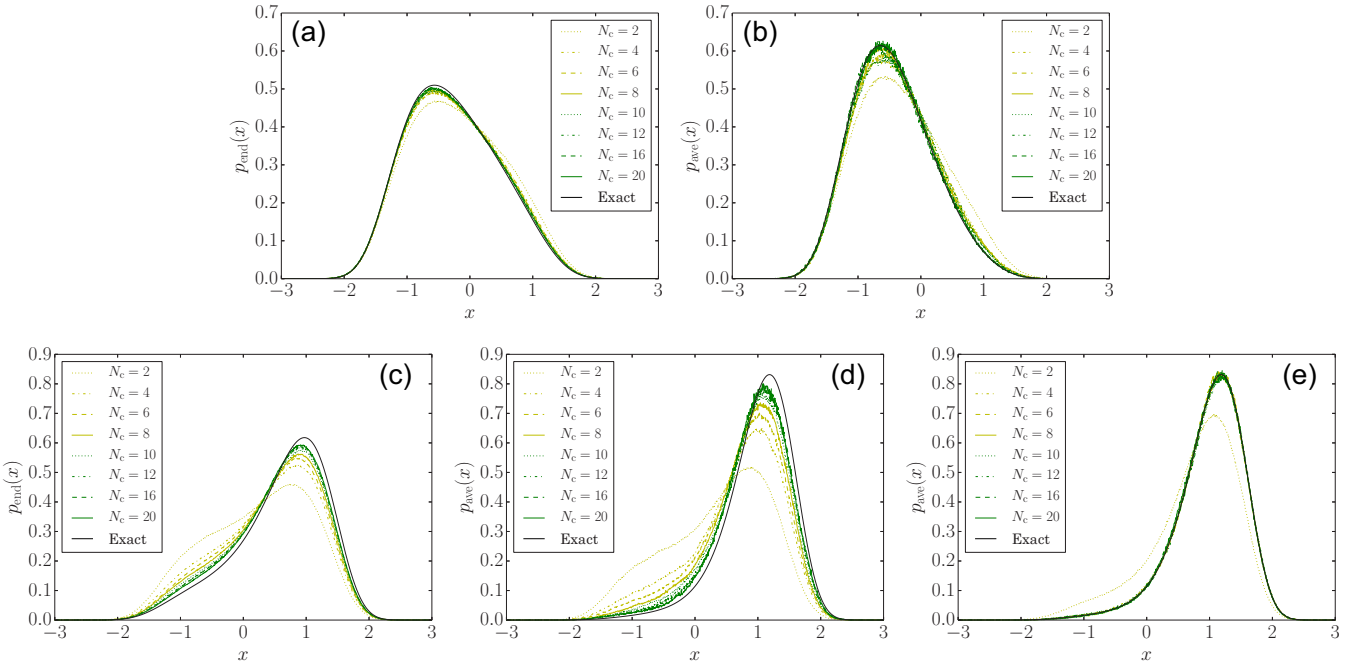

FIG. 2. (a–d) Distributions $p_{\rm end}(x)$ and $p_{\rm ave}(x)$, defined in (8) and (9), calculated from the population dynamics method, with various numbers of clones N_c. The different panels correspond to a different value of h ($h = \pm 1$) or a different distribution function [$p_{\rm end}(x)$ or $p_{\rm ave}(x)$]: (a) $p_{\rm end}(x)$ for $h = -1$, (b) $p_{\rm ave}(x)$ for $h = -1$, (c) $p_{\rm end}(x)$ for $h = 1$, and (d) $p_{\rm ave}(x)$ for $h = 1$. For all panels, we set $\epsilon = 1$ and $\tau = 30$. The numerically exact result is plotted as a black line. We repeat the simulation $1200/N_c$ times, and the result is the average of them (this procedure means that we vary N_c while keeping a fixed computational cost). The results of the population dynamics converge to the analytical ones as N_c increases. (e) $p_{\rm ave}(x)$ for $h = 1$ (improved estimation) calculated from a population dynamics method with control-with-feedback, as described in Secs. IV C and V. Results are shown after two iterations of the feedback procedure. The exact distribution $p_{\rm ave}(x)$ is again shown as a black line. The comparison between (d) and (e) indicates that the convergence with respect to N_c is improved significantly by the control-with-feedback method. The variance m_2 and the relative entropy D_2 defined in (12) and (13) both measure how much large values of N_c are required for the cloning procedure to be reliable. For panel (b), (d) and (e), these values are ($m_2 = 0.068$, $D_2 = 0.039$), ($m_2 = 0.33$, $D_2 = 0.17$), and ($m_2 = 0.0064$, $D_2 = 0.0032$) respectively.

value of its future multiplicity is $p_{\rm ave}(x)/p_{\rm end}(x)$. Since the clone positions $x^a(t)$ are distributed as $p_{\rm end}$, averaging this future multiplicity over configurations x yields $\int p_{\rm end} \cdot (p_{\rm ave}/p_{\rm end})\, dx = \int p_{\rm ave}(x)\, dx = 1$, which reflects the fact that the population size is constant in time.

In practice, the distribution of the multiplicity $m_a(t,\tau)$ is very broad, and *typical* multiplicities are far from their *average* values. There are many clones for which no descendants survive until time τ [see Fig. 1(a)], in which case $m_a(t,\tau) = 0$. In order to maintain an average multiplicity of 1, these zero-multiplicity clones are balanced by a small number of clones with larger multiplicity. It is useful to define $\tilde{N}_c(t,\tau)$ as the number of clones that are present in the population at time t, for which $m_a(t,\tau) > 0$. Numerical results for $\tilde{N}_c(t,\tau)$ are shown in Fig. 3; this quantity decreases rapidly as t decreases away from τ, showing that many clones have no surviving descendants, and it follows that the multiplicities of the surviving clones must be large. From (11), one sees that if $\tilde{N}_c(t,\tau)$ is small, numerical estimates of $p_{\rm ave}$ contain only a small number of independent samples, which can lead to large numerical uncertainties within the algorithm.

Moreover, the presence of large multiplicities within the cloning scheme can lead to large *systematic errors*, which cannot be reduced by averaging over repeated runs of the same algorithm. On running the system with a fixed population, the future multiplicity of any clone is bounded above by the population size N_c. We will show in the next section that this constraint has serious implications for systems in the small noise limit. More generally, in order to characterize whether a system requires a large population or not, it is useful to define two numbers that measure how different are the distributions $p_{\rm ave}$ and $p_{\rm end}$. These are

$$m_2 = \int p_{\rm end}(x) \left\{ \left[\frac{p_{\rm ave}(x)}{p_{\rm end}(x)} \right]^2 - 1 \right\} dx \qquad (12)$$

and

$$D_2 = \int p_{\rm ave}(x) \log \left[\frac{p_{\rm ave}(x)}{p_{\rm end}(x)} \right] dx. \qquad (13)$$

Given that $p_{\rm ave}(x)/p_{\rm end}(x)$ is the expected future multiplicity of a clone at x, we recognize m_2 as the variance of this multiplicity, with respect to the distribution $p_{\rm end}$ of clone positions (recall that the average multiplicity with respect to this distribution is equal to unity). Similarly D_2 is the relative entropy of $p_{\rm ave}$ with respect to $p_{\rm end}$ [47]: this is related to the controlling forces that will be introduced in Sec. IV. Large values of m_2 and D_2 indicate that $p_{\rm end}$ and $p_{\rm ave}$ are different from each other, in which case larger values of N_c will be required for accurate results within population dynamics. For the two cases $h = \pm 1$ shown in Fig. 2, we have for $h = -1$ that $(m_2, D_2) = (0.068, 0.039)$, while for $h = +1$, $(m_2, D_2) =$

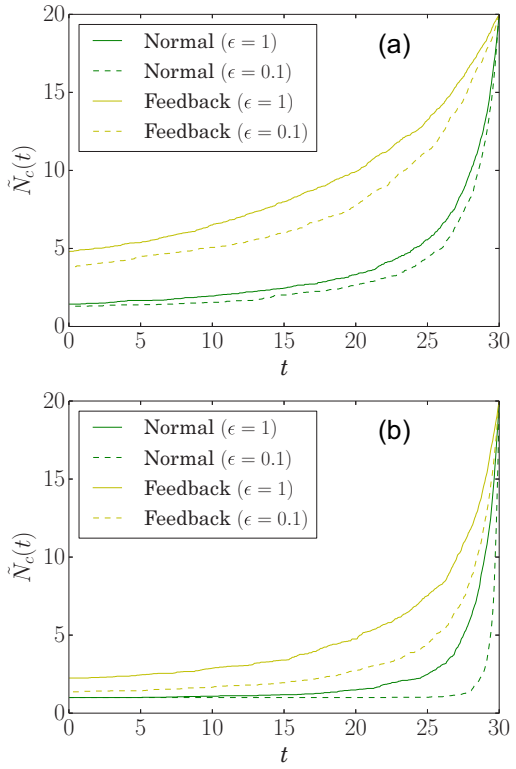

FIG. 3. The number of independent (distinct) clones $\tilde{N}_c(t)$ obtained from the normal population dynamics method [green line (dark gray) line] for $h = -1$ (a) and $h = 1$ (b). The line type corresponds to the value of noise intensity: $\epsilon = 1$ (solid line) and $\epsilon = 0.1$ (dashed line). We set $N_c = 20$ and $\tau = 30$. When the distributions p_{ave} and p_{end} are very different from each other, we expect that $\tilde{N}_c(t)$ decreases rapidly as t decreases from τ: to illustrate this, note that (for $\epsilon = 1$) $m_2 = 0.068$ for $h = -1$ and $m_2 = 0.33$ for $h = 1$: the same ordering is preserved for smaller ϵ. We also plot $\tilde{N}_c(t)$ obtained from the controlled population dynamics [yellow line (light gray) line] with the control-with-feedback explained in Secs. IV C and V. The larger values of $\tilde{N}_c(t)$ obtained with the control-with-feedback method lead to smaller statistical uncertainties in the results.

(0.33,0.17), reflecting the larger populations required for accurate results when $h = +1$. Obtaining general estimates of the actual population size N_c required for convergence is an important goal for future work.

C. Sampling problems for weak noise

The effect described in the previous section is particularly severe for systems where the random (noise) force in (1) is small. To illustrate this case, we set $B(x) = \sqrt{2\epsilon}B_0(x)$, consistent with the numerical example of Sec. II C (for which $B_0 = 1$). The small noise limit is then $\epsilon \to 0$. We define $x^* = \arg\max_x[p_{\text{ave}}(x)]$ as the most likely value of x, within the distribution p_{ave}. The population dynamics requires that the typical multiplicity of a clone with position x^* should be (at least) of order $m^* \equiv p_{\text{ave}}(x^*)/p_{\text{end}}(x^*)$. This clearly cannot be achieved unless $N_c \gtrsim m^*$, which provides an estimate of the number of clones required for accurate results.

This multiplicity m^* increases exponentially as the noise intensity of the system becomes small. In this limit, the dynamics

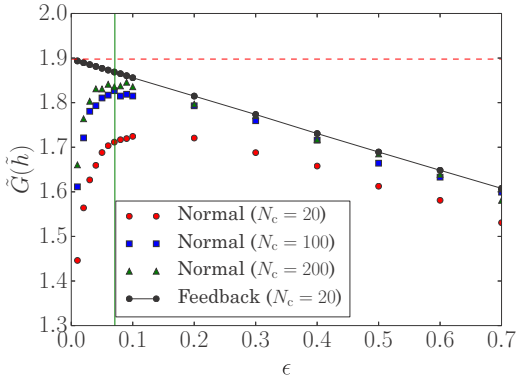

FIG. 4. Estimates of $\tilde{G}(\tilde{h} = 1)$, as ϵ is varied. We compare results from the normal population dynamics and from the control-with-feedback method explained in Secs. IV C and V. The analytical result for $\lim_{\epsilon\to 0}\tilde{G}(\tilde{h})$ is shown as a red dashed line, and the characteristic value of the noise intensity ϵ^*, defined in (14), is plotted as a green vertical solid line. The standard method fails for ϵ smaller than ϵ^*, but the control-with-feedback method (black continuous line and black circles) converges to the correct value even for $\epsilon < \epsilon^*$.

of the system runs increasingly slowly, so it is natural to rescale either the time variable or (equivalently) the biasing field h as $\tilde{G}(\tilde{h}) \equiv \epsilon G(h)$ with $\tilde{h} \equiv h\epsilon$. (This scaling also appears in the hydrodynamic limit of microscopic models [48].) In this limit, p_{ave} and p_{end} satisfy a large deviation principle with respect to the noise intensity ϵ: $p_{\text{ave}}(x) \sim \mathrm{e}^{-I_{\text{ave}}(x)/\epsilon}$ and $p_{\text{end}}(x) \sim \mathrm{e}^{-I_{\text{end}}(x)/\epsilon}$. Hence, $m^* \sim \mathrm{e}^{I_{\text{end}}(x^*)/\epsilon}$, where we used $I_{\text{ave}}(x^*) = 0$. This indicates that we need an exponentially large N_c as ϵ becomes small. More quantitatively, we define a characteristic noise intensity ϵ^* by

$$\epsilon^* \equiv \frac{1}{I_{\text{end}}(x^*)}. \tag{14}$$

For $\epsilon < \epsilon^*$, we expect that population dynamics can not be used practically, because of the exponentially large N_c required.

As a numerical example, we again consider the Brownian particle introduced in Sec. II C. We numerically estimate ϵ^* by using a quadratic approximation of the large deviation function $I_{\text{end}}(x)$. We plot it as a green vertical line in Fig. 4. In the same figure, we show the result of the population dynamics for $\tilde{G}(\tilde{h})$ as ϵ is reduced, with a red constant line corresponding to the analytical value of $\tilde{G}(\tilde{h})$ in the $\epsilon \to 0$ limit (see Appendix D 3 for its determination). Below the characteristic value ϵ^*, the population dynamics method converges very poorly as N_c increases.

IV. POPULATION DYNAMICS WITH A FEEDBACK CONTROL

A. Controlled dynamics

To resolve the sampling issues described in the previous section, we introduce a "control strategy," which modifies the original model (1), in order to make the rare events of interest more likely. (These large deviation problems have dual representations in terms of optimal control problems [49–54], which provide a natural interpretation of the method presented

here.) The modified model is

$$\dot{x}_t = F(x_t) + w(x_t) + B(x_t)\xi_t, \tag{15}$$

where $w(x)$ is a controlling force which we write as

$$w(x) = h\kappa\lambda_c(x) - \kappa\nabla V(x), \tag{16}$$

where V acts as a potential. A straightforward calculation shows that averages with respect to the biased distribution P_h can be rewritten as averages with respect to this modified dynamics, but with a bias on a different observable Λ^w, which replaces Λ. That is, by defining

$$\Lambda^w = \frac{1}{\tau}\int_0^\tau \lambda^w(x_t)\,dt \tag{17}$$

with

$$\lambda^w = \lambda_{\rm d} + \frac{1}{h}\left[(F + w/2)\cdot\kappa^{-1}w - \frac{1}{2}{\rm Tr}(H_V\kappa)\right], \tag{18}$$

in which H_V is a Hessian matrix with elements $(\partial^2 V/\partial x_i\partial x_j)$, we have

$$P_h[X] = P_w[X]{\rm e}^{V(x_\tau)-V(x_0)} \tag{19}$$

with

$$P_w[X] = \pi_0(x_0)\exp\left[-\int_0^\tau \mathcal{L}^w(x_t,\dot{x}_t)\,dt + h\tau\Lambda^w(\tau)\right], \tag{20}$$

where $\mathcal{L}^w$ is the action corresponding to the controlled Langevin equation (15) obtained by replacing $F \mapsto F + w$ in (6). See Appendix B for details of the derivation. We stress that these relations are satisfied for *any* control w.

Averages with respect to P_w are denoted by $\langle\cdot\rangle_w$ and can be calculated using the population dynamics method with the modified model (15). Physically, Eq. (19) says that rare events for the system (1) have an alternative characterization as rare events for the controlled system (15). More precisely, from (19), the averages $\langle\cdot\rangle_h$ and $\langle\cdot\rangle_w$ are not equal, but their associated probability distributions differ only through boundary terms at initial and final times. For large τ, we focus on properties far from initial and final times, in which case the averages $\langle\cdot\rangle_h$ and $\langle\cdot\rangle_w$ are equivalent.

This equivalence implies that

$$p^w_{\rm ave} = p_{\rm ave}, \tag{21}$$

where $p^w_{\rm ave}$ is defined as in (8) but for the controlled population dynamics (15). On the other hand, when we consider properties close to the final time τ (which are not relevant for the large deviations of time-averaged quantities), the two averages $\langle\cdot\rangle_h$ and $\langle\cdot\rangle_w$ are different in general. For example, the end-time distribution $p^w_{\rm end}$ for the controlled dynamics differs from its uncontrolled counterpart as

$$p^w_{\rm end} \propto p_{\rm end}{\rm e}^{-V(x)}, \tag{22}$$

as read from (19) [or see Appendix B 2 for a detailed derivation of (21) and (22)]. Thus the control w allows $p^w_{\rm end}$ to be varied, while always keeping $p^w_{\rm ave}$ constant (and hence leaving unchanged the bulk properties of P_h, which are relevant for the large deviations of time-averaged quantities).

B. Optimal control

These results apply for any control force w, but a (unique) optimal choice w^* can be defined by the condition

$$p^{w^*}_{\rm ave} = p^{w^*}_{\rm end}. \tag{23}$$

From (12) and (13), this result implies that for the controlled population dynamics, $m_2 = D_2 = 0$: all clones have expected future multiplicity of unity, regardless of their position. In fact, this case also implies that $\lambda^{w^*}(x)$ is independent of x (see Appendix B 2), so that there is no cloning or deletion of clones in the optimally controlled population dynamics algorithm. That is, all multiplicities are equal to unity (not just their expected values). The result is that the optimally controlled process [50–54] generates directly the path measure P_h, up to the corrections given in (19) [38,55–58]. Note also that D_2, as defined in (13) for the original population dynamics, is also related to an average of the optimal control potential V^* (where V^* is the potential V corresponding to the optimal control w^*), since $\log[p_{\rm ave}(x)/p_{\rm end}(x)] = -V^*(x) - \log[\int {\rm e}^{-V^*(x')}p_{\rm end}(x')\,dx']$.

The optimal control can be estimated by using the population dynamics with *any* nonoptimal control force w (or its corresponding potential V). We perform the population dynamics and generate sample paths from P_w. From the definition of the optimal force (23) with the relations between $p^w_{\rm end,ave}$ and $p_{\rm end,ave}$ [(21) and (22)], we obtain

$$V^*(x) = V(x) + \log\frac{p^w_{\rm end}(x)}{p^w_{\rm ave}(x)}. \tag{24}$$

Since all terms on the right-hand side of (24) can be measured from the population dynamics with a nonoptimal control w, this allows an estimate of V^*, and hence of w^*.

C. Control-with-feedback for population dynamics

Based on (24), we arrive at the following iteration and feedback scheme for efficient analysis of large deviations of $\Lambda(\tau)$. Starting with the original population dynamics of Ref. [7], we obtain estimates $p^0_{\rm end}$ and $p^0_{\rm ave}$ of $p_{\rm end}$ and $p_{\rm ave}$, and we use (24) to obtain an estimate of the optimal control potential V^*, which we denote by V^1. We then repeat the population dynamics calculation with a control force $w = w^1$ derived from the potential V^1. We use results from this new calculation together with (24) to obtain a new (more accurate) estimate of the optimal control. Iterating this scheme, the estimate of V^* at iteration r is V^r. As $V^r \to V^*$, we have from (24) that $p^w_{\rm end} \to p^w_{\rm ave}$, and hence the sampling problems described in Sec. III B are reduced. This improves the accuracy of the population dynamics method.

Given sufficiently many clones N_c, the original method of [7] can already provide accurate results, but we have demonstrated that for finite N_c there may be large systematic errors. The strength of our scheme is that on repeated iteration, the control potential V approaches the optimal control V^*, and the errors within the method are reduced. Thus, the numerical accuracy of the method increases as the scheme is iterated.

For the implementation of this iteration scheme, we require a computational representation of the function $V(x)$ and its gradient ∇V. From (24), a natural choice might be to represent $p_{\rm ave}$ and $p_{\rm end}$ by histograms, based on a discretization of the

configuration space. However, this choice does not facilitate estimation of ∇V, and it is also unfeasible in high-dimensional systems. We therefore use a potential V that is defined in terms of a set of basis functions ζ_i, with coefficients c_i:

$$V(x) = \sum_{i=1}^{k} c_i \zeta_i(x), \tag{25}$$

where k is the size of the basis set.

At stage r of our iterative scheme, the coefficients c are denoted by $\boldsymbol{c}^r = (c_i^r)_{i=1}^k$. In the absence of prior information about the optimal control V^*, the first stage of the method ($r = 0$) uses the original population dynamics, so $c_i^0 = 0$ for all i. In stage $r+1$, we update these coefficients according to (24) so that the potential V^{r+1} in the next stage is the best available estimate of V^*. There is considerable freedom in how to obtain this estimate: we take

$$\boldsymbol{c}^{r+1} = \text{argmin}_{\boldsymbol{c}} \int_{\Omega_r} \left[V^r(x) + \log \frac{p_{\text{end}}^{w,r}(x)}{p_{\text{ave}}^{w,r}(x)} - \sum_{i=0}^{k} c_i \zeta_i(x) \right]^2 dx, \tag{26}$$

where $p_{\text{end}}^{w,r}$ is the numerical estimate for p_{end}^{w} obtained at iteration r, and similarly $p_{\text{ave}}^{w,r}$. The state space Ω_r is defined as the space where $p_{\text{ave}}^{w,r} > 0$ [note that $p_{\text{end}}^{w,r}(x) > 0$ whenever $p_{\text{ave}}^{w,r}(x) > 0$, from the definition of how to construct $\tilde{x}^a(t)$ as shown in Fig. 1(b)].

We emphasize that, for *any* basis set ζ_i (with *any* truncation number k), Eq. (19) is satisfied, meaning that if the number of clones N_c and the time τ are large enough, the result of *any* controlled population dynamics always leads to the same results, which can also be obtained from the original (uncontrolled) population dynamics. However, the choice of the expansion functions $\zeta_i(x)$ (and the value of the truncation number k) does affect the computational cost, through the number of clones required for convergence, as discussed in Sec. III B.

D. Advantages of the control-with-feedback for population dynamics, and relation to other methods

Compared to the original population dynamics method, the addition of control forces and the use of iteration and feedback increase the complexity of the method presented here. Here, we summarize the improvements that these changes achieve. Typically, existing methods either exploit an ensemble (population) of copies of the system [19–24], or they use modified (controlled) dynamical rules to drive the system towards rare events of interest [12–18,53], or they use path-sampling methods [27,59]. All these methods are useful, but the population-based methods can suffer convergence problems, due to the very large populations required in some problems. On the other hand, the controlled methods require accurate estimation of an optimal control force that is typically a high-dimensional and complex object, which can be difficult to represent computationally (see, for example, Ref. [60]). Path sampling methods are most effective when the ensemble P_h has time-reversal symmetry, which limits their applicability in nonequilibrium settings. The method proposed here is a mixture of the population- and control-based methods, as illustrated schematically in Fig. 5.

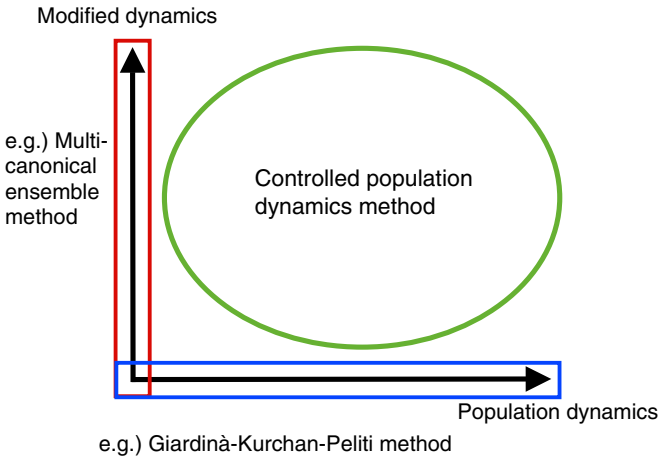

FIG. 5. A schematic map illustrating the methodological situation of the controlled population dynamics.

In terms of the applicability of this new method, we expect the following general behavior. For complex high-dimensional problems, accurate representation of the optimal control V^* is likely to be difficult, but we expect even approximate representations of V^* to significantly improve the performance of the population dynamics method. Thus, the controlled method should reduce the computational cost of problems that are already tractable using population dynamics, allowing access (for example) to larger system sizes and larger values of the bias parameter h. On the other hand, for relatively simple problems such as the particle in a quartic potential of Sec. II C, the original population dynamics fails for small noise (Fig. 2), but we would expect that a solution by the controlled method of Ref. [12] might already be possible. However, for a similar model in three or more dimensions, we expect that the method of Ref. [12] would already be challenging, due to the difficulty of representing exactly the effective potential. Here we combine that control strategy with population dynamics: we arrive at a flexible method that exploits the strengths of both approaches and that we anticipate will be effective in a wide variety of problems.

V. NUMERICAL EXAMPLE

To illustrate the control-with-feedback method, we consider the numerical example from Sec. II C, and we take the effective description in (25) to be a quartic polynomial: $\zeta_i(x) \equiv x^i$ (that is, "x raised to the power i") and $k = 4$. For the first iteration of the method we take $(c_i^0)_{i=1}^k = 0$. Note that this potential-parametrization of V cannot capture the exact V^*, neither for $\epsilon > 0$ nor in the limit $\epsilon \to 0$ (see Appendix D). This emphasizes that the control-with-feedback method does not require a perfect representation of the optimal control in order to improve the convergence of the population dynamics method.

Figure 2 shows estimates of the distribution p_{ave} obtained using the original cloning method [Figs. 2(a)–2(d)], compared with the results obtained using control-with-feedback

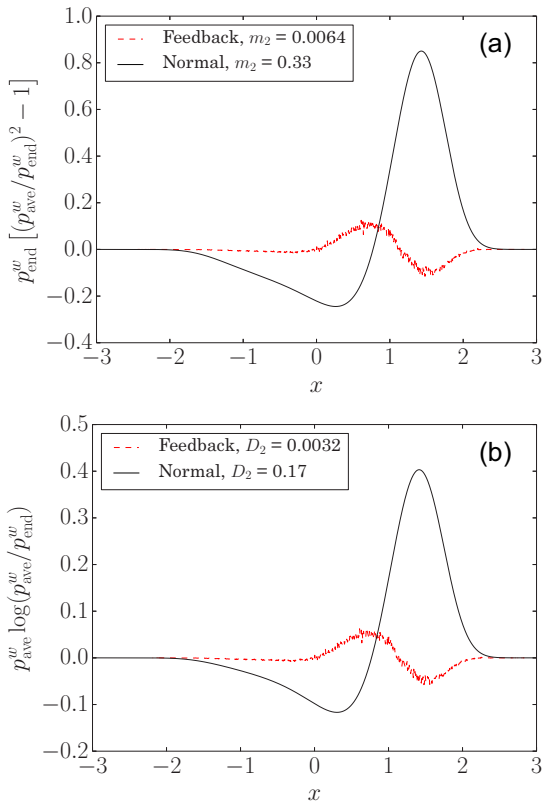

FIG. 6. The integrands of m_2 and D_2 defined as (a) $p^w_{\rm end}(x)\{[\frac{p^w_{\rm ave}(x)}{p^w_{\rm end}(x)}]^2 - 1\}$ and (b) $p^w_{\rm ave}(x)\log[\frac{p^w_{\rm ave}(x)}{p^w_{\rm end}(x)}]$ [see (12) and (13)] for the standard (normal) population dynamics method ($w = 0$) and for the control-with-feedback method (w: obtained from the control-with-feedback method). For the control-with-feedback method, we set $N_c = 20$. In the legends, we put the values corresponding to m_2 and D_2.

procedure proposed here [Fig. 2(e)]. (Two iterations of the feedback were used, which allow an accurate estimate of the optimal control potential V^*.) The comparison between Figs. 2(d) and 2(e) shows that the number of clones required to obtain convergence to the exact result is much reduced using the control-with-feedback method.

In the weak-noise limit $\epsilon \to 0$, one can see this advantage more clearly. In this limit, a sampling issue arises because of the exponential increase of the required number of copies N_c, as discussed in Sec. III C. Figure 3(a) shows numerical results for $\tilde{G}(\tilde{h})$, as ϵ is reduced. The normal population dynamics converges very poorly for small noise, $\epsilon < \epsilon^*$. However, the controlled population dynamics does not fail at small ϵ because it maintains $p^w_{\rm end} \approx p^w_{\rm ave}$ [61].

We then consider *statistical* errors. Figure 3(b) shows the number of distinct clone positions in the population, $\tilde{N}_c(t)$. Again, the control-with-feedback method performs better than the original method, in that it averages over a larger sample of distinct positions, reducing the statistical errors.

Finally, in order to illustrate how the control-with-feedback method improves the standard population dynamics method, in Fig. 6, we show the integrands of m_2 and D_2 defined in (12) and (13) [62]. As discussed in Secs. III B and III C, the standard population dynamics has sampling issues, which are captured by the deviations of m_2 and D_2 from 0. In the figure, we can see that the control-with-feedback method greatly reduces the values of m_2 and D_2 close to 0, ensuring that $p^w_{\rm end}$ and $p_{\rm ave}$ are closer than in the original cloning, thus yielding better performances as seen throughout this section.

VI. OUTLOOK

We have shown that the performance of the population dynamics algorithm for sampling large deviations [7] can be improved by introducing a controlling force w. Given the optimal choice for this force, the rare events of interest in large deviation theory can be characterized as typical trajectories of the controlled system without any cloning. In complex systems with many degrees of freedom it is likely that the optimal w cannot be determined exactly, but even nonoptimal controls can still significantly improve both the statistical and the systematic errors associated with the population dynamics method (see Sec. V). It is straightforward to improve existing population dynamics codes to include this approach: we expect that it will significantly expand the range of systems for which numerical calculations can be performed, including open quantum systems [63,64], or more complex molecular dynamics models than those considered so far [27,59].

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of Fondation Sciences Mathématiques de Paris, EOTP NEMOT15RPO, PEPS LABS, and LAABS Inphyniti CNRS project. The research leading to these results has received funding from the European Research Council under the European Union's seventh Framework Programme (FP7/2007-2013 Grant No. 616811) (F.B. and T.N.). V.L. acknowledges support by the National Science Foundation under Grant No. NSF PHY11-25915 during a stay at KITP, UCSB.

APPENDIX A: POPULATION DYNAMICS METHOD

In this Appendix, complementing Secs. II B and III A, we explain the details of the population dynamics algorithm.

1. Population dynamics algorithm

The population dynamics is a numerical technique designed to evaluate a large deviation function associated to the CGF of a time-averaged observable $\Lambda(t)$. Each step of the algorithm consists of a first substep in which the normal (unbiased) dynamics of the system is simulated for a time ΔT, followed by an elimination-multiplication substep. (The elimination-multiplication substep is also called a cloning step, or a mutation-selection step.) In detail, the method is

(1) Generate N_c initial conditions, for example, drawn from the stationary state of the unbiased ($h = 0$) dynamics.

(2) Repeat the following procedure M times. (The iteration index is $m = 0, 1, \ldots, M-1$.)

(a) For each copy of the system, perform the normal dynamics from $t = m\,\Delta T$ to $(m+1)\Delta T$. We denote each trajectory by $x^a(t)$. (Throughout this section, $a = 1, 2, \ldots, N_c$.) During the simulation, for each trajectory,

calculate

$$s_a = \exp\{h[(t+\Delta T)\,\Lambda(t+\Delta T) - t\Lambda(t)]\}. \quad \text{(A1)}$$

(b) For each trajectory a, calculate an integer n_a as

$$n_a = \left\lfloor \frac{s_a}{\sum_b s_b} N_c + \eta \right\rfloor, \quad \text{(A2)}$$

where η is a random number uniformly distributed on [0,1] and $\lfloor\cdot\rfloor$ denotes the lower integer part. Calculate and store the quantity $S_m = \sum_b s_b$.

(c) Multiply or eliminate each trajectory a so that it appears n_a times in the new population. (For example, if $n_a = 0$ then trajectory a is deleted. If $n_a = 5$, then we retain trajectory a and we introduce four new copies of that trajectory.)

(d) Eliminate or multiply trajectories within the population, chosen randomly and uniformly, so that the total number of surviving trajectories is N_c.

(e) Go back to (a), using the current set of configurations $x^a[(m+1)\Delta T]$ as initial conditions for the next iteration of the normal dynamics.

Note that if the population were not kept constant in step 2c above, then the population would expand by a factor of S_m/N_c. It follows that the CGF can be estimated as

$$G(h) \simeq \frac{1}{M\,\Delta T} \sum_m \log \frac{S_m}{N_c}. \quad \text{(A3)}$$

Also, averages over the population at the final time τ are estimates of averages with respect to $p_{\rm end}$:

$$\int f(x) p_{\rm end}(x)\,dx \simeq \frac{1}{N_{\rm c}} \sum_{a=1}^{N_{\rm c}} f[x^a(\tau)], \quad \text{(A4)}$$

which follows from the definition of $p_{\rm end}$. When estimating $p_{\rm end}$, we can improve the statistics by using the history of $x^a(t)$. That is, assuming an ergodicity property, we can replace $f[x^a(\tau)]$ by its time average, leading to

$$\int f(x) p_{\rm end}(x)\,dx \simeq \frac{1}{\tau N_{\rm c}} \int_0^\tau \sum_{a=1}^{N_{\rm c}} f[x^a(t)]\,dt. \quad \text{(A5)}$$

This means that the empirical distribution of $x^a(t)$ is an estimator for $p_{\rm end}$, as shown in (10).

In order to generate the sample paths corresponding to the biased measure P_h, we also need to copy the history of trajectory (not just the current configuration of x) in the selection-mutation procedure in step 2(b) of the algorithm. This fact is directly derived from the definition of P_h. Thus, the $x^a(t)$ defined above do not correspond to sample paths of P_h. The paths are obtained as $\tilde{x}^a(t)$, which are defined as those trajectories that survive until the final time τ (see Fig. 1). In numerical simulations, there are several ways to generate (or reconstruct) these trajectories, as we now explain.

2. Generating continuous sample paths $\tilde{x}^a(t)$ for the biased dynamics

A simple way to characterize $\tilde{x}^a(t)$ is the following: If we do not require full sample paths but wish only to evaluate the biased average of an additive observable $A(\tau) = \int_0^\tau a[x(t)]\,dt$, a simple method [65] consists in attaching a value of the observable A to every trajectory and, at every time step, to update its value and copy or delete it together with the trajectory. Then an evaluation of the biased average of A is given by an average of the numerical values of A: this average runs over all trajectories that are present *at the final time*. For example, when we divide the configuration space into small bins and take $a_i[x(t)] = 1$ if $x(t)$ is in bin i, $A_i(\tau)/\tau$ is an estimate of $p_{\rm ave}$, integrated across the ith bin.

For the small systems where we can store all of the trajectories in the population dynamics, we can generate full sample paths corresponding to $\tilde{x}^a$. The procedure is as follows: we first generate all the trajectories and then select those that survive until the final time τ. Considering the N_c copies at final time, indexed by $1 \leqslant a \leqslant N_c$, one can follow the ancestors of every copy. Upon every coalescence observed backwards in time (corresponding to multiplications of clones in the original forwards simulation), one increments a counter $m_a(t,\tau)$ by the number of trajectories which have coalesced. At the end of the procedure, the counters $[m_a(t,\tau)]_{1\leqslant a\leqslant N_c}$ represent, at time t, the number of descendants of a copy a at final time τ.

APPENDIX B: DERIVATION OF THE RATIO OF PATH PROBABILITY DENSITY (19)

In this Appendix, complementing Sec. IV we derive the relation between $P_h[X]$ and $P_w[X]$ [Eq. (19)]. We show the derivation in two ways, one based on path probability densities (stochastic differential equations) and the other on Fokker-Plank equations.

1. Derivation using path probability density

We denote a trajectory of the system by $X = [x(t)]_{0\leqslant t\leqslant\tau}$. From the definitions of $P_h[X]$ and $P_w[X]$, we have

$$\frac{P_w[X]\mathrm{e}^{-h\tau\Lambda^w(\tau)}}{P_h[X]\mathrm{e}^{-h\tau\Lambda(\tau)}} = \exp\left[\int_0^\tau (\dot{x}-F)\cdot\kappa^{-1}w\,dt - \frac{1}{2}\int_0^\tau w\cdot\kappa^{-1}w\,dt\right]. \quad \text{(B1)}$$

The integrand on the right-hand side is written as

$$(\dot{x}-F)\cdot\kappa^{-1}w - \frac{1}{2}w\cdot\kappa^{-1}w = \dot{x}\cdot(-\nabla V + h\lambda_{\rm c}) - \left(F+\frac{1}{2}w\right)\cdot\kappa^{-1}w, \quad \text{(B2)}$$

where we have used the expression of $w(x)$ as given in the main text ($w(x) = \kappa[-\nabla V(x) + h\lambda_{\rm c}(x)]$). We then consider the integral of the first term on the right-hand side:

$$\int_0^\tau \dot{x}\cdot(-\nabla V)\,dt. \quad \text{(B3)}$$

Since the trajectory X is generated from the stochastic differential equation (15) and we use the Itô convention, the time-derivative of $V(x(t))$ is given by Itô's formula

$$\frac{d}{dt}V = \dot{x}\cdot\nabla V + \frac{1}{2}\mathrm{Tr}[B^T H_V B]. \quad \text{(B4)}$$

Here H_V is a Hessian matrix defined as $(H_V)_{i,j} = \frac{\partial V}{\partial x_i \partial x_j}$. Combining (B4) and (B3) we have

$$\int_0^\tau \dot{x}\cdot(-\nabla V)\,dt = -V[x(\tau)] + V[x(0)] + \int_0^\tau \frac{1}{2}\mathrm{Tr}[B^T H_V B]\,dt. \quad \text{(B5)}$$

Thus, from (B1), (B2), and (B5), we get

$$\frac{P_w[X]\mathrm{e}^{-h\tau\Lambda^w(\tau)}}{P_h[X]\mathrm{e}^{-h\tau\Lambda(\tau)}} = \mathrm{e}^{-V[x(\tau)]+V[x(0)]}\exp\left\{\int_0^\tau\left[\frac{1}{2}\mathrm{Tr}[B^T H_V B] + h\dot{x}\cdot\lambda_c - \left(F+\frac{w}{2}\right)\cdot\kappa^{-1}w\right]dt\right\}. \quad \text{(B6)}$$

Finally, by noticing $\mathrm{Tr}[B^T H_V B] = \mathrm{Tr}[H_V\kappa]$ and using the definitions of Λ^w and Λ, the right-hand side is $\mathrm{e}^{-V[x(\tau)]+V[x(0)]}\mathrm{e}^{h\tau\Lambda(\tau)-h\tau\Lambda^w(\tau)}$. Hence one arrives at Eq. (19).

2. Derivation using time-evolution operator

An alternative derivation of (19) is obtained by using a "tilted" generator (or master operator) for the biased ensemble of trajectories. Let $u^h(x,\tau)$ be the (unnormalized) probability density at time τ, obtained as a marginal of the path distribution P_h. As discussed, for example, in Appendix A.2 of Ref. [38], this distribution evolves in time according to a generalized Feynman-Kac formula as

$$\frac{\partial}{\partial\tau}u^h = L^h[u^h] \quad \text{(B7)}$$

with

$$L^h[f] \equiv L^F_{\mathrm{FP}}[f] + h(\lambda_d + \lambda_c\cdot F)f + \frac{h^2}{2}(\lambda_c\cdot\kappa\lambda_c)f - h\nabla\cdot(\kappa\lambda_c f). \quad \text{(B8)}$$

Here the Fokker-Planck operator L^F_{FP} is

$$L^F_{\mathrm{FP}}[f] = -\nabla\cdot[Ff] + \frac{1}{2}\sum_{i,j}\frac{\partial^2}{\partial x_i\partial x_j}\kappa_{ij}f, \quad \text{(B9)}$$

where the superscript F on L^F_{FP} indicates that the particle feels the physical force F introduced in (1).

For the controlled population dynamics, the analog of u^h is $u^w(x,\tau)$, which evolves as $\frac{\partial}{\partial\tau}u^w = L^w[u^w]$, with

$$L^w[f] \equiv L^{F+w}_{\mathrm{FP}}[f] + h\lambda^w f. \quad \text{(B10)}$$

The relation (19) follows from a duality relation between L^h and L^w:

$$L^h[f] = \mathrm{e}^V L^w[f\mathrm{e}^{-V}]. \quad \text{(B11)}$$

This relation may be verified directly from (B8) and (B10), noting that the potential V is related to the control w via the definition $w = h\kappa\lambda_c - \kappa\nabla V$.

From (B7), we note that the operator $U^h_\tau = \mathrm{e}^{\tau L^h}$ corresponds to integration forward in time over a duration τ. Similarly $U^w_\tau = \mathrm{e}^{\tau L^w}$, and from (B11) we have $U^h_\tau[f] = \mathrm{e}^V U^w_\tau[f\mathrm{e}^{-V}]$. Setting $f(x) = \delta(x-x_0)$, then $u^h(x,\tau|x_0,0) = U^h_\tau[f]$ is the (unnormalized) probability density at x, for a particle that was at x_0 a time τ earlier. Defining similarly $u^w(x,\tau|x_0,0)$, (B11) implies

$$u^h(x,\tau|x_0,0) = \mathrm{e}^{V(x)}u^w(x,\tau|x_0,0)\mathrm{e}^{-V(x_0)}. \quad \text{(B12)}$$

Hence one arrives at (19).

This approach also provides insight into the distributions p_{ave} and p_{end}, as discussed in Refs. [7,56]. One easily sees that

$$p_{\mathrm{end}}(x) = \lim_{\tau\to\infty}\frac{u^h(x,\tau|x_0,0)}{\int_{x'}u^h(x',\tau|x_0,0)}, \quad \text{(B13)}$$

which is independent of x_0. Similarly,

$$p_{\mathrm{ave}}(x) = \lim_{\tau\to\infty}\frac{\int_{x_1}u^h(x_1,\tau/2|x)u^h(x,0|x_0,-\tau/2)}{\int_{x',x_1}u^h(x_1,\tau/2|x')u^h(x',0|x_0,-\tau/2)}. \quad \text{(B14)}$$

For large τ, the propagator u^h is dominated by the largest eigenvalue of L^h, as

$$u^h(x,\tau|x_0,0) \simeq p_{\mathrm{end}}(x)\mathrm{e}^{G(h)\tau}q(x_0), \quad \text{(B15)}$$

where $p_{\mathrm{end}}(x)$ is the dominant right eigenvector of L^h [required for consistency with (B13)], the associated eigenvalue is $G(h)$, and $q(x)$ is the dominant left eigenvector. The approximate equality in (B15) is valid for large times, up to corrections of order $\mathrm{e}^{-\lambda\tau}$, where λ is the spectral gap of L^h. Combining (B13)–(B15) we have $p_{\mathrm{ave}}(x) \propto p_{\mathrm{end}}(x)q(x)$.

This approach also shows why p_{ave} is not affected by the control force w: the dominant left and right eigenvectors of L^h are q and p_{end} so (B11) means that the dominant eigenvectors of L^w are $q^w = q\mathrm{e}^V$ and $p^w_{\mathrm{end}} = \mathrm{e}^{-V}p_{\mathrm{end}}$. Hence it is clear that $p^w_{\mathrm{ave}} = q^w p^w_{\mathrm{end}} = qp_{\mathrm{end}} = p_{\mathrm{ave}}$.

In the special case where w is given by the optimal control w^* (that is defined as the control w satisfying the condition $p^w_{\mathrm{ave}} = p^w_{\mathrm{end}}$ in the main text), one can show that the controlled system is described by the auxiliary process [56] (or the "driven process" [38]), which is a Markov process whose path probability density is equivalent to P_h in its stationary regime. (Indeed, $p^{w^*}_{\mathrm{ave}} = p^{w^*}_{\mathrm{end}}$ implies $q^{w^*} = 1$, which expresses that L^{w^*} conserves probability.) In this case, one has [38]

$$\mathrm{e}^{-V}L^h[f\mathrm{e}^V] = L^{F+w^*}_{\mathrm{FP}}[f] + G(h)f, \quad \text{(B16)}$$

where $G(h)$ is a constant (independent of x): this is the CGF. Comparing with (B11) one sees that $\lambda^{w^*}(x)$ is independent of x, from which it follows that the population dynamics in this case has no cloning or deletion of clones (this property is true for all finite N_c: all clones have equal weights at all times).

APPENDIX C: AN EXAMPLE OF THE FEEDBACK-ALGORITHM

Here, in order to complement Sec. IV C, we explain the algorithm used within the feedback population dynamics. The procedure is a combination of the population dynamics and an iterative construction of a control potential $V(x)$ that is close to the optimal control V^*. There is considerable flexibility in the precise definitions of the estimators used in this algorithm, but these choices have proven effective in the simple model problem considered here.

(1) Generate N_c initial conditions, for example, drawn from the stationary distribution of the original (unbiased) system.

(2) Repeat the following feedback procedure R times (the iteration index is $r = 0,1,\ldots,R-1$). We denote by $V^r(x)$ the control potential $V(x)$ for iteration r, and we take $V^0(x) = 0$.

(a) Perform the population dynamics for the system as explained in Appendix A, using a time interval $M\tau_0$. The unbiased evolution within the method includes the control force w^r that is obtained from the control potential V^r, and the elimination-multiplication step uses the corresponding biasing factor Λ^{w^r}. The time τ_0 between elimination-multiplication steps should be larger than the correlation time of the system. From each time segment (indexed by m), estimate the distributions

$$p_1^{m,r}(x) = \frac{1}{N_c\tau_0}\sum_{a=1}^{N_c}\int_{m\tau_0}^{(m+1)\tau_0}\delta[x - x^a(t)]\,dt \qquad \text{(C1)}$$

and

$$p_0^{m,r}(x) = \frac{1}{N_c(\tau_0 - t_{\rm end})} \times \sum_{a=1}^{N_c}\int_{m\tau_0}^{(m+1)\tau_0 - t_{\rm end}}\delta[x - \tilde{x}^a(t)]\,dt, \qquad \text{(C2)}$$

where the trajectories $\tilde{x}$ are defined on the time interval $[m\tau_0,(m+1)\tau_0]$, as specified in Appendix A 2. The shift parameter $t_{\rm end}$ is chosen so that p_0 is an accurate estimator for $p_{\rm ave}$, by excluding times t that are too close to the final time $(m+1)\tau_0$. If τ_0 is large enough, all results should depend weakly on $t_{\rm end}$.

(b) Having completed M time segments within the population dynamics, evaluate $p_{\rm end}^{w,r}(x)$ and $p_{\rm ave}^{w,r}(x)$ as

$$p_{\rm end}^{w,r}(x) = \frac{1}{M}\sum_m p_1^{m,r}(x), \qquad \text{(C3)}$$

$$p_{\rm ave}^{w,r}(x) = \frac{1}{M}\sum_m p_0^{m,r}(x). \qquad \text{(C4)}$$

(c) Finally, from these distribution functions, calculate $V^{r+1}(x)$ in terms of a sum of basis functions, according to Eq. (26). In practice, note that it is not necessary to keep track of the full distributions p_0 and p_1, but only those statistics that are required to solve the minimisation in (26). Also, it is sometimes convenient to take $V^{r+1}(x) = V^r(x)(1-\alpha) + V_{\rm new}(x)\alpha$, where $V_{\rm new}(x)$ is the control potential specified by the right-hand side of (26), and α is a parameter (with $0 < \alpha \leqslant 1$) that acts to suppress large fluctuations in V.

(3) Go back to step 2 and perform the next iteration $(r+1)$, with the control potential V^{r+1}, and initial conditions for the clones given by their current states $x^a(M\tau_0)$.

APPENDIX D: LANGEVIN SYSTEM WITH QUARTIC POTENTIAL

In this final Appendix, in order to complement Sec. V, we explain the property of the system we considered there: the parameters are given by $d = 1$, $F(x) = -x^3$, $B(x) = \sqrt{2\epsilon}$, $\lambda_{\rm d}(x) = \lambda(x) \equiv x + x^2$ and $\lambda_{\rm c}(x) = 0$. We focus on the small-noise limit $\epsilon \to 0$. Throughout this section, h corresponds to $\tilde{h}$ in the main text (see below).

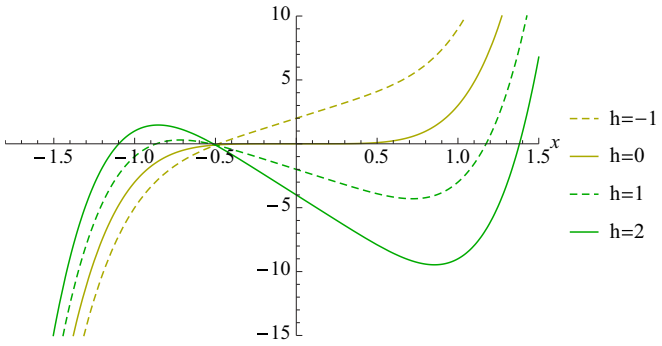

FIG. 7. Plots of the polynomial $3x^5 - 4hx - 2h$ for several h. The roots of this polynomial determine the concentration points of $p_{\rm ave}(x)$ for $\epsilon \to 0$ in the model system considered.

The main features of the limit $\epsilon \to 0$ are the following:

(1) The distribution $p_{\rm ave}(x)$ concentrates on a point $x_{\rm ave}$ that is a root of the polynomial

$$3x^5 - 4hx - 2h = 0.$$

This function is sketched in Fig. 7. For $h > 0$, the concentration is at the positive root ($x_{\rm ave} > 0$); for $h = 0$ one has $x_{\rm ave} = 0$. For negative h, the point $x_{\rm ave}$ decreases quickly from zero and localizes at $x_{\rm ave} \approx \frac{1}{2}$.

(2) There is a second-order dynamical phase transition at $h = 0$, which appears as divergence of the second derivative of the dynamical free energy, $G''(h)$ (see Fig. 8).

(3) The distribution $p_{\rm end}(x)$ concentrates on a point $x_{\rm end}$, with $x_{\rm end} \neq x_{\rm ave}$ in general. This leads to poor convergence of the population dynamics method for small ϵ, as discussed in the main text.

(4) Even though the system is simple, the analytical expressions of $p_{\rm ave}$ and $p_{\rm end}$ are not straightforward. In particular, the perfect potential $V^*(x)$ corresponding to $w^*(x)$ is not expressed exactly as the quartic polynomial expansion used to perform a numerical evaluation of $w(x)$; however, as described in the main text, this does not affect the effectiveness of the numerical procedure.

Below, relying on the Euler-Lagrange equation, we derive the analytical results of $G(h)$, $p_{\rm ave}$ and $p_{\rm end}$ in $\epsilon \to 0$, from which these features are obtained.

1. Euler-Lagrange equation (Instanton equation)

We consider the following finite time CGF:

$$G_{\tau,\epsilon}(h) = \frac{\epsilon}{\tau}\log\langle e^{(\tau/\epsilon)h\Lambda(\tau)}\rangle_{\rm st}, \qquad \text{(D1)}$$

where $\langle\,\rangle_{\rm st}$ means the average with respect to the path with a stationary initial condition. [Hereafter, we denote this initial distribution function by $P_{\rm st}(x)$.] The function $G_\epsilon(h) \equiv \lim_{\tau\to\infty} G_{\tau,\epsilon}(h)$ corresponds to $\tilde{G}(\tilde{h})$ in the main text. By

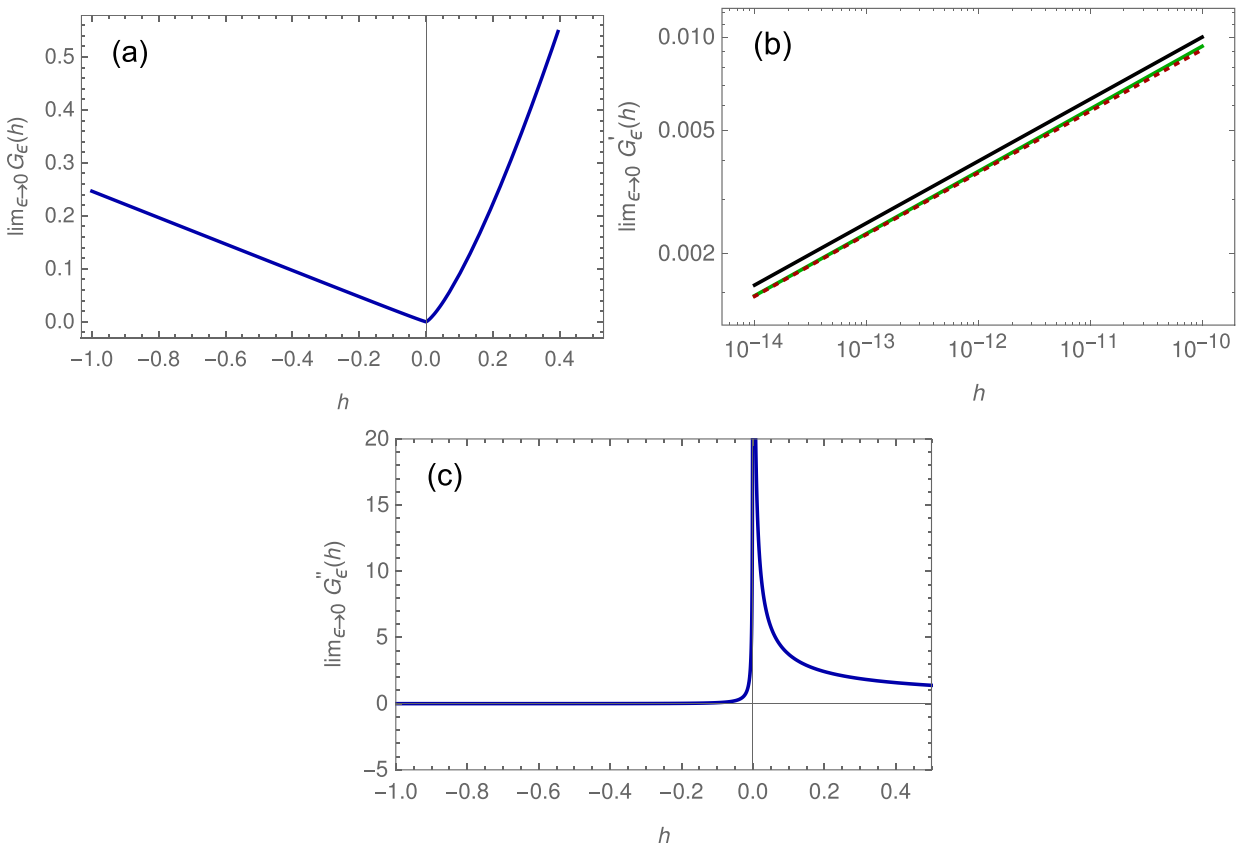

FIG. 8. (a) Generating function $\lim_{\epsilon\to 0} G_\epsilon(h)$. (b) The first derivatives, where the green (dark gray) solid line represents $(d/dh)\lim_{\epsilon\to 0} G_\epsilon(h)$, the red dotted line represents $-(d/dh)\lim_{\epsilon\to 0} G_\epsilon(-h)$, and the black solid line represents a straight line $h^{1/5}$. We find that the first derivative converges to 0 as a power law $\sim h^{1/5}$ (as can also be checked analytically). (c) The second derivative of $\lim_{\epsilon\to 0} G_\epsilon(h)$ with respect to h. These are calculated from (D11)–(D12). We can see that the second derivative shows a singularity at $h = 0$, although the first derivative converges to 0. This represents a second-order dynamical phase transition.

taking $\epsilon \to 0$, we obtain the following variational principle:

$$\lim_{\epsilon\to 0} G_{\tau,\epsilon}(h) = -\frac{1}{\tau}\min_{x_0,x_\tau}\left[\min_{\substack{(x(t))_{t=0}^{\tau} \\ x(0)=x_0, x(\tau)=x_\tau}} \int_0^\tau L(\dot{x}(t),x(t))\,dt + F_{\rm free}(x_0)\right] \tag{D2}$$

with the Lagrangian $L(\dot{x},x)$ defined as

$$L(\dot{x},x) \equiv \tfrac{1}{4}[\dot{x} - F(x)]^2 - h\lambda(x), \tag{D3}$$

and the free energy function $F_{\rm free}(x_0)$ defined as

$$F_{\rm free}(x_0) \equiv -\lim_{\epsilon\to 0} \epsilon \log P_{\rm st}(x_0) = \tfrac{1}{4}x_0^4 + \text{const.} \tag{D4}$$

Then, the corresponding Euler-Lagrange equation (Hamilton equation), which is obtained from minimizing this action, is

$$\dot{x} = -x^3 + 2p, \tag{D5}$$

$$\dot{p} = 3px^2 - h(2x+1) \tag{D6}$$

with the required initial and the final conditions as

$$p(0) = \left.\frac{\partial F_{\rm free}(x)}{\partial x}\right|_{t=0} = x(0)^3, \tag{D7}$$

$$p(\tau) = 0. \tag{D8}$$

We analyze these equations numerically and analytically in Ref. [66]. The following results are based on that study.

2. Steady solutions

Here we consider the steady solutions of these instantons, which is defined as the solution obtained from $\dot{x}_{\rm st} = \dot{p}_{\rm st} = 0$ in (D5) and (D6). These conditions lead to

$$p_{\rm st} = \tfrac{1}{2}x_{\rm st}^3 \tag{D9}$$

and

$$3x_{\rm st}^5 - 4hx_{\rm st} - 2h = 0. \tag{D10}$$

We plot the left-hand side of (D10) as a function of x in Fig. 7 for several fixed h. The figure shows that this equation has three solutions, when h is larger than a certain value (larger than 0).

3. Cumulant generating function

From the variational principle (D2), even in the case where there are multiple instanton solutions, the CGF can be calculated. This is based on the observation that the instanton solution corresponding to the minimum is time independent

[67]. More precisely, by combining this observation with the variational principle (D2), we get

$$\lim_{\epsilon\to 0} G_\epsilon(h) = \max_{x_{\rm st}} G_{\rm st}(x_{\rm st}) \tag{D11}$$

with

$$G_{\rm st}(x_{\rm st}) \equiv -\tfrac{1}{4}x_{\rm st}^6 + h\left(x_{\rm st}^2 + x_{\rm st}\right). \tag{D12}$$

We plot the $\epsilon \to 0$ result, $\lim_{\epsilon\to 0} G_\epsilon(h)$, in Fig. 8, from which we can see that the generating function has a kink at the origin, which is the sign of the dynamical phase transition in this system, appearing in the zero-temperature limit. Asymptotic analysis allows to find $G(h) \sim A_\pm |h|^{1/5}$ with $A_\pm$ depending on the sign of h, as illustrated on Fig. 8.

4. Analytical expressions of $p_{\rm end}(x)$ and $p_{\rm ave}(x)$ in $\epsilon \to 0$

Finally, we write the explicit analytical expressions of $p_{\rm end}(x)$ and $p_{\rm ave}(x)$ in the $\epsilon \to 0$ limit. We consider the biased (unnormalized) probability density u^h introduced at the beginning of Sec. B 2. We also consider the same function but with fixed initial condition $u^h(x,\tau|x_0,\tau)$. By using these function, we introduced two logarithmic functions defined as

$$W_{\rm F}(x,t) \equiv \epsilon \log u^h(x,t), \tag{D13}$$

$$W_{\rm B}(x,t) \equiv \epsilon \log \int u^h(y,t|x,0)\,dy. \tag{D14}$$

From the generalized Feynman-Kac formula (B7), we obtain the time evolution equation for them as

$$\frac{\partial}{\partial t} W_{\rm F}(x,t) = -\epsilon \frac{\partial}{\partial x} F(x) - F(x)\frac{\partial}{\partial x} W_{\rm F}(x,t) + \epsilon\left(\frac{\partial}{\partial x}\right)^2 W_{\rm F}(x,t) + \left[\frac{\partial}{\partial x} W_{\rm F}(x,t)\right]^2 + h\lambda(x) \tag{D15}$$

and

$$\frac{\partial}{\partial t} W_{\rm B}(x,t) = F(x)\frac{\partial}{\partial x} W_{\rm B}(x,t) + \epsilon\left(\frac{\partial}{\partial x}\right)^2 W_{\rm B}(x,t) + \left[\frac{\partial}{\partial x} W_{\rm B}(x,t)\right]^2 + h\lambda(x). \tag{D16}$$

These equations can be solved in $\epsilon = 0$ with t large limit. Indeed, by setting $W_{\rm F}(x,t) = tG(h) + W_{\rm F}(x)$ and $W_{\rm B}(x,\tau - t) = (\tau - t)G(h) + W_{\rm B}(x)$ with $G(h) \equiv \lim_{\epsilon\to 0} G_\epsilon(h)$ in

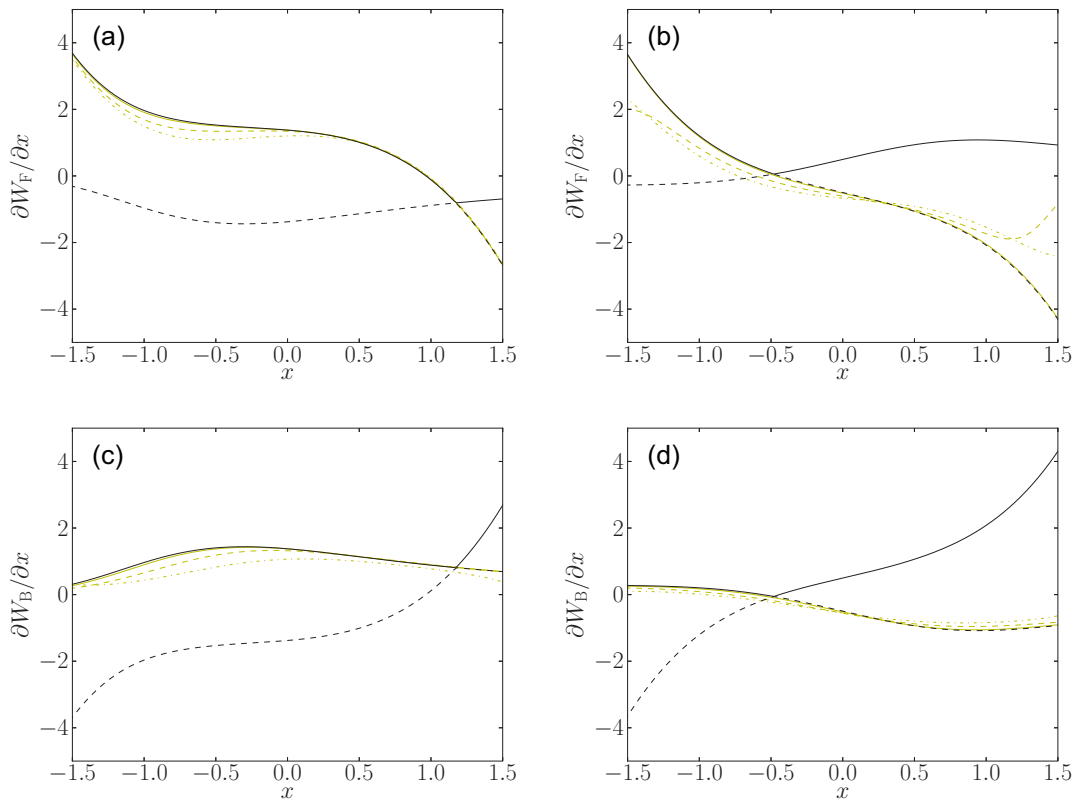

FIG. 9. The functions $\partial W_{\rm F}(x,t)/\partial x$ (a, b) and $\partial W_{\rm B}(x,t)/\partial x$ (c, d) obtained in the large t limit by solving numerically (D15) and (D16) [yellow lines (light gray) lines]. We set $h = 1$ (a, c) and $h = -1$ (b, d). The line types correspond different values of ϵ: dash-dotted, dashed, and solid lines correspond to $\epsilon = 1$, 0.5, 0.1, respectively. To illustrate the determination of the $\pm$ sign of C_h in the analytical results (D17) and (D18), we also plot on each subfigure those results with the choice of $C_h = 1$ (for all x) as black solid lines and the choice of $C_h = -1$ (for all x) as black dashed lines. As the noise goes to zero, we observe the convergence of the functions $\partial W_{\rm F}(x,t)/\partial x$ and $\partial W_{\rm B}(x,t)/\partial x$ determined numerically at large t towards the analytical line (D17) and (D18), where the $+$ sign in $\pm$ is taken for $x < x_{\rm min}$ and the $-$ sign is taken for $x > x_{\rm min}$.

these expressions, we obtain the equations to determine $W_{\rm F}(x)$ and $W_{\rm B}(x)$ as

$$\frac{\partial W_{\rm F}(x)}{\partial x} = \frac{1}{2}\{F(x) + C_h(x)\sqrt{F(x)^2 - 4h\lambda(x) - \min_y[F(y)^2 - 4h\lambda(y)]}\} \tag{D17}$$

and

$$\frac{\partial W_{\rm B}(x)}{\partial x} = \frac{1}{2}\{-F(x) + C_h(x)\sqrt{F(x)^2 - 4h\lambda(x) - \min_y[F(y)^2 - 4h\lambda(y)]}\} \tag{D18}$$

with

$$C_h(x) = 1 \quad (x < x_{\rm min}), \tag{D19}$$

$$C_h(x) = -1 \quad (x > x_{\rm min}), \tag{D20}$$

where

$$x_{\rm min} \equiv {\rm Argmin}_x[F(x)^2 - 4h\lambda(x)]. \tag{D21}$$

Equations (D17) and (D18) are the key result in this subsection. From them, we indeed get

$$p_{\rm end}(x) \sim \exp\left\{(1/\epsilon)\int^x \frac{1}{2}[F(y) + C_h(y)\sqrt{F(y)^2 - 4h\lambda(y) - \min_z[F(z)^2 - 4h\lambda(z)]}]\, dy\right\} \tag{D22}$$

and

$$p_{\rm ave}(x) \sim \exp\left\{(1/\epsilon)\int^x C_h(y)\sqrt{F(y)^2 - 4h\lambda(y) - \min_z[F(z)^2 - 4h\lambda(z)]}\; dy\right\}. \tag{D23}$$

Also from the same equations, we get the most probable x in $p_{\rm end}(x)$ and $p_{\rm ave}(x)$ with $\epsilon \to 0$. We denote them by $x_{\rm end}$ and $x_{\rm ave}$, respectively. Then, from (D22) and (D23), we find that these values satisfy

$$x_{\rm ave} = {\rm Argmax}_{x_{\rm st}}\, G_{\rm st}(x_{\rm st}), \tag{D24}$$

where $G_{\rm st}(h)$ is defined in (D11), and

$$\frac{F(x_{\rm ave})^2}{4h} = \lambda(x_{\rm ave}) - \lambda(x_{\rm end}). \tag{D25}$$

Since $\frac{F(x_{\rm ave})^2}{4h} \neq 0$, $x_{\rm ave}$ and $x_{\rm end}$ are different from each other. In other words, $p_{\rm ave}$ and $p_{\rm end}$ concentrate on different values of their argument in the $\epsilon \to \infty$ limit, as demonstrated in the main text.

For checking the validity of the obtained expressions, we numerically solve the equations (D15) and (D16) during a sufficiently large time interval t. We set $h = 1$ [Figs. 9(a) and 9(c)] and $h = -1$ [Figs. 9(b) and 9(d)]. The different colors represent the different values of ϵ: yellow, blue, and red lines correspond to $\epsilon = 1,\ 0.5,\ 0.1$, respectively. In the same figure, we plot the analytical lines (D17) and (D18), with $C_h = 1$ (for all x) (black solid line) and $C_h = -1$ (for all x) (black dashed line). We can see the convergence of the numerical lines (with decreasing ϵ) towards the analytical lines (D17) and (D18), where a $+$ sign is chosen for $x < x_{\rm min}$ and a $-$ sign is chosen for $x > x_{\rm min}$.

[1] S. Auer and D. Frenkel, Nature (London) **409**, 1020 (2001).
[2] R. P. Sear, J. Phys.: Condens. Matter **19**, 033101 (2007).
[3] W. Ren, E. Vanden-Eijnden, P. Maragakis, and W. E, J. Chem. Phys. **123**, 134109 (2005).
[4] M. Berhanu *et al.*, Europhys. Lett. **77**, 59001 (2007).
[5] F. Bouchet and E. Simonnet, Phys. Rev. Lett. **102**, 094504 (2009).
[6] D. R. Easterling, G. A. Meehl, C. Parmesan, S. A. Changnon, T. R. Karl, and L. O. Mearns, Science **289**, 2068 (2000).
[7] C. Giardinà, J. Kurchan, and L. Peliti, Phys. Rev. Lett. **96**, 120603 (2006).
[8] J. Tailleur and J. Kurchan, Nature Phys. **3**, 203 (2007).
[9] P. L'Ecuyer, V. Demers, and B. Tuffin, in *Proceedings of the 2006 Winter Simulation Conference* (IEEE, Piscataway, NJ, 2006), pp. 137–148.
[10] R. J. Allen, P. B. Warren, and P. R. ten Wolde, Phys. Rev. Lett. **94**, 018104 (2005).
[11] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, Annu. Rev. Phys. Chem. **53**, 291 (2002).
[12] T. Nemoto and S.-I. Sasa, Phys. Rev. Lett. **112**, 090602 (2014).
[13] B. A. Berg and T. Neuhaus, Phys. Lett. B **267**, 249 (1991).
[14] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001).
[15] L. Ortiz and L. Kaelbling, in *Proceedings of the Sixteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-2000)* (Morgan Kaufmann, San Francisco, CA, 2000), pp. 446–454.
[16] P. Dupuis and H. Wang, Ann. Appl. Prob. **15**, 1 (2005).
[17] O. Cappé, R. Douc, A. Guillin, J.-M. Marin, and C. Robert, Stat. Comp. **18**, 447 (2008).
[18] H. P. Chan and T. L. Lai, Ann. Appl. Prob. **21**, 2315 (2011).
[19] J. B. Anderson, J. Chem. Phys. **63**, 1499 (1975).
[20] D. Aldous and U. Vazirani, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science, 1994* (IEEE, Piscataway, NJ, 1994), pp. 492–501.
[21] P. Grassberger, Comput. Phys. Commun. **147**, 64 (2002).

[22] P. Del Moral and J. Garnier, Ann. Appl. Probab. **15**, 2496 (2005).
[23] F. Cérou, A. Guyader, T. Lelièvre, and D. Pommier, J. Chem. Phys. **134**, 054108 (2011).
[24] J. Rolland, F. Bouchet, and E. Simonnet, J. Stat. Phys. **162**, 277 (2016).
[25] H. Touchette, Phys. Rep. **478**, 1 (2009).
[26] A. Dembo and O. Zeitouni, *Large Deviations: Techniques and Applications* (Springer, New York, 1998).
[27] L. O. Hedges, R. L. Jack, J. P. Garrahan, and D. Chandler, Science **323**, 1309 (2009).
[28] M. Picciani, M. Athènes, J. Kurchan, and J. Tailleur, J. Chem. Phys. **135**, 034108 (2011).
[29] J. K. Weber, R. L. Jack, and V. S. Pande, J. Am. Chem. Soc. **135**, 5501 (2013).
[30] A. S. J. S. Mey, P. L. Geissler, and J. P. Garrahan, Phys. Rev. E **89**, 032109 (2014).
[31] P. I. Hurtado, C. P. Espigares, J. J. del Pozo, and P. L. Garrido, J. Stat. Phys. **154**, 214 (2014).
[32] B. Derrida, J. Stat. Mech. (2007) P07023.
[33] P. I. Hurtado and P. L. Garrido, J. Stat. Mech. (2009) P02032.
[34] All the results presented here can be generalized to models with n interacting particles, by replacing $d \to nd$ and including interaction forces in $F(x)$. Markov jump processes are also workable.
[35] C. W. Gardiner, *Handbook of Stochastic Methods for Physics, Chemistry, and the Natural Sciences* (Springer, Berlin, 1983).
[36] More precisely, when the Langevin equation (1) and the time-averaged quantity (4) are written in the Stratonovich convention, one first transforms these equations to ones in Itô by using a transformation formula [35], and then redefine F and $\Lambda_{\rm d}$ so that these transformed equations have the same forms as (1), (2), (3), and (4) (and can be analyzed by the same method that is discussed throughout this paper).
[37] It is also related to quantities in equilibrium thermodynamics. When we set λ_c to be an external force, $\tau\,\Lambda_{\rm c}(\tau)$ is the work done by the external system.
[38] R. Chetrite and H. Touchette, Ann. Henri Poincaré **16**, 2005 (2015).
[39] J. L. Lebowitz and H. Spohn, J. Stat. Phys. **95**, 333 (1999).
[40] V. Lecomte, C. Appert-Roland, and F. van Wijland, J. Stat. Phys. **127**, 51 (2007).
[41] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, Phys. Rev. Lett. **98**, 195702 (2007).
[42] T. Bodineau and B. Derrida, Phys. Rev. E **72**, 066110 (2005).
[43] The matrix κ is assumed to be invertible, but the generalization to noninvertible κ is direct.
[44] H. Touchette, J. Stat. Phys. **159**, 987 (2015).
[45] V. Lecomte and J. Tailleur, J. Stat. Mech. (2007) P03004.
[46] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, J. Phys. A **42**, 075007 (2009).
[47] D_2 is also called the Kullback-Leibler divergence [68]. It is positive and measures the distance between the distributions $p_{\rm ave}$ and $p_{\rm end}$. It is zero if and only if those distributions are equal almost everywhere.
[48] C. Kipnis and C. Landim, *Scaling Limits of Interacting Particle Systems* (Springer, New York, 1999).
[49] W. Fleming and S. Mitter, Stochastics **8**, 63 (1982).
[50] R. L. Jack and P. Sollich, Eur. Phys. J.: Special Topics **224**, 2351 (2015).
[51] R. Chetrite and H. Touchette, J. Stat. Mech. (2015) P12001.
[52] W. H. Fleming, Stochastic control and large deviations, in *Future Tendencies in Computer Science, Control and Applied Mathematics* (Springer, Berlin, 1992), pp. 291–300.
[53] C. Hartmann and C. Schütte, J. Stat. Mech. (2012) P11004.
[54] V. Y. Chernyak, M. Chertkov, J. Bierkens, and H. J. Kappen, J. Phys. A **47**, 022001 (2013).
[55] R. M. L. Evans, Phys. Rev. Lett. **92**, 150601 (2004).
[56] R. L. Jack and P. Sollich, Prog. Theor. Phys. Suppl. **184**, 304 (2010).
[57] T. Nemoto and S.-I. Sasa, Phys. Rev. E **84**, 061113 (2011).
[58] C. Maes and K. Netocny, Europhys. Lett. **82**, 30003 (2008).
[59] T. Speck, A. Malins, and C. P. Royall, Phys. Rev. Lett. **109**, 195703 (2012).
[60] R. L. Jack and P. Sollich, J. Phys. A **47**, 015003 (2014).
[61] The numerical calculations exploit an "annealing procedure" in which ϵ is reduced in steps of size $\delta\epsilon = 0.01$, with the best estimate of w^* from the each step being used as the first guess for w^* in the subsequent step.
[62] These integrands, which we denote by $\tilde{m}_2(x)$ and $\tilde{D}_2(x)$, are proportional to each other when $\delta \equiv |p_{\rm ave}(x) - p_{\rm end}(x)|$ is small. These satisfy $\tilde{m}_2(x) = 2\tilde{D}_2(x) + \mathcal{O}(\delta^3)$, which is why Figs. 6(a) and 6(b) show similarities.
[63] J. P. Garrahan and I. Lesanovsky, Phys. Rev. Lett. **104**, 160601 (2010).
[64] J. M. Hickey, S. Genway, I. Lesanovsky, and J. P. Garrahan, Phys. Rev. A **86**, 063824 (2012).
[65] J. Tailleur and V. Lecomte, in *Modeling and simulation of new materials: Proceedings of modeling and simulation of new materials: Tenth granada lectures*, edited by J. Marro, P. L. Garrido, and P. I. Hurtado, AIP Conf. Proc. No. 1091 (AIP, New York, 2009), p. 212.
[66] T. Nemoto and F. Bouchet (unpublished).
[67] We can prove this fact from a linear stability analysis. See Ref. [66] for details.
[68] S. Kullback and R. A. Leibler, Ann. Math. Stat. **22**, 79 (1951).