Sampling rare trajectories using stochastic bridges

Javier Aguilar ●, Joseph W. Baron ●, Tobias Galla, and Raúl Toral ●

Instituto de Física Interdisciplinar y Sistemas Complejos IFISC (CSIC-UIB), Campus UIB, 07122 Palma de Mallorca, Spain

The numerical quantification of the statistics of rare events in stochastic processes is a challenging computational problem. We present a sampling method that constructs an ensemble of stochastic trajectories that are constrained to have fixed start and end points (so-called stochastic bridges). We then show that by carefully choosing a set of such bridges and assigning an appropriate statistical weight to each bridge, one can focus more processing power on the rare events of a target stochastic process while faithfully preserving the statistics of these rare trajectories. Further, we also compare the stochastic bridges we produce to the Wentzel-Kramers-Brillouin (WKB) optimal paths of the target process, derived in the limit of low noise. We see that the generated paths, encoding the full statistics of the process, collapse onto the WKB optimal path as the level of noise is reduced. We propose that the method can also be used to judge the accuracy of the WKB approximation at finite levels of noise.

I. INTRODUCTION

The most uncommonly occurring events in stochastic systems are often the most consequential. Instances where this unlikely-yet-important combination occurs include fade-outs of epidemics [1,2], the extinction of species in ecology [3,4], the dynamics of biological switches [5–10], the escape of a Brownian particle from a double-well potential [11,12], large fluctuations in chemical reactions [13], and the detection or prediction of rare natural disasters such as earthquakes, storms, or heavy rains [14,15]. The broad range of these applications justifies the considerable recent effort expended on developing sampling algorithms for rare events in models of stochastic phenomena [16–20].

Rare events can often be conceived of as paths in phase space connecting long-lived states. A number of different approaches exist to generate these transition paths for a given system. The celebrated Wentzel-Kramers-Brillouin (WKB) method, for example, is not only used to compute quasistationary distributions and nonequilibrium landscapes [7,10,21–24], but it also delivers paths describing rare events. This approach relies on a saddle-point approximation in the limit of weak noise. As a consequence, little can be learned from the WKB approach about the statistics of transition paths in stochastic systems with finite noise. Instead, the WKB instanton only provides information about the most likely path by which a system transits from one long-lived state to another [1,4,25–29].

Transition path sampling algorithms [30–32] and forward flux techniques [31,33,34] to sample rare events account for stochasticity with finite amplitude. Transition path sampling starts from an initial trajectory connecting two long-lived states and then uses a Metropolis scheme to systematically update this path. Forward flux techniques divide phase space (or a reduced reaction coordinate space) into patches. Transition paths are then constructed as a sequence of small segments connecting these patches. Other techniques such as so-called weighted ensemble methods [35,36] also rely on a segmentation of phase space. Methods that introduce artificial temperatures and sample rare trajectories associated to rare values of a target macroscopic quantity are also of extended use [37–39].

While these powerful tools are widely used to sample rare events, each of these approaches also has limitations. Transition path sampling methods require detailed balance [32], and forward flux algorithms are known to draw statistically biased trajectories [40]. Recent work has focused on combining the strengths of these two strategies [41].

In this paper we use so-called "stochastic bridges" [42], which pass through specified start and end points by construction, to quantify the statistics of rare trajectories. Stochastic bridges are used in physics [43–47], finance [48,49], and information processing [50]. Further, Langevin bridges have also been applied to generate trajectories connecting long-lived states of stochastic differential equations (SDEs) [51,52].

The method we present here can be summarized as follows: For a given target stochastic process, we define a bespoke stochastic bridge process. We show that the statistics of the target process can be recovered by associating a statistical weight with each stochastic bridge. This allows us to dedicate more computational effort to rare trajectories without introducing bias or interdependence. The method is flexible; the target process is fully general, detailed balance is not required, no small-noise approximation is made, and it is not limited to the sampling of trajectories associated to a target macroscopic quantity, nor does it require the introduction of artificial parameters such as temperatures.

We show further that the stochastic bridges produced provide the full statistics of the ensemble of transition paths

between long-lived states of the target process. This allows one to sample fluctuations around the WKB instanton, and thus to judge if the WKB approximation scheme is accurate at various levels of noise.

II. MOTIVATION AND SIMPLE EXAMPLE

We first consider a Markov process in discrete time, $t = 0, 1, 2, \ldots, T$, with a discrete set of states which we label x_t. The process is defined by the transition probabilities $W^t_{x \to y} = P(x_{t+1} = y | x_t = x)$ and a probability distribution $P_0(x_0)$ for the initial state x_0. As indicated by the superscript t we allow for an explicit time dependence of the transition probabilities. We will refer to the ordered sequence of states visited in a realization of the process as a *path*. We write this as $\mathcal{T} = (x_0, x_1, x_2, \ldots, x_T)$, noting that the same state can be visited multiple times along a path. The probability to observe a particular path $\mathcal{T}$ is

$$\mathcal{P}(\mathcal{T}) = P_0(x_0) W^0_{x_0 \to x_1} W^1_{x_1 \to x_2} \cdots W^{T-1}_{x_{T-1} \to x_T}. \tag{1}$$

These probabilities fully characterize the process. For example, the probability of finding the system in state x at time t is the marginal $P(x, t) = \sum_{\mathcal{T}} \mathcal{P}(\mathcal{T}) \delta_{x, x_t}$, where $\delta_{x,y}$ is the Kronecker delta.

A random walk on the set of non-negative integers is a simple example of such a process. We assume that the walker departs from a fixed state x_0. The transition rates are $W_{x \to x+1} = p$ and $W_{x \to x-1} = 1 - p$ ($0 \leqslant p \leqslant 1$). Typical paths of this process are illustrated in Fig. 1(a). In the figure the random walk is biased towards lower integers. The probability that a realization terminates at a state $x_T > x_0$ is then low, in particular when x_T is much larger than x_0, or when p is much smaller than $1/2$. It is then difficult to sample paths ending at values $x_T > x_0$ in direct simulations of the biased walk.

III. METHOD

Our strategy for sampling paths of a given target dynamics connecting specific start and end points, x_0 and x_T, respectively, is based on what we will call an "associated bridge process." This process operates backwards in time, that is, paths of this bridge process are generated from x_T to x_0. We first describe this for the case of discrete time and discrete states (generalizations are discussed below).

We define the associated bridge process via transition probabilities $\tilde{W}^t_{x \leftarrow y} = P(x_t = x | x_{t+1} = y)$. That is to say, $\tilde{W}^t_{x \leftarrow y}$ is the probability that the target process was at x at time t, given that state y is visited at time $t + 1$. To define the associated bridge process we also need to specify a probability distribution $\tilde{P}_T(x)$ for the state x_T, which must satisfy $\tilde{P}_T(x) = 0$ if $P(x, T) = 0$.

By virtue of Bayes' theorem we have

$$\tilde{W}^t_{x \leftarrow y} = W^t_{x \to y} \frac{P(x, t)}{P(y, t+1)}. \tag{2}$$

As a consequence, the probabilities $\tilde{W}^t_{x \leftarrow y}$ will in general be time dependent, even if the transition probabilities of the original model do not depend on time.

The rates $\tilde{W}^t_{x \leftarrow y}$ defined in Eq. (2) might seem reminiscent of a Doob transform of the target process [51,55–57].
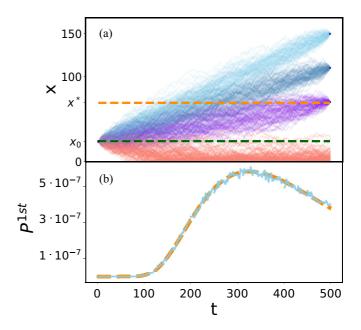

FIG. 1. (a) Red (bottom) lines show trajectories generated from simulations of the biased random walk described in the text (for $p = 0.45$), departing from $x_0 = 25$. We use reflecting boundary conditions ($W_{0 \to 1} = 1$). None of these trajectories were found to cross $x^* = 70$ before time $t = 500$. Lines in the upper part of (a) show stochastic bridges with fixed initial state $x_0 = 24$ and final states $x_T = 150$, $x_T = 130$, and $x_T = 70$, respectively (top to bottom). (b) Distribution of first-passage times at $x = x^*$ computed from our approach (see Sec. S2 of SM [53] for details). Also shown (dashed line) is the analytical result from Ref. [54].

However, it is important to note that there are significant differences between the two. In particular, the probabilities $P(x, t)$ and $P(y, t + 1)$ on the right-hand side of Eq. (2) are not conditioned on any final state x_T, so improving the overall efficacy of the sampling method [see Sec. S1 of the Supplemental Material (SM) [53] for further details].

We now focus on a fixed path $\mathcal{T} = (x_0, \ldots, x_T)$ of the target process. The path can also occur in the associated bridge process where it is traversed starting at x_T and ending at x_0. The probability to observe path $\mathcal{T}$ in the associated bridge process is

$$\tilde{\mathcal{P}}(\mathcal{T}) = \tilde{P}_T(x_T) \tilde{W}^T_{x_{T-1} \leftarrow x_T} \cdots \tilde{W}^1_{x_0 \leftarrow x_1}. \tag{3}$$

Combining Eqs. (1)–(3) we obtain a relation between the probabilities of finding path $\mathcal{T}$ in the target and associated processes respectively,

$$\mathcal{P}(\mathcal{T}) = \tilde{\mathcal{P}}(\mathcal{T}) \frac{P(x_T, T)}{\tilde{P}_T(x_T)}. \tag{4}$$

Equations (2) and (4) are the key components of our approach. We use the process defined by Eq. (2) to generate paths $\mathcal{T}$, i.e., we sample from $\tilde{\mathcal{P}}(\mathcal{T})$. Using Eq. (4) we then read off the probability with which each sample path occurs in the target process.

If we choose $\tilde{P}_T(x) = P(x, T)$ then $\mathcal{P}(\mathcal{T}) = \tilde{\mathcal{P}}(\mathcal{T})$ and paths that are rare in the target process will also be rare in the associated bridge process. Equations (2) and (4) then do not constitute an efficient sampling method for rare paths. If,

on the other hand, we choose $P_0(x) = \delta_{x,x_0}$ and $\tilde{P}_T(x) = \delta_{x,x_T}$, the trajectories generated from the process in Eq. (2) will constitute stochastic bridges connecting x_T and x_0. For such paths Eq. (4) then reduces to

$$\mathcal{P}(\mathcal{T}) = \tilde{\mathcal{P}}(\mathcal{T})P(x_T, T). \tag{5}$$

While the most demanding part of the procedure in terms of computing time is the calculation of $P(x, t)$ for $t = 0, 1, \ldots, T$, this can usually be obtained efficiently. Depending on the nature of the stochastic process (continuous/discrete states and time), $P(x, t)$ can be found with a number of established numerical methods (e.g., Refs. [58–61]).

We note that the knowledge of $P(x, t)$ does not disclose the whole statistical characterization of the enquired stochastic processes. For example, the maximum span of a trajectory or the description of transition paths is not accessible with the only knowledge of $P(x, t)$. Instead, any quantity of interest can be obtained from a representative set of paths. There is, therefore, a genuine value in estimating $P(\mathcal{T})$ using $P(x, t)$ as an input.

The complexity can be reduced further for escape paths from long-lived states. The distribution P in Eq. (2) can then be replaced by the time-independent quasistationary distribution P^{QS} describing the metastable state, which can often be approximated analytically (see, e.g., Refs. [3,25,58,62,63]). Analytical expressions for P^{QS} will make available a general expression for the rates of one-dimensional and one-step stochastic bridge process (see Sec. S4 of SM [53]). This is of particular interest since the explicit form of the generator for stochastic bridges is only known for a few systems (see, e.g., Refs. [45,51,64]). If the transition probabilities of the target process are time independent, so are then the probabilities of the associated process $\tilde{W}_{x \leftarrow y} = W_{x \rightarrow y} P^{QS}(x)/P^{QS}(y)$.

We now return to the example of a biased random walk. Figure 1(a) shows three ensembles of bridges, all starting at a fixed value of x_0 at $t = 0$ and each ensemble ending at a different choice of x_T at time $t = T$. These paths were generated from the associated process in Eq. (2), where we have chosen $\tilde{P}(x_T)$ as delta functions at the desired end points.

For a fixed final time T and choice of x_T we can determine whether the generated trajectories have crossed a fixed x^* ($x_0 < x^* < x_T$) by time T. For each trajectory we record the first crossing time. Repeating the process for different values of x_T and weighting trajectories according to Eq. (4) we then obtain the distribution of first-passage times through x^* (see Sec. S2 of the SM [53] for details). A comparison of our simulations with the analytical predictions in Ref. [54] is shown in Fig. 1(b), confirming the validity of our sampling method.

IV. CONTINUOUS-TIME PROCESSES

The method can also be used when time or the state space of the target process are continuous. The only modification is an adjustment of the expression in Eq. (2),

$$\tilde{\omega}_{x \leftarrow y}^t = \omega_{x \rightarrow y}^t \frac{P(x, t)}{P(y, t + dt)} = \omega_{x \rightarrow y}^t \frac{P(x, t)}{P(y, t)} + O(dt), \tag{6}$$

where $\omega_{x \rightarrow y}^t$ are the transition rates of the target dynamics, and $\tilde{\omega}_{x \leftarrow y}^t$ those of the associated process. For discrete states a process with time-dependent rates can be simulated for example using Lewis' thinning algorithm [65,66].

We note that $\tilde{\omega}_{x \leftarrow y}^t$ is a Gaussian bridge process [42,44,50] when the original process is Gaussian (for details see Sec. S3 of the SM [53]).

V. APPLICATIONS: MODELS OF AN EPIDEMIC AND OF A SIMPLE GENETIC SWITCH

In the context of two examples, we now compare the WKB trajectories of the target process to the associated bridge trajectories. The WKB optimal path of the target process is the most likely path that the system will take (in the limit of small noise) given the end points x_0 and x_T. At finite noise levels our method captures stochastic fluctuations about the WKB instanton. As the level of noise is reduced, we find that the realizations of the associated bridge process approach the WKB trajectory.

We first focus on the extinction of an epidemic (or "fadeout") in the individual-based susceptible-infected-susceptible (SIS) model. Conceptually this is similar to extinctions of species in ecology [67]. The model describes N individuals, of which n are infected. The population evolves in continuous time via infection and recovery processes, with rates

$$\omega_{n \rightarrow n+1} = \beta \frac{n(N - n)}{N}, \quad \omega_{n \rightarrow n-1} = \gamma n. \tag{7}$$

The parameters β and γ characterize speed of infection and recovery, respectively.

When $\beta > \gamma$, the model is known to evolve to a quasistationary "endemic" state in which the number of infected individuals fluctuates around $N(1 - \gamma/\beta)$ [68]. The population will remain in this metastable state until a large fluctuation drives the epidemic to extinction. The mean time to reach this absorbing state grows exponentially with the population size N [69]. For large populations it is therefore difficult to observe extinction in direct simulations of the SIS dynamics. Such paths can however be generated straightforwardly using the associated bridge process with rates given by Eq. (6). To evaluate these rates we use the analytical solution for the quasistationary distribution in $P^{QS}(n)$ [69] (see also Sec. S4 of the SM [53]).

An ensemble of extinction paths is shown in Fig. 2(a), along with the WKB instanton to extinction [25,63]. As illustrated in Fig. 2(a), the extinction paths at finite N fluctuate around the WKB instanton. We also show the distribution of transition times (τ) towards the infection-free state in Fig. 2(b). This timescale characterizes the duration of the transition towards extinction once the system has left the endemic state, and is not to be confused with the lifetime of the metastable state itself [3,21,25,62,63], nor with the fixed total duration of the stochastic bridges (T) (see Sec. S4 of the SM [53]). This distribution is relatively broad for small populations, but becomes more and more concentrated on the WKB estimate for larger N.

As a second example, we focus on a model of cell differentiation discussed in Ref. [7]. The two real variables $x_1 \geqslant 0$ and $x_2 \geqslant 0$ in the model describe protein concentrations, governed
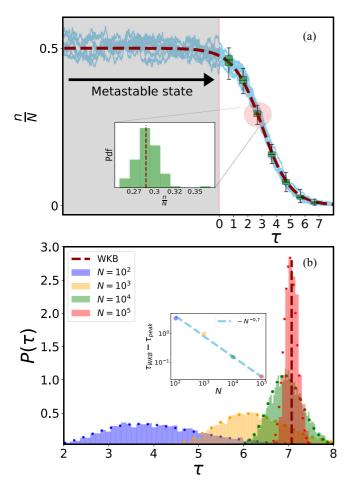
FIG. 2. Extinction paths for the SIS model ($\beta = 2, \gamma = 1$). (a) Paths leading to extinction from a common starting point $x_0 = \frac{n}{N} = 1 - \frac{\gamma}{\beta} = 0.5$ to $x_T = 0$ for $N = 10^3$, which have been shifted in time to match the WKB instanton, shown as a dashed line (see Sec. S4 of the SM [53] for additional information about this procedure). The time $\tau = 0$ corresponds to the point where the WKB instanton crosses $n/N = 0.48$. Boxes indicate the median and first quartiles, and error bars the observed range of the ensemble of stochastic paths. The inset shows the distribution of n/N at time $\tau = 2.8$. (b) Distribution of transition times for extinction trajectories from the quasistationary state towards the absorbing state (see Sec. S4 of the SM [53] for details). Dots show fits to log-normal distributions. The inset shows that the modes τ_{peak} of these fits approach the value predicted from the WKB instanton, with $|\tau_{\text{WKB}} - \tau_{\text{peak}}| \sim N^{-0.7}$.

by the SDEs

$$\dot{x}_1 = \frac{x_1^n}{S^n + x_1^n} + \frac{S^n}{S^n + x_2^n} - x_1 + \sqrt{2D}\xi_1(t),$$

$$\dot{x}_2 = \frac{x_2^n}{S^n + x_2^n} + \frac{S^n}{S^n + x_1^n} - x_2 + \sqrt{2D}\xi_2(t), \quad (8)$$

where $\xi_1(t)$ and $\xi_2(t)$ are Gaussian noise variables with mean zero and $\langle \xi_i(t)\xi_j(t')\rangle = \delta_{ij}\delta(t - t')$. The noise describes effects external to the gene circuit and its strength is governed by the model parameter $D \geqslant 0$. The deterministic terms on the right-hand side of Eqs. (8) represent self-activation, mutual inhibition, and degradation, respectively [23].

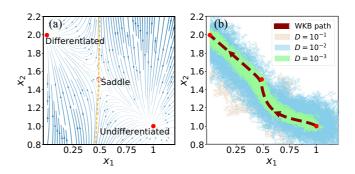FIG. 3. (a) Phase portrait of the system in Eqs. (8) for $D = 0$. Arrows indicate the deterministic flow. Due to the $x_1 \leftrightarrow x_2$ symmetry we only show one of the two fixed points describing differentiated states, and one saddle point. The corresponding separatrix is indicated as a dashed line. (b) Stochastic transition paths from the undifferentiated to one the differentiated states obtained with our sampling method. The dashed line is the WKB instanton. In both panels, $n = 4$ and $S = 0.5$.

In the deterministic limit ($D = 0$) Eqs. (8) have three stable fixed points: (i) one with $x_1 = x_2 = 1$, (ii) one with $x_1 > x_2$, and (iii) a third one obtained from the second by exchanging x_1 and x_2. The first fixed point describes an undifferentiated cell state, and the other two differentiated states [7]. The deterministic model also has two saddle points on the separatrices between the basins of attraction of the fixed points. The resulting phase portrait is shown in Fig. 3(a).

For $D > 0$, noise-driven transitions from the undifferentiated to either one of the differentiated states become possible. Similar to the SIS model the typical escape time grows exponentially with the inverse noise strength [70]. For small noise amplitudes, numerical integration of Eqs. (8) is therefore unlikely to generate transition paths within realistic computing times.

To sample these rare paths we first discretize time using an Euler-Maruyama scheme [61,71]. The resulting process is described by Gaussian transition rates. These are then used in Eq. (6) together with the quasistationary distribution describing the undifferentiated state (see Sec. S5 of the SM [53] for details). As a result of this procedure we obtain an ensemble of trajectories starting from the undifferentiated state and ending in the differentiated state $x_2 > x_1$ (see also Sec. S6 in the SM [53]). In Fig. 3(b) we show ensembles of transition paths for different noise amplitudes. As demonstrated in the figure the stochastic trajectories approach the WKB instanton as $D \to 0$.

The level of noise at which the WKB approximation is useful very much depends on the application. The above examples illustrate that the comparison between the generated stochastic bridges and the WKB instantons allows us to appraise the accuracy of WKB results.

VI. CONCLUSIONS

In this paper, we have presented a method to sample rare trajectories. The core component of the approach is a process in reverse time that generates stochastic bridges connecting desired start and end points. Using Eq. (4) we can then calculate the probability with which these paths occur in the

target process under consideration. This allows us to quantify the statistics of rare events such as first passage times. Our approach does not require the noise in the model to be weak, and it generates uncorrelated and unbiased transition paths.

Our approach goes beyond the WKB method, and delivers an ensemble of transition paths along with their statistical weights. This enables us to obtain entire distributions of first-passage times or other characteristics in simulations.

We envisage that the method that we have developed will have applications in myriad systems where sampling rare events is important. We imagine that it can also be used as a numerical aid to intuit when the WKB method will be accurate and useful. The approach presented here can also be extended

to sample stochastic trajectories constrained to pass through more than two desired points.

ACKNOWLEDGMENTS

[1] M. Assaf and B. Meerson, WKB theory of large deviations in stochastic populations, J. Phys. A: Math. Theor. **50**, 263001 (2017).

[2] A. Kamenev and B. Meerson, Extinction of an infectious disease: A large fluctuation in a nonequilibrium system, Phys. Rev. E **77**, 061107 (2008).

[3] M. Mobilia and M. Assaf, Fixation in evolutionary games under non-vanishing selection, Europhys. Lett. **91**, 10002 (2010).

[4] D. A. Kessler and N. M. Shnerb, Extinction rates for fluctuation-induced metastabilities: a real-space WKB approach, J. Stat. Phys. **127**, 861 (2007).

[5] C. Li and J. Wang, Landscape and flux reveal a new global view and physical quantification of mammalian cell cycle, Proc. Natl. Acad. Sci. USA **111**, 14130 (2014).

[6] N. J. L. Alamilla and I. S. Holek, Reconstructing the free-energy landscape associated to molecular motors processivity, Biophys. Chem. **167**, 16 (2012).

[7] J. Wang, K. Zhang, L. Xu, and E. Wang, Quantifying the Waddington landscape and biological paths for development and differentiation, Proc. Natl. Acad. Sci. USA **108**, 8257 (2011).

[8] J. N. Onuchic, Z. Luthey-Schulten, and P. G. Wolynes, Theory of protein folding: the energy landscape perspective, Annu. Rev. Phys. Chem. **48**, 545 (1997).

[9] P. B. Warren and P. R. Ten Wolde, Chemical models of genetic toggle switches, J. Phys. Chem. B **109**, 6812 (2005).

[10] B. Bhattacharyya, J. Wang, and M. Sasai, Stochastic epigenetic dynamics of gene switching, Phys. Rev. E **102**, 042408 (2020).

[11] P. Talkner and H.-B. Braun, Transition rates of a non-Markovian Brownian particle in a double well potential, J. Chem. Phys. **88**, 7537 (1988).

[12] A. Simon and A. Libchaber, Escape and Synchronization of a Brownian Particle, Phys. Rev. Lett. **68**, 3375 (1992).

[13] M. I. Dykman, E. Mori, J. Ross, and P. Hunt, Large fluctuations and optimal paths in chemical kinetics, J. Chem. Phys. **100**, 5735 (1994).

[14] A. Gabrielov, I. Zaliapin, W. I. Newman, and V. I. Keilis-Borok, Colliding cascades model for earthquake prediction, Geophys. J. Int. **143**, 427 (2000).

[15] C. Frei and C. Schär, Detection probability of trends in rare events: Theory and application to heavy precipitation in the alpine region, J. Clim. **14**, 1568 (2001).

[16] N. Malik and U. Ozturk, Rare events in complex systems: Understanding and prediction, Chaos **30**, 090401 (2020).

[17] F. Bouchet, J. Rolland, and J. Wouters, Rare event sampling methods, Chaos **29**, 080402 (2019).

[18] M. Khasin and M. I. Dykman, Control of rare events in reaction and population systems by deterministically imposed transitions, Phys. Rev. E **83**, 031917 (2011).

[19] R. Hurtado-Gutiérrez, F. Carollo, C. Pérez-Espigares, and P. I. Hurtado, Building Continuous Time Crystals from Rare Events, Phys. Rev. Lett. **125**, 160601 (2020).

[20] F. Carollo, J. P. Garrahan, I. Lesanovsky, and C. Pérez-Espigares, Making rare events typical in Markovian open quantum systems, Phys. Rev. A **98**, 010103(R) (2018).

[21] A. D. Wentzell, *Random Perturbations of Dynamical Systems* (Springer, Berlin, 1998).

[22] J. Wang, L. Xu, and E. Wang, Potential landscape and flux framework of nonequilibrium networks: robustness, dissipation, and coherence of biochemical oscillations, Proc. Natl. Acad. Sci. USA **105**, 12271 (2008).

[23] J. Wang, K. Zhang, and E. Wang, Kinetic paths, time scale, and underlying landscapes: A path integral framework to study global natures of nonequilibrium systems and networks, J. Chem. Phys. **133**, 125103 (2010).

[24] L. Ye, Z. Song, and C. Li, Landscape and flux quantify the stochastic transition dynamics for p53 cell fate decision, J. Chem. Phys. **154**, 025101 (2021).

[25] P. Ashcroft, The WKB method: A user-guide, in *The Statistical Physics of Fixation and Equilibration in Individual-Based Models* (Springer, Berlin, 2016), p. 127.

[26] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications* (Springer, Berlin, 1998).

[27] M. Heymann and E. Vanden-Eijnden, The geometric minimum action method: A least action principle on the space of curves, Commun. Pure Appl. Math. **61**, 1052 (2008).

[28] T. Grafke, R. Grauer, and T. Schäfer, The instanton method and its numerical implementation in fluid mechanics, J. Phys. A: Math. Theor. **48**, 333001 (2015).

[29] T. Grafke and E. Vanden-Eijnden, Numerical computation of rare events via large deviation theory, Chaos **29**, 063118 (2019).

[30] C. Dellago, P. G. Bolhuis, F. S. Csajka, and D. Chandler, Transition path sampling and the calculation of rate constants, J. Chem. Phys. **108**, 1964 (1998).

[31] C. Dellago and P. G. Bolhuis, Transition path sampling and other advanced simulation techniques for rare events, in *Advanced Computer Simulation Approaches for Soft Matter Sciences III*, edited by C. Holm and K. Kremer (Springer, Berlin, 2009), pp. 167–233.

[32] P. G. Bolhuis and C. Dellago, Trajectory-based rare event simulations, in *Reviews in Computational Chemistry*, edited by K. B. Lipkowitz (Wiley, Hoboken, NJ, 2010), Vol. 27, Chap. 3, pp. 111–210.

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

[34] J. T. Berryman and T. Schilling, Sampling rare events in nonequilibrium and nonstationary systems, J. Chem. Phys. **133**, 244101 (2010).

[35] G. A. Huber and S. Kim, Weighted-ensemble Brownian dynamics simulations for protein association reactions, Biophys. J. **70**, 97 (1996).

[36] R. M. Donovan, J. J. Tapia, D. P. Sullivan, J. R. Faeder, R. F. Murphy, M. Dittrich, and D. M. Zuckerman, Unbiased rare event sampling in spatial stochastic systems biology models using a weighted ensemble of trajectories, PLoS Comput. Biol. **12**, e1004611 (2016).

[37] G. Claussen, A. K. Hartmann, and S. N. Majumdar, Convex hulls of random walks: Large-deviation properties, Phys. Rev. E **91**, 052104 (2015).

[38] A. K. Hartmann, S. N. Majumdar, H. Schawe, and G. Schehr, The convex hull of the run-and-tumble particle in a plane, J. Stat. Mech.: Theory Exp. (2020) 053401.

[39] J. Tailleur and V. Lecomte, Simulation of large deviation functions using population dynamics, in *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, Melville, NY, 2009), pp. 212–219.

[40] T. S. Van Erp, Dynamical rare event simulation techniques for equilibrium and nonequilibrium systems, in *Kinetics and Thermodynamics of Multistep Nucleation and Self-Assembly in Nanoscale Materials*, edited by G. Nicolis and D. Maes, Advances in Chemical Physics Vol. 151 (Wiley, 2012).

[41] P. Buijsman and P. Bolhuis, Transition path sampling for nonequilibrium dynamics without predefined reaction coordinates, J. Chem. Phys. **152**, 044108 (2020).

[42] D. Gasbarra, T. Sottinen, and E. Valkeila, Gaussian bridges, in *Stochastic Analysis and Applications*, edited by F. E. Benth, G. Di Nunno, T. Lindstrøm, B. Øksendal, and T. Zhang (Springer, Berlin, 2007), pp. 361–382.

[43] O. Bénichou, P. Krapivsky, C. Mejía-Monasterio, and G. Oshanin, Joint distributions of partial and global maxima of a Brownian bridge, J. Phys. A: Math. Theor. **49**, 335002 (2016).

[44] M. Delorme and K. J. Wiese, Extreme-value statistics of fractional Brownian motion bridges, Phys. Rev. E **94**, 052105 (2016).

[45] B. D. Bruyne, S. N. Majumdar, and G. Schehr, Generating constrained run-and-tumble trajectories, J. Phys. A: Math. Theor. **54**, 385004 (2021).

[46] E. Barkai, E. Aghion, and D. A. Kessler, From the Area under the Bessel Excursion to Anomalous Diffusion of Cold Atoms, Phys. Rev. X **4**, 021036 (2014).

[47] B. De Bruyne, S. N. Majumdar, and G. Schehr, Generating discrete-time constrained random walks and Lévy flights, Phys. Rev. E **104**, 024117 (2021).

[48] F. Mori, S. N. Majumdar, and G. Schehr, Time Between the Maximum and the Minimum of a Stochastic Process, Phys. Rev. Lett. **123**, 200201 (2019).

[49] F. Mori, S. N. Majumdar, and G. Schehr, Distribution of the time between maximum and minimum of random walks, Phys. Rev. E **101**, 052111 (2020).

[50] L. A. Mengütürk, Gaussian random bridges and a geometric model for information equilibrium, Physica A **494**, 465 (2018).

[51] S. N. Majumdar and H. Orland, Effective Langevin equations for constrained stochastic processes, J. Stat. Mech.: Theory Exp. (2015) P06039.

[52] H. Orland, Generating transition paths by Langevin bridges, J. Chem. Phys. **134**, 174114 (2011).

[53] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevE.105.064138 for additional information and remarks, which includes Refs. [72–76].

[54] W. Kager, The hitting time theorem revisited, Am. Math. Mon. **118**, 735 (2011).

[55] J. L. Doob, Conditional Brownian motion and the boundary limits of harmonic functions, Bul. Soc. Math. France **85**, 431 (1957).

[56] D. A. Levin and Y. Peres, *Markov Chains and Mixing Times*, 2nd ed. (American Mathematical Society, Providence, RI, 2010).

[57] R. Chetrite and H. Touchette, Nonequilibrium Markov processes conditioned on large deviations, Ann. Henri Poincaré **16**, 2005 (2015).

[58] R. Dickman, Numerical analysis of the master equation, Phys. Rev. E **65**, 047701 (2002).

[59] H. Risken, *The Fokker-Planck Equation* (Springer, Berlin, 1996).

[60] J. A. Kromer, L. Schimansky-Geier, and R. Toral, Weighted-ensemble Brownian dynamics simulation: Sampling of rare events in nonequilibrium systems, Phys. Rev. E **87**, 063311 (2013).

[61] R. Toral and P. Colet, *Stochastic Numerical Methods: An Introduction for Students and Scientists* (Wiley-VCH, Weinheim, Ger., 2014).

[62] B. Meerson and P. V. Sasorov, WKB theory of epidemic fade-out in stochastic populations, Phys. Rev. E **80**, 041130 (2009).

[63] M. Assaf and B. Meerson, Extinction of metastable stochastic populations, Phys. Rev. E **81**, 021116 (2010).

[64] J. Grela, S. N. Majumdar, and G. Schehr, Non-intersecting Brownian bridges in the flat-to-flat geometry, J. Stat. Phys. **183**, 49 (2021).

[65] P. W. Lewis and G. S. Shedler, Simulation of nonhomogeneous Poisson processes by thinning, Nav. Res. Logist. Q. **26**, 403 (1979).

[66] Y. Ogata, On Lewis' simulation method for point processes, IEEE Trans. Inf. Theory **27**, 23 (1981).

[67] M. S. Bartlett, *Stochastic Population Models in Ecology and Epidemiology* (Wiley, New York, 1960).

[68] J. Marro and R. Dickman, *Nonequilibrium Phase Transitions in Lattice Models* (Cambridge University Press, Cambridge, UK, 2005).

[69] I. Nåsell, The quasi-stationary distribution of the closed endemic SIS model, Adv. Appl. Probab. **28**, 895 (1996).

[70] P. Hänggi, P. Talkner, and M. Borkovec, Reaction-rate theory: Fifty years after Kramers, Rev. Mod. Phys. **62**, 251 (1990).

[71] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations* (Springer, Berlin, 1992).

[72] S. Särkkä and A. Solin, *Applied Stochastic Differential Equations*, Vol. 10 (Cambridge University Press, Cambridge, UK, 2019).

[73] K. L. Chung and J. B. Walsh, *Markov Processes, Brownian Motion, and Time Symmetry*, Grundlehren der

mathematischen Wissenschaften, Vol. 249 (Springer, New York, 2005).

[74] L. J. Allen and A. M. Burgin, Comparison of deterministic and stochastic SIS and SIR models, Dept. Math. Stat. Tech. Rep. Ser., 98 (1998).

[75] H. Touchette, The large deviation approach to statistical mechanics, Phys. Rep. **478**, 1 (2009).

[76] C. W. Gardiner, *Handbook of Stochastic Methods*, Vol. 3 (Springer, Berlin, 1985).

SUPPLEMENTAL MATERIAL
Sampling rare trajectories using stochastic bridges

Javier Aguilar, Joseph W. Baron, Tobias Galla, Raúl Toral

Instituto de Física Interdisciplinar y Sistemas Complejos IFISC (CSIC-UIB), Campus UIB, 07122 Palma de Mallorca, Spain.

S1. RELATION WITH DOOB'S h-TRANSFORM

Despite a superficial resemblance, Eq. (2) is different from Doob's so-called h-transform [56, 72]. In this section we provide details regarding these differences. We also explain in what sense the method based on the associated process in reverse time [Eq. (2)] is more efficient for our purposes than the use of the Doob's h-transform.

A. Doob's h-transform

We first focus on discrete-time processes with transition probabilities $W^t_{x \to y}$. For a given function function $h(\cdot, \cdot)$ Doob's h-transform [55, 57, 72, 73] is then a process defined by the transition rates

$$\hat{W}^t_{x \to y} = W^t_{x \to y} \frac{h(y, t+1)}{h(x, t)}. \tag{S1}$$

The positive function h can be chosen arbitrarily provided that it fulfills

$$h(x, t) = \sum_y W^t_{x \to y} \, h(y, t+1). \tag{S2}$$

This condition ensures the normalisation $\sum_y \hat{W}^t_{x \to y} = 1$.

If time is continuous a very similar definition applies

$$\hat{w}^t_{x \to y} = w^t_{x \to y} \frac{h(y, t)}{h(x, t)}, \tag{S3}$$

with the condition

$$h(x, t) = \int dy \, w^t_{x \to y} h(y, t). \tag{S4}$$

B. Construction of stochastic bridges from the Doob transform

A suitable choice of the function $h(x, t)$ in Eq. (S1) allows one to generate conditioned Markov processes. Consider for example

$$h(x, t) = P(x_T | x_t = x), \tag{S5}$$

where $P(x_T | x_t = x)$ is the probability that a trajectory of the original process visits x_T at time T, given that it was at x at time t. We then have Eq. (S1) as follows,

$$\hat{W}^t_{x \to y} = W^t_{x \to y} \frac{P(x_T | x_{t+1} = y)}{P(x_T | x_t = x)} = P(x_{t+1} = y | x_t = x, x_T). \tag{S6}$$

We have used the definition of conditional probabilities, and the fact that $P(x_T | x_t = x, x_{t+1} = y) = P(x_T | x_{t+1} = y)$ for Markov processes.

The last expression in Eq. (S6) indicates that $\hat{W}^t_{x \to y}$ is the probability that a trajectory of the original process ultimately arriving at x_T hops from to y in the next step if it is at x at time t. Therefore, the $\hat{W}^t_{x \to y}$ are the transition rates one would obtain from the ensemble of trajectories of the original process that end at x_T.

The transition rate $\hat{W}^t_{x \to y}$ is non-zero only when $P(x_T | x_{t+1} = y) > 0$. This means that at any time t the process defined by these rates can only jump to a state y if the desired final state x_T can be reached from y in the remaining time. This statement hold in particular at time $t = T - 1$. Therefore all trajectories of the process must end in x_T.

C.   Relation to sampling method in the main paper

The transition rates $\tilde{W}_{x \leftarrow y}$ defined in Eq. (2) of the main paper are different to those of the Doob transform $\hat{W}^t_{x \to y}$ defined in Eq. (S6), but they produce an ensemble of paths with exactly the same statistics. That is, the rates $\hat{W}^t_{x \to y}$ define a process that runs forwards in time from x_0 to x_T and the rates $\tilde{W}_{x \leftarrow y}$ define a process that runs backwards in time from x_T to x_0, but each method produces the same set of paths with equal weights.

However, note that in order to compute the rates in Eq. (S6) and simulate one bridge (with fixed start and end points), one requires the conditional probabilities $P(x_T | x_t = y)$ for all values of y and a given value of x_T. This set of probabilities can be obtained efficiently from the backward master equation. Eq. (2) instead involves the conditional probabilities $P(x_t | x_0)$ for all values of x_t and a fixed value of x_0. These conditional probabilities can be obtained most efficiently from the forward master equation.

The crucial difference between the two implementations is in computational efficiency in the case where the end point x_T is not fixed, i.e. when we want to produce stochastic bridges with different end points as in Fig. 1 (a). In our algorithm, we only have to integrate the forward master equation once to obtain the probabilities $P(x_t | x_0)$. However, if we were to use the Doob transform approach, we would have to integrate the backward master equation many times to obtain different sets of probabilities $P(x_T | x_t = y)$ for each final point x_T. When one wishes to sum over many final points, as we did to produce Fig. 1 b, a many-fold increase in efficiency is obtained by using our associated bridge process over the Doob transform.

S2.   RECOVERING THE STATISTICS OF THE TARGET PROCESS FROM THE ASSOCIATED BRIDGE PROCESS

A.   General procedure

In Eq. (4) in the main text, we demonstrate how the probability of observing a particular path of the associated bridge process can be related to observing that same path in the target process. In this section, we describe how this result can be used to deduce statistics of the target process from the statistics of the associated bridge process.

We write $\mathcal{O}$ for an observable related to an individual path. For example, this could be the first time the path crosses a given barrier, or the largest excursion from the starting point, or indeed quantities relating to multiple points in time (e.g. the 'range' of a path, i.e., the difference between

maximum and minimum position attained by a random walker). We write $\mathcal{O}(\mathcal{T})$ for the value of $\mathcal{O}$ associated with path $\mathcal{T}$.

Suppose now that for each possible end state x_T we sample M paths $\mathcal{T}_{x_T}^{(i)}$ ($i = 1, \ldots, M$) ending at x_T. The statistical weight of each of the paths ending at a particular value x_T is $P(x_T, T)/M$ by virtue of Eq. (4), and accounting for the fact that we have generated M paths for each value of x_T. The overall mean of $\mathcal{O}$ over all paths of the original system (without conditioning on a particular end state) is then obtained by summing over all possible end states x_T,

$$\langle \mathcal{O} \rangle = \sum_{x_T} \left[P(x_T, T) \frac{1}{M} \sum_{i=1}^{M} \mathcal{O}(\mathcal{T}_{x_T}^{(i)}) \right]. \tag{S7}$$

For simplicity, we have chosen to sample the same number of trajectories M from all the possible end states x_T. It is possible to sample different numbers of paths M_{x_T} for different end states x_T. Eq. (S7) then becomes $\langle \mathcal{O} \rangle = \sum_{x_T} \left[P(x_T, T) \frac{1}{M_{x_T}} \sum_{i=1}^{M_{x_T}} \mathcal{O}(\mathcal{T}_{x_T}^{(i)}) \right]$.

B. Application to random walk [Fig. 1(b) in the main manuscript]

We now give further details of this procedure for the example in Fig. 1(b), namely the first passage time distribution $P^{1\text{st}}(x^*, t)$ of the biased random walk. The quantity $P^{1\text{st}}(x^*, t)$ is the probability that a walker first reaches x^* at time t, i.e., $P^{1\text{st}}(x^*, t) = P(x(t) = x^* | x(t') \neq x^* \; \forall t' < t)$.

The procedure that we follow to compute $P^{1\text{st}}(x^*, t)$ for a fixed t is as follows.

1. For a fixed value of x_T, use the rates in Eq. (2) to construct M trajectories $\mathcal{T}_{x_T}^i$ ($i = 1, \ldots, M$) of the random walk all ending at x_T.

2. Determine how many of these M trajectories first reach state x at time t. Call this number $N_{x_T}(x^\star, t)$.

3. Repeat items 1 and 2 for all possible end states x_T.

4. The quantity $P^{1\text{st}}(x^*, t)$ is then given by

$$P^{1\text{st}}(x^*, t) = \sum_{x_T} P(x_T, T) \frac{N_{x_T}(x^*, t)}{M}. \tag{S8}$$

If the random walk is run up to time T, then all $x_0 + T + 1$ final states from $x = 0$ to $x = x_0 + T$ are possible (recall that the state zero is reflecting for the example in Fig. 1). In order to limit this number at large times T, we introduce a reflecting boundary $W_{L \to L-1} = 1$ in our simulations. Provided that $L \gg x^*$ this will not materially affect the empirical first passage time distribution through x^* (in our simulations we use $x^* = 70$, $L = 200$). Further, we have repeated the process for multiple final times $T = 250, 300, 350, 400, 450, 500$, and have then combined results.

S3. REVERSED LANGEVIN EQUATION FOR GAUSSIAN PROCESSES

In this Section we show that the method described in the main text reproduces known results for Gaussian bridges [43]. Consider the following linear stochastic differential equation

$$\dot{x} = a(t)x + \sqrt{2D(t)}\xi(t). \tag{S9}$$

where we allow for a possible time-dependence of the drift coefficient $a(t)$ and of the diffusion coefficient $D(t)$. We choose the initial condition $x(t = 0) = x_0$. The distribution of x at time t can be obtained directly from the Fokker-Planck equation describing this process [59], and reads

$$P(x,t) = \frac{e^{-\frac{(x-\mu(t))^2}{2\sigma(t)^2}}}{\sigma(t)\sqrt{2\pi}} \equiv G\left(x; \mu(t), \sigma(t)\right), \tag{S10}$$

where $\mu(t) = \langle x(t) \rangle$ and $\sigma(t)^2 = \langle x(t)^2 \rangle - \langle x(t) \rangle^2$ and we introduce the shorthand $G(\cdot; \mu, \sigma)$ for a Gaussian distribution with mean μ and standard deviation σ. We have $\mu(t = 0) = x_0$, $\sigma(t = 0) = 0$.

We next discretize time into intervals of length Δt using the Euler-Maruyama prescription [61],

$$x_{t+\Delta t} = x_t + a(t)x_t\Delta t + \sqrt{2D(t)\Delta t}\Gamma_t, \tag{S11}$$

where Γ_t are uncorrelated Gaussian variables of mean zero and variance one. Alternatively, this scheme can be understood as the generation of a random value $x_{t+\Delta t}$ using the Gaussian transition probabilities

$$W^t_{x_t \to x_{t+\Delta t}} = G\left(x_{t+\Delta t}; x_t + a(t)x_t\Delta t, \sqrt{2\Delta t D(t)}\right). \tag{S12}$$

Inserting Eqs. (S12) and (S10) in Eq. (2) in the main paper we find

$$\tilde{W}^t_{x_t \leftarrow x_{t+\Delta t}} = G\left(x_{t+\Delta t}; x_t + a(t)x_t\Delta t, \sqrt{2\Delta t D(t)}\right) \times \frac{G\left(x_t; \mu(t), \sigma(t)\right)}{G\left(x_{t+\Delta t}; \mu(t+\Delta t), \sigma(t+\Delta t)\right)}. \tag{S13}$$

The right-hand side is the exponential of a quadratic polynomial in the variable x_t and hence $\tilde{W}^t_{x_t \leftarrow x_{t+\Delta t}}$, has a Gaussian functional form. We note that $\mu(t)$ and $\mu(t + \Delta t)$ feature on the right-hand side, and similarly $\sigma(t)$ and $\sigma(t + \Delta t)$. Expanding the coefficients of the polynomial in powers of Δt one finds for the reversed transition probabilities:

$$\tilde{W}^t_{x_t \leftarrow x_{t+\Delta t}} = G\left(x_t; x_{t+\Delta t} + \tilde{f}(x_{t+\Delta t})\Delta t, \sqrt{2\Delta t D(t)}\right) + \mathcal{O}((\Delta t)^2), \tag{S14}$$

where

$$\tilde{f}(x) = -a(t)x + (\mu(t) - x)\frac{2D(t)}{\sigma^2(t)}. \tag{S15}$$

Restoring continuous time, these transition probabilities in turn correspond to the stochastic differential equation for a Gaussian process $\tilde{x}(t)$:

$$\dot{\tilde{x}} = -a(t)\tilde{x} + [\mu(t) - \tilde{x}]\frac{2D(t)}{\sigma^2(t)} + \sqrt{2D(t)}\xi(t). \tag{S16}$$

which generalizes previous known results in the literature, e.g. [57]. The drift term for the process $\tilde{x}$ consists of two contributions. The first, $-a(t)$, is the drift term of the original process [Eq. (S9)], but with reversed sign. The second contribution, $[\mu(t) - \tilde{x}]\frac{2D(t)}{\sigma^2(t)}$ pulls trajectories towards the mean value $\mu(t)$ with a strength that at time t is inversely proportional to the variance $\sigma^2(t)$ of the original process. The pull becomes infinite as $t \to 0$ given that $\sigma(t = 0) = 0$, due to the fixed initial condition x_0 of the process $x(t)$. As a consequence all trajectories $\tilde{x}$ take the value $\tilde{x}(t = 0) = \mu(0) = x_0$.

Example: Brownian motion. As an example we consider simple Brownian motion,

$$\dot{x} = \sqrt{2D}\xi(t), \tag{S17}$$

with initial condition $x(0) = 0$. We then have $\mu(t) = 0$ for all t, and $\sigma(t) = \sqrt{2Dt}$. The stochastic differential equation for the associated reverse process $\tilde{x}$ is

$$\dot{\tilde{x}} = -\frac{\tilde{x}}{t} + \sqrt{2D}\xi(t), \tag{S18}$$

This equation is to be integrated backwards, that is, from $t = T$ to $t = 0$. For any starting point $\tilde{x}(t = T) = x_T$ for this integration, the resulting trajectory $\tilde{x}$ will necessarily end in $\tilde{x}(t = 0) = 0$.

S4. SUSCEPTIBLE-INFECTED-SUSCEPTIBLE MODEL AND WKB METHOD

A. Quasi-stationary distribution and reverse process

To determine the quasi-stationary distribution of the SIS model we follow the standard approach described for example in [1, 25–29, 63]. We write $x = n/N$ for the fraction of infected individuals in the following. Using the large-deviation ansatz $P(n, t) \propto e^{-NS_0(x,t)}$ in the master equation describing the model and expanding to first order in $\frac{1}{N}$, one obtains a partial differential equation for $S_0(x, t)$ [25],

$$\frac{\partial S_0(x, t)}{\partial t} = -\sum_{\ell=+1, -1} \frac{\omega_{x \to x + \frac{\ell}{N}}}{N} \left(e^{\ell \frac{\partial S_0(x,t)}{\partial x}} - 1 \right). \tag{S19}$$

Seeking stationary solutions $S_0(x) = \lim_{t \to \infty} S_0(x, t)$, one finds the following ordinary differential equation for $S_0(x)$:

$$\frac{dS_0(x)}{dx} = \log \left(\frac{\omega_{x \to x - \frac{1}{N}}}{\omega_{x \to x + \frac{1}{N}}} \right). \tag{S20}$$

Using the rates of the SIS model [Eqs. (7) in the main paper] this can be solved directly, leading to

$$S_0(x) = x \left(1 - \log \frac{\beta}{\gamma} \right) + (1 - x) \log(1 - x) - 1 + \frac{\gamma}{\beta} + \log \frac{\beta}{\gamma}, \tag{S21}$$

where the integration constant was fixed by imposing that $S_0(x)$ is zero at its minimum. The quasi-stationary distribution is then

$$P^{QS}(x) = \mathcal{N} e^{-NS_0(x)}, \tag{S22}$$

with a suitable normalisation constant $\mathcal{N}$. Expressions consistent with S22 can be derived from other means [69, 74]. In these works it is shown that this approximation is valid for any $n > 0$.

To compute the rates of the reversed process, we use Eqs. (6) after replacing the time-dependent probabilities by the quasi-static distribution in Eq.(S22). We find,

$$\tilde{\omega}_{n+1 \leftarrow n} = \gamma(n + 1)\frac{P^{QS}(n + 1)}{P^{QS}(n)},$$

$$\tilde{\omega}_{n-1 \leftarrow n} = \beta \frac{(n - 1)(N - n + 1)}{N} \frac{P^{QS}(n - 1)}{P^{QS}(n)}, \tag{S23}$$

These rates were used for the generation of the stochastic trajectories in Fig. 2.

B. WKB instanton

Again following along the lines of [1, 25–29, 63], we introduce $p = \frac{\partial H(x,p)}{\partial x}$ and $H(x,p) = \sum_{\ell=+1,-1} \frac{\omega_{x \to x + \frac{\ell}{N}}}{N} (e^{\ell p} - 1)$. Eq. (S19) is then recognised as a Hamilton-Jacobi equation

$$\frac{\partial S_0(x,t)}{\partial t} = -H\left(x, \frac{\partial S_0(x,t)}{\partial x}\right). \tag{S24}$$

Using the rates of the SIS model in Eq. (7) one obtains

$$H(x,p) = \beta x(1-x)\left(e^p - 1\right) + \gamma x \left(e^{-p} - 1\right), \tag{S25}$$

and the associated Hamilton equations

$$\dot{x} = \frac{\partial H(x,p)}{\partial p} = \beta x(1-x)e^p - \gamma x e^{-p},$$
$$\dot{p} = -\frac{\partial H(x,p)}{\partial x} = \beta(2x-1)\left(e^p - 1\right) - \gamma \left(e^{-p} - 1\right). \tag{S26}$$

These equations have four fixed points $(x = 0, p = 0)$, $(x = 0, p = \log(\gamma/\beta))$, $(x = x_{\text{st}}, p = 0)$, $\left(x = \frac{4\beta - \gamma(1 + \sqrt{1 + 8\beta/\gamma})}{8\beta}, p = \log(\frac{\gamma(\sqrt{1+8\beta/\gamma}-1)}{2\beta})\right)$, with $x_{\text{st}} \equiv 1 - \gamma/\beta$. Fig. S1 shows the solutions of Eqs. (S26) in phase space.

Following [75], we will refer to solutions of Hamilton's equations with two fixed endpoints as instantons. The WKB instanton is the most likely trajectory for the system to take in the limit $N \to \infty$ given any start and end points x_0 and x_T. For the case of the SIS model, it is found by solving Eq. (S26).

The instanton connecting $(x = x_{\text{st}}, p = 0)$ with $(x = 0, p = \log(\gamma/\beta))$ is the dominant path to extinction, shown as a dashed line in Fig. 2a. This trajectory fulfills Eqs. (S26) and $H(x(t), p(t)) = 0$ (since $H(x = 0, p) = 0$ and H is conserved along trajectories obeying Hamilton's equations). The constraint $H = 0$ leads to

$$\dot{x} = -x\left[(1-x)\beta - \gamma\right], \tag{S27}$$

with solution

$$x(t) = \frac{x_{\text{st}}}{1 + e^{\beta x_{\text{st}}(t-t_0)}}. \tag{S28}$$

The quantity t_0 is an integration constant and can be chosen arbitrarily (the WKB instanton in Fig. 2 is for $t_0 = 3.178$ so that $x(0) = 0.48$). Here, it is useful to distinguish between two phases of the dynamics of the SIS model. After some initial transient the systems reaches quasi-stationarity in an endemic state. The fraction of infected individuals fluctuates about $x_{\text{st}} = 1 - \gamma/\beta$. This is a long-lived state, the residence time is a random variable with a mean which increases exponentially with the population size N. Once the system has left this long-lived state by a random fluctuation, it will approach the absorbing state $(n = 0)$ on a much shorter time scale of order N^0.

We stress that the WKB instanton in Eq. (S28) describes the relatively quick transition towards absorption after the system has left the quasi-stationary state. The preceding residence time in the quasi-stationary state is separate and random, and the WKB instanton makes no statement about this time (hence the open parameter t_0). The mean residence time in the stationary state can be computed separately[21, 25].
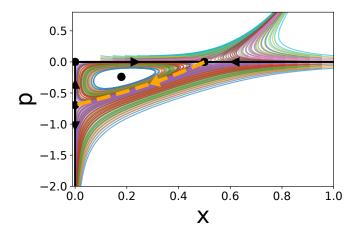
FIG. S1. Trajectories of Eqs. (S26) with initial conditions $x(t = 0) \in [0.1, 1]$ and $p(t = 0) \in [-0.1, 0.1]$. Dots are placed at the fixed points of Eqs. (S26). The heteroclinic connecting $(x = x_{\mathrm{st}}, p = 0)$ with $(x = 0, p = \log(\gamma/\beta))$ (dashed line) is the optimal path to extinction plotted in Fig. 2. This heteroclinic is the only solution connecting the line $x = x_{st}$ with the line $x = 0$ with all $p(t)$ finite.

C.   Transition towards extinction

The most likely path to extinction, obtained in the limit $N \to \infty$, is the heteroclinic orbit in (S28) connecting the fixed points $(x = x_{\mathrm{st}}, p = 0)$ and $(x = 0, p = \log(\gamma/\beta))$ of Eqs. (S26). This path approaches x_{st} for $t \to -\infty$, and $x = 0$ for $t \to \infty$. The typical time scale associated with the transition from the meta-stable state and the absorbing state can however be characterised by the time required for the WKB instanton to reach $x = \epsilon$ starting from $x = x_{\mathrm{st}} - \epsilon$. This time difference is independent of the choice of t_0 and is easily obtained from Eq.(S28) as

$$\tau = \frac{1}{\beta x_{\mathrm{st}}} \log \left(\frac{x_{\mathrm{st}} - \epsilon}{\epsilon^2} \right). \tag{S29}$$

In Fig. (2b) we have used the value $\epsilon = 0.02$, leading to $\tau = 7.09$.

The time at which a trajectory of the stochastic SIS model leaves the vicinity of x_{st} and crosses $x = x_{\mathrm{st}} - \epsilon$ is a random variable, determined by how long the system resides in the meta-stable state around x_{st} (see above). The system then moves towards extinction at $x = 0$. In order to align these transition paths with the WKB instanton in Eq.(S28), we shift each trajectory in time. More precisely we arbitrarily fix $t_0 = 3.178$ for the WKB instanton (so that $x(0) = 0.48$), and then choose the time shift for each trajectory of the stochastic systems so as to minimise $\int dt \, [x_{\mathrm{stochastic}}(t) - x_{\mathrm{instanton}}(t)]^2$. This is the procedure we used to produce Fig. 2 (a) in the main paper. See Fig. (S2) with instances of the raw outcome of the simulation, with no temporal shift. This figure highlight the differences of the random transition time (τ) with the fixed total duration of stochastic bridges (T).

S5.   FOKKER-PLANCK EQUATION FOR CELL DIFFERENTIATION PROBLEM

The quasi-stationary distribution for the genetic switch model is obtained from a numerical integration of the Fokker-Planck equation for the dynamics in Eqs. (8). This Fokker-Planck reads:

$$\partial_t P(x_1, x_2; t) = -\vec{\nabla} \cdot (P(x_1, x_2; t) \vec{F}(x_1, x_2)) + D\nabla^2 P(x_1, x_2; t), \tag{S30}$$
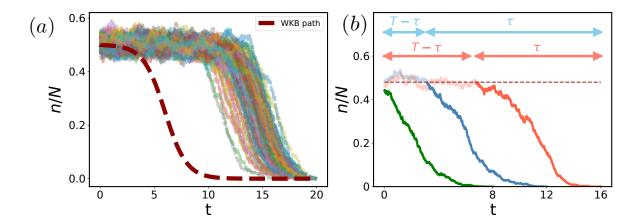
FIG. S2. In (a), WKB instanton (dashed line) and stochastic bridges for the SIS process without temporal shift. The WKB instanton take an infinite time to go from the long-lived state $\frac{I}{N} = 0.5$ to the absorbing state. We plotted such trajectory for a duration of 20 time units. The stochastic bridges have a fixed total duration $T = 20$, however, each of them execute the transition in a well-defined random time τ. The source of discrepancies between the WKB instanton and the stochastic bridges comes from the choice of T, which is arbitrary. We shifted the stochastic bridges in order get a fair comparison of the shapes of the stochastic bridges and the WKB instanton. In (b), three particular realizations of the stochastic bridges with different total duration ($T = 8, 12, 16$). Stochastic bridges with big T are likely to execute the inquired transition. This is the case of the stochastic bridges with $T = 16$ and $T = 12$ showed in the figure. Big values of T will not affect the transition times sampled (τ). If T is too small, the bridges will not have enough time to execute the transition (e.g. the bridge showed with $T = 8$). Thus, the only requirement to choose T is that it should be way longer than the typical time scales of transitions.

where $\vec{F}(x_1, x_2) = [f(x_1, x_2), f(x_2, x_1)]$ with $f(u,v) = \frac{u^n}{S^n + u^2} + \frac{S^n}{S^n + v^n} - u$. We study the system within a square domain $(x_1, x_2) \in [0, L] \times [0, L]$ with reflecting boundary conditions. In Fig. (3b) we used $L = 3$. We solved Eq. (S30) numerically.

The numerical integration of Eq. (8) requires a discretisation of space and time. We used a forward time centered space method. The time step and spatial discretisation element need to fulfill the conditions [61] $\Delta t < \frac{(\Delta x)^2}{D}$ and $\Delta t < \sqrt{\frac{\Delta x}{|\max_{x_1, x_2}\{\vec{F}(x_1, x_2)\}|}}$. The reflecting boundary conditions are imposed by requiring that the normal component of the probability flux $\vec{J}(x_1, x_2; t) = P(x_1, x_2; t)\vec{F}(x_1, x_2) - D\vec{\nabla}P(x_1, x_2; t)$ vanishes at the boundaries [76]. We find that the numerical solution of Eq. (S30) converges to a stationary distribution $P^{\mathrm{QS}}(x_1, x_2)$ relatively quickly. In Fig. (3b) we used $P(x_1, x_2; T = 10 | x_1^0 = 1, x_2^0 = 1, 0)$ as a proxy for the quasi-stationary distribution. We carried out a sensitivity analysis testing different final times $T = 10, 12, 15, 20$. This did not result in significant differences for the statistics of paths.

S6. GENERATION OF REVERSE TRAJECTORIES FOR STOCHASTIC DIFFERENTIAL EQUATIONS

In this section, we describe the strategy to numerically sample paths of the associated bridge process with a target dynamics described by an SDE of the form:

$$\dot{x}(t) = f(x, t) + g(x, t)\xi(t). \tag{S31}$$

Here, the noise $\xi(t)$ has mean zero and correlations $\langle \xi(t)\xi(t') \rangle = \delta(t - t')$, and the SDE is to be interpreted in the Itô sense. For simplicity, we describe the procedure for a univariate process, but generalization is possible to the multivariate case. Again, the Euler-Maruyama scheme to draw trajectories of the SDE is given by

$$x_{t+\Delta t} = x_t + f(x_t, t)\Delta t + g(x_t, t)\sqrt{\Delta t}\Gamma_t, \tag{S32}$$

where Γ_t are uncorrelated Gaussian variables of mean zero and variance one. The transition rates for this discretisation are

$$W^t_{x_t \to x_{t+\Delta t}} = G(x_{t+\Delta t}; x_t + f(x_t, t)\Delta t, g(x_t, t)\sqrt{\Delta t}), \tag{S33}$$

where $G(x; \mu, \sigma)$ is the Gaussian distribution defined in Eq. (S10). This distribution can be easily sampled using, for example, the Box-Muller-Wiener or any other suitable algorithm [61].

Using Eq. (2) in the main paper, the transition rates for the associated process are

$$\tilde{W}^t_{x_t \leftarrow x_{t+\Delta t}} = G\left(x_{t+\Delta t}; x_t + f(x_t, t)\Delta t, g(x_t, t)\sqrt{\Delta t}\right) \times \frac{P(x_t, t)}{P(x_{t+\Delta t}, t + \Delta t)}. \tag{S34}$$

In contrast with the case of a linear SDE [Eq. (S9)], $G\left(x_{t+\Delta t}; x_t + f(x_t, t)\Delta t, g(x_t, t)\sqrt{\Delta t}\right)$ is in general not a Gaussian distribution for x_t (as x_t appears as an argument of the functions f and g).

To sample x_t from this distribution we use a proposal-rejection technique [61]. For the proposal step we use the approximation

$$G\left(x_{t+\Delta t}; x_t + f(x_t, t)\Delta t, g(x_t, t)\sqrt{\Delta t}\right) \approx G\left(x_t; x_{t+\Delta t} - f(x_{t+\Delta t}, t + \Delta t)\Delta t, g(x_{t+\Delta t}, t + \Delta t)\sqrt{\Delta t}\right)$$
$$\equiv G^{\mathrm{prop}}(x_t). \tag{S35}$$

We then draw a proposed value x_t from the Gaussian distribution $G^{\mathrm{prop}}(x_t)$. This proposal is then accepted with a probability proportional to $\frac{\tilde{W}^t_{x_t \leftarrow x_{t+\Delta t}}}{G^{\mathrm{prop}}(x_t)}$.

This method generates values of the random variable x_t distributed according to Eq. (S34). There is no further bias or error other than those resulting from the Euler-Maruyama discretisation.

In summary, the proposal-rejection algorithm for the associated reverse process is as follows:

If the system is in state $x_{t+\Delta t} = y$ at time $t + \Delta t$ then:

1. Draw a random number x distributed according to the Gaussian distribution $G\left(x; y - f(y, t + \Delta t)\Delta t, g(y, t + \Delta t)\sqrt{\Delta t}\right)$.

2. With probability

$$H(x) = C \frac{G\left(y; x + f(x, t)\Delta t, g(x, t)\sqrt{\Delta t}\right)}{G\left(x; y - f(y, t + \Delta t)\Delta t, g(y, t + \Delta t)\sqrt{\Delta t}\right)} \frac{P(x, t)}{P(y, t + \Delta t)}, \tag{S36}$$

accept $x_t = x$. If rejected, go to step 1. Here, C is chosen such that $\max_x H(x) \leq 1$.

We have found in the examples that $C \approx 1$ and the average acceptance probability is close to 1.