Action functional gradient descent algorithm for estimating escape paths in stochastic chemical reaction networks

Praful Gagrani [1,2] and Eric Smith [1,3,4,5,6]

[1]*Department of Physics, University of Wisconsin–Madison, Madison, Wisconsin 53706, USA*
[2]*Wisconsin Institute for Discovery, University of Wisconsin-Madison, Madison, Wisconsin 53715, USA*
[3]*Department of Biology, Georgia Institute of Technology, Atlanta, Georgia 30332, USA*
[4]*Earth-Life Science Institute, Tokyo Institute of Technology, Tokyo 152-8550, Japan*
[5]*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, New Mexico 87501, USA*
[6]*Ronin Institute, 127 Haddon Place, Montclair, New Jersey 07043, USA*

We first derive the Hamilton-Jacobi theory underlying continuous-time Markov processes, and then we use the construction to develop a variational algorithm for estimating escape (least improbable or first passage) paths for a generic stochastic chemical reaction network that exhibits multiple fixed points. The design of our algorithm is such that it is independent of the underlying dimensionality of the system, the discretization control parameters are updated toward the continuum limit, and there is an easy-to-calculate measure for the correctness of its solution. We consider several applications of the algorithm and verify them against computationally expensive means such as the shooting method and stochastic simulation. While we employ theoretical techniques from mathematical physics, numerical optimization and chemical reaction network theory, we hope that our work finds practical applications with an inter-disciplinary audience including chemists, biologists, optimal control theorists and game theorists.

I. INTRODUCTION

Stochastic modeling has played an increasingly central role in science since the advent of high speed computing. There are, however, many categories of events that are vital to the organization of long-term dynamics for which the rate diminishes exponentially with system size ("rare" events) and computer simulations become unaffordable. A typical approach in these scenarios is to employ importance sampling approaches in the stochastic simulation to efficiently simulate the rare event of interest (see Refs. [1–3]). In this work, for the particular case of stochastic chemical reaction networks, we provide a deterministic alternative for estimating the likelihood of rare events.

Stochastic nonlinear dynamical processes exhibiting multistability, or multiple coexisting attractors, have found several scientific applications ranging from modeling climate change to population biology and the origin of life [4–7], and are an active subject of interdisciplinary research. A question of critical importance for any practical application of such a system is, how often does it transition out of a stable attractor and what are the least-improbable paths by which such a transition occurs? In literature, the dynamics that arise due to the system's transitioning from one stable attractor to another and the optimal path of transition is referred to as "switching dynamics" and "escape path," respectively.

Many phenomena in biology, chemistry, physics or engineering can be modeled as a chemical reaction network (CRN). It is also well known that CRNs can exhibit a host of dynamics, including limit cycles and multistability [8].

Perhaps less well-known is the role of Hamilton-Jacobi ray theory underlying stochastic CRN, and the interpretation of escape paths as particular characteristic curves that arise as a solution to the associated Hamilton-Jacobi equation. In our work, we review the necessary formalism and employ techniques from mathematical physics and numerical optimization to estimate the escape paths for a multistable CRN. More precisely, we first recognize that the escape paths are variational (locally least-action) solutions to the action functional, and subsequently use the functional to perform a gradient descent that converges on the desired path.

It must be noted that neither the recognition of the escape path as a variational solution nor using the action functional to perform gradient descent is novel to us (for instance, see Ref. [9]); however, we differ from the earlier work in a few ways. In the past, the use of Hamilton-Jacobi theory to find the escape path for CRN has only been made using the "shooting-method," which relies on integrating Hamilton's equations of motion, rather than a functional gradient descent approach that we present here (see Ref. [10]). However, while the action functional has been used for a gradient descent, to the best of our knowledge, the form of the Hamiltonian assumed has always been separable in momentum and position (see Refs. [11,12]). This form of the Hamiltonian, with other constraints on the functional form, makes it amenable to finding an analytic form for the associated Lagrangian which drastically simplifies the gradient descent. While this assumption is typical for a Hamiltonian describing mechanical energy in physics, it is far too restrictive for a generic Hamiltonian in

chemistry, as can be seen in the following:

$$H_{\text{ME}}(p, q) = K(p) + U(q),$$

$$H_{\text{CRN}}(p, q) = \sum_{y_\alpha, y_\beta} (e^{(y_\beta - y_\alpha)p} - 1) k_{y_\alpha \to y_\beta} q^{y_\alpha}, \quad (1)$$

where p and q are the momentum and position coordinates, $K(p)$ and $U(q)$ are the kinetic and potential energy, and H_{ME} and H_{CRN} [Eq. (16)] are the Hamiltonians for mechanical energy and chemical reaction networks, respectively. In our algorithm, however, we start from the Hamiltonian, numerically solve for the Lagrangian at each iteration and use it to calculate the descent direction. While our algorithm is designed for stochastic CRNs, it is amenable to generalization for other classes of stochastic Hamiltonian dynamical systems.

The layout of the paper is as follows. In Sec. II, we derive the Hamilton-Jacobi theory for stochastic processes starting from a master equation and apply it to stochastic CRN. We show that the Hamilton-Jacobi theory of the nonequilibrium potential (NEP) arises naturally as a result of a variational principle applied to the *action functional*. In Sec. III, we propose an *action functional gradient descent* (AFGD) algorithm that finds the variational solution in the space of paths constrained at end points for a given Hamiltonian value constraint. We also explain how the algorithm can be used to find least-improbable escape paths out of a stable attractor and assign a value to the NEP along them. The details of the implementation are provided in Appendix D, and a MATLAB implementation is made available at Ref. [13]. In Sec. IV, we demonstrate the applications of the AFGD algorithm on several CRNs. We first consider the Selkov model, and compare the result of the algorithm against the escape trajectory found by the "shooting-method" (Fig. 9). We then define a class of high-dimensional birth-death models, namely "N-Schlögl model," and compare the results of the algorithm against a stochastic simulation for the 2-Schlögl model (Fig. 4). We also use the algorithm on the six-dimensional (6D) 6-Schlögl model, and compare the result against the integration of Hamilton's equations of motion (Fig. 21). Finally, in Sec. V, we conclude with a discussion of our contribution and potential avenues of future research.

II. FROM MASTER EQUATION TO HAMILTON-JACOBI EQUATION FOR CRN

The application of Hamilton-Jacobi theory to chemical reaction networks (CRNs) has a long history [10,14]. In this section, we review the necessary formulation needed to understand the switching dynamics of stochastic chemical reaction networks (Sec. II C) and to devise a variational algorithm for predicting the transitions (Sec. III). In Sec. II A we derive the nonequilibrium potential for continuous-time Markov population processes and explain its role in estimating the probability of stochastic events. In Sec. II B, we derive the Hamiltonian for a CRN and prove that the NEP is a Lyapunov function along the deterministic trajectories under mass-action kinetics. In Appendix A, we investigate the relevance our work might have to a stochastic modeling practitioner by posing a general practical problem, giving an overall picture

of the solution and explaining where our algorithm fits in. In Appendix B, we employ the Hamilton-Jacobi formalism to recover the ACK theorem and NEP for complex-balanced systems. In Appendix C, we define a "nonequilibrium action," of which both the master equation and Schrödinger equation can be seen as a variational solution, derive the path integral formula and action functional for stochastic population processes, and calculate the first and second variational derivatives of the action functional. There are many resources that provide an alternative treatment of the same subject, for instance, see Refs. [15,16] and references therein.

A. Hamilton-Jacobi equation and nonequilibrium potential (NEP) for stochastic dynamics

A continuous-time Markovian population process is specified by its *master equation*,

$$\frac{\partial}{\partial t} \rho(n, t) = \sum_{n'} \mathbb{T}_{nn'} \rho(n', t),$$

$$\sum_n \rho(n, t) = 1 \quad \text{for all } t,$$

$$\sum_n \mathbb{T}_{nn'} = 0 \quad \text{for all } n', \quad (2)$$

where t is the time variable, n is a D-dimensional discrete vector in the positive integer lattice $\mathbb{Z}_{\geqslant 0}^D$ denoting a position in the state space (to simplify notation we drop the arrow in $\vec{n}$), and $\rho(n, t)$ is a time-evolving probability distribution function (PDF) over the state space. $\mathbb{T}$ is referred to as a transition operator, and in this paper we only consider time independent transition operators.

For what follows, it is useful to recast Eq. (2) into the following form and define a Hamiltonian operator $\hat{\underline{H}}$ that acts on the PDF ρ:

$$\frac{\partial}{\partial t} \rho(n, t) = \hat{\underline{H}}\left(-\frac{\partial}{\partial n}, n\right) \rho(n, t), \quad (3)$$

where $\partial/\partial n$ is the infinitesimal-shift operator.[1]

The resemblance of Eq. (3) to the Schrödinger equation is not a mere coincidence, and we show in Appendix C 1 how they can both be derived from a common variational problem of extremizing, what is termed as, the "nonequilibrium action" (NEA) in Ref. [17]. It should come as no surprise then that we can employ the same machinery developed in mathematical physics for quantum mechanics to derive results about stochastic dynamics. This line of reasoning has a long history and we refer the readers to Refs. [6,18–20] for a "second-quantization" treatment via abstract linear algebra.

We can integrate Eq. (3) starting from a PDF $\rho(n_0, 0)$ at time 0 to get a distribution $\rho(n_T, T)$ at time T, indexed by n_0

---

[1]A function in x can be shifted by y by the application of a *shift operator* given by $e^{y\frac{\partial}{\partial x}}$, as can be seen by

$$f(x + y) = \sum_{n=0}^\infty \frac{y^n}{n!} \frac{\partial^n f(x)}{\partial x^n} = e^{y\frac{\partial}{\partial x}} f(x).$$

and n_T, respectively:

$$\rho(n_T, T) = e^{\int_0^T \underline{\hat{\mathbf{H}}} dt} \rho(n_0, 0)$$

$$= \int [dn] \int [dp] e^{-\mathcal{A}[n(t), p(t)]} \rho(n_0, 0), \quad (4)$$

where $\mathcal{A}$ is the *action functional*

$$\mathcal{A}[n(t), p(t)] = \int_0^T \left[p \frac{dn}{dt} - \underline{H}(p, n) \right] dt, \quad (5)$$

and p is a momentum variable canonically conjugate to the position variable n. The second line in Eq. (4) is the path-integral formula, a proof of which can be found in Appendix C 2, and $[dn][dp]$ is the path-integral measure. Note that in Eq. (5), the Hamiltonian function $\underline{H}(p, n)$ has the same functional form as the Hamiltonian operator $\underline{\hat{H}}$ except the operator $-\partial/\partial n$ is replaced by the variable p.

In the first line of Eq. (4), the exponentiation of the Hamiltonian operator is to be understood as a time-ordered matrix product over small intervals dt. Its role is to time-evolve the initial distribution $\rho(0)$ and accumulate probability through all possible chains of states until time T to obtain a new distribution $\rho(T)$. In the second line, upon taking appropriate limits signifying a continuous-time parameter, we obtain another description of the same process where we sum over all the paths starting at n_0 at time 0 and ending at n_T at time weighted with an appropriate measure. The negative log of the measure of a particular path is given by the action functional (in other words, $\mathcal{A}$ takes as input a path and returns its log-improbability), and integrating under the path-integral measure amounts to summing over all paths. For a didactic introduction to the topic and its application to stochastic dynamics, we refer the reader to Ref. [19].

It is useful to pass from a discrete to a continuous state space by descaling n with a scale factor V and considering large V limit. To align our presentation here with the mathematical physics literature, we denote the continuous variable by $q = n/V$ and refer to the continuous state space as configuration space. There are several interpretations the variables can take; the one of interest to us is where n is a population vector (thus, $n \geqslant 0$), q is the concentration vector and V is a scale of the the total population. Following Eq. (4), we can write the time evolution of a distribution in the transformed coordinate q as

$$\rho(q, T) = \int [dq] \int [dp] e^{-V \mathcal{A}[q(t), p(t)]} \rho(q(0), 0),$$

$$\mathcal{A}[q(t), p(t)] = \int_0^T \left[p \frac{dq}{dt} - H(p, q) \right] dt, \quad (6)$$

where $\mathcal{A}$ is the action functional in the new coordinates and has a different Hamiltonian $H(p, q)$ which relates to $\underline{H}(p, n)$ by[2]

$$V H(p, q) \asymp \underline{H}(p, qV) = \underline{H}(p, n). \quad (7)$$

_____

[2]We say $f(V) \asymp g(V)$ [read as $f(V)$ is asymptotic to $g(V)$] if

$$\lim_{V \to \infty} \frac{f(V)}{g(V)} = 1.$$

It must be noted that Eq. (7) is an additional constraint on the form of the Hamiltonian, but it is one that is satisfied by the CRN Hamiltonian which we are interested in for the scope of this paper. Moreover, if the asymptotic equality was an equality, then the condition is often referred to by saying "the Hamiltonian is extensive in the position coordinate," or "the Hamiltonian is a homogenous function of degree one in the position coordinate." We will return to this point in the next subsection, once we write the explicit Hamiltonian for a CRN.

Together the position and momentum coordinates (q, p) constitute a point in the *phase space* [21]. To evaluate the distribution at time T using Eq. (6), in principle we need to sum over all the phase-space paths that end at q at time T. Evaluating the full sum, however, is not necessary to obtain the leading large-V asymptotics. Through Laplace's or saddle-point method, for large V the path integral is dominated by the saddle point of the functional and the distribution at time T can be approximated by

$$\rho(q, T) \asymp N(V) e^{-V \mathcal{A}[q^*(t), p^*(t)]} \rho(q^*(0), 0),$$

$$\text{where } (q^*(t), p^*(t)) \text{ is such that}$$

$$\delta \mathcal{A}[q^*(t), p^*(t)] = 0 \text{ and } q^*(T) = q, \quad (8)$$

and $N(V)$ is a normalization factor expounded upon in the subsequent paragraph. The constraint in the last line specifies the stationarity or optimality condition on the saddle point path $(q^*(t), p^*(t))$, also referred to as the path of stationary action, and we henceforth refer to it simply as the *optimal path* (for a MinMax formulation of the saddle-point see Sec. III A). We require here that the Hamiltonian function is convex in the momentum variable p, which we show is the case for CRN Hamiltonians with mass-action kinetics in Sec. II B.

There are two caveats to Eq. (8) that we now point out. First, the normalization factor $N(V)$ in the first line depends not only on the scale factor V but also on the optimal path. There is an in-principle method, although costly in practice, to obtain $N(V)$ from the action functional itself, outlined in Ref. [22], Ch. 7 of Ref. [23] and Ch. 4 of Ref. [6] (for an application, see Ref. [24]). Second, since the path integral is a sum over all paths, there are several stationary paths starting from different $q^*(0)$ that reach q at T, and one must perform a sum over all of them. This becomes increasingly relevant in the limit $T \to \infty$, where the system can bounce back and forth multiple times between a $q^*(0)$ and q (see Refs. [6,23]). There is, however, a twofold resolution to this problem. First, the normalization factor $N(V)$ is subexponential in V, and can be well-approximated by unity large for V (also see the introduction to Sec. III). Second, the initial distribution $\rho(q, 0)$ is typically taken to be peaked at a given value, thus yielding negligible contribution from every $q^*(0)$ that is not near the peak of the initial condition. Thus, we only pick the least-improbable direct paths that start from near the peak of the initial distribution and reach q in Eq. (8).

In the first line of Eq. (8), the only free variables are configuration q and final time T, which together characterize a time evolving distribution. The optimality condition picks a path $(q^*(t), p^*(t))$, and it is only the probability mass at $q^*(0)$ in the initial distribution at time 0, i.e., $\rho(q^*(0), 0)$, weighted appropriately by the action functional, that contributes to the

final distribution at configuration q and time T. The weight factor $e^{-V\mathcal{A}[q^*,p^*]}$ accounts for the fact that not all the probability mass from $q(0)$ goes to q in time T, and in this sense takes the interpretation of a conditional probability.

Since the conditional probability term is only in the exponential, it is useful to write the probability distribution itself as an exponential function $\rho(q,t) = e^{-VS(q,t)}$. Here $S(q,t)$ is the *action function* (not be confused with $\mathcal{A}$ which is the action functional) and its time evolution is given as follows:

$$
\begin{aligned}
S(q,T) &= \mathcal{A}[q^*(t), p^*(t)] + S(q^*(0),0) \\
&= \int_0^T \left[p^* \frac{dq^*}{dt} - H(p^*,q^*) \right] dt + S(q^*(0),0),
\end{aligned}
\tag{9}
$$

where in the second line, we expand the action functional using Eq. (6). Thus, the action function evolves by the optimal value of the action functional, and since it is a function we can also consider its total differential

$$
\begin{aligned}
S(q,T) &= \int_0^T dS + S(q^*(0),0) \\
&= \int_0^T \left[\frac{\partial S(q^*,t)}{\partial q^*} \frac{dq^*}{dt} + \frac{\partial S(q^*,t)}{\partial t} \right] dt \\
&\quad + S(q^*(0),0).
\end{aligned}
\tag{10}
$$

Comparing Eqs. (9) and (10), we get

$$
\begin{aligned}
\frac{\partial S(q^*,t)}{\partial q^*} &= p^*, \\
\frac{\partial S(q^*,t)}{\partial t} &= -H(p^*,q^*), \\
\frac{\partial S(q^*,t)}{\partial t} &= -H\left(\frac{\partial S}{\partial q^*}, q^*\right),
\end{aligned}
\tag{11}
$$

where the first two lines are obtained by comparison and the last line is obtained by substituting the first line in the second line. The nonlinear PDE in the last equation is called the Hamilton-Jacobi equation. For an introduction to the subject and its history, we refer the readers to Refs. [21,25] (what we call as the action function is also called Hamilton's principal function). We briefly remark that the definite integral of the action functional along the optimal path $\int_0^T dS$ also acts as a divergence function in information geometry (see Refs. [20,26]) and is denoted as $S(q^*(T)||q^*(0))$ yielding the relation $S(q_T^*) = S(q_T^*||q_0^*) + S(q_0^*)$, as mentioned in Eq. (A2).

The role of *momentum* for stochastic systems can be best understood by the Hamilton Jacobi equation. The first line in Eq. (11) shows that momentum is the gradient of the action function, which in turn is descaled negative log probability of the time evolving distribution given some initial conditions. Using this, we see that $p^*(0)$ must be the gradient of the action function S at $q^*(0)$, and that an optimal path in phase space corresponds to a contour in the probability distribution arising by scaling and exponentiating the action.

The third line in Eq. (2) ensures that the transition operator always has a cokernel, namely the row vector consisting of all ones $[1, \ldots, 1]$. The Perron-Frobenius theorem for stochastic matrices proves that the operator also has a unique right kernel, referred to as the stationary distribution which we denote

here by π. π thus satisfies $\sum_n \mathbb{T}_{n'n}\pi(n,t) = 0$ or equivalently

$$
\underline{\hat{\mathrm{H}}}\pi(n,t) = 0.
$$

From Eqs. (2) and (3) one can see that the stationary distribution is indeed time independent as $\partial \pi(n,t)/\partial t = 0$, thus justifying its name.

The descaled log-improbability or action function of the stationary distribution is commonly referred to as the NEP and denoted by $\mathcal{V}$ (see Ref. [27]),

$$
\pi(n) \asymp e^{-V\mathcal{V}(q)},
\tag{12}
$$

where the normalization constant has been set to unity following the discussion below Eq. (8). As the action of a stationary distribution, it must satisfy

$$
\begin{aligned}
\frac{\partial \mathcal{V}}{\partial t} &= 0, \\
H\left(\frac{\partial \mathcal{V}}{\partial q}, q\right) &= 0,
\end{aligned}
\tag{13}
$$

where the second line is the Hamilton-Jacobi equation from Eq. (11) that must be satisfied by the NEP.

In the remainder of this section, we will develop and demonstrate the applications of Hamilton-Jacobi theory for CRNs. Finding the NEP $\mathcal{V}$ for a general CRN is of importance for getting any numerical estimates on its behavior, and it is the task that we will concern ourselves with in this paper starting Sec. III. As explained above, to calculate the difference of the action function or NEP between two points we first need to find the optimal path connecting them, which is precisely what the action functional gradient descent (AFGD) algorithm is designed to do.

B. Hamilton-Jacobi theory for stochastic chemical reaction networks (CRN)

In this subsection, we start by introducing the mathematical formulation of CRNs and derive its Hamiltonian function in the concentration coordinate. Next, we show that the Hamiltonian function is convex in the momentum coordinate, a necessary condition for the existence of optimal paths, and write the equations that the optimal paths must satisfy. Then we classify the optimal paths constituting the stationary distribution into two categories, namely relaxation paths and escape paths, and show that the former yield the deterministic mass-action kinetics (MAK) and the latter yields the NEP. Finally we show that, for a general CRN with multiple fixed points, the NEP is always a Lyapunov function with respect to MAK.

A CRN is defined by the triple $\{\mathcal{S}, \mathcal{C}, \mathcal{R}\}$, where $\mathcal{S}, \mathcal{C}, \mathcal{R}$ are the set of species, complexes, and reactions, respectively:

$$
\begin{aligned}
\mathcal{S} &= \{S_1, \ldots, S_i, \ldots, S_{|\mathcal{S}|}\}, \\
\mathcal{C} &= \{y_1, \ldots, y_\alpha, \ldots, y_{|\mathcal{C}|} : y_\alpha \in \mathbb{N}^{|\mathcal{S}|}\}, \\
\mathcal{R} &= \left\{ y_\alpha \xrightarrow{k_{y_\alpha \to y_\beta}} y_\beta : k_{y_\alpha \to y_\beta} \geqslant 0 \right\},
\end{aligned}
$$

where Roman letters (i, j) and Greek letters (α, β) are used to denote species and complex indices, respectively. A complex is a multiset of species, and is denoted by the column vector y_α representing the stoichiometry of the multiset. The state of

a CRN is characterized by a population vector $n \in \mathbb{Z}_{\geqslant 0}^{|\mathcal{S}|}$, and the time evolution of a probability distribution over the state space is given by (for a rigorous microphysical derivation, see Ref. [28])

$$
\begin{aligned}
\frac{\partial \rho(n,t)}{\partial t} &= \sum_{y_\alpha, y_\beta} \frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}} \left(\prod_{i=1}^{|\mathcal{S}|} \frac{(n_i - y_{\beta,i} + y_{\alpha,i})!}{(n_i - y_{\beta,i})!} \rho(n - y_\beta + y_\alpha, t) - \prod_{i=1}^{|\mathcal{S}|} \frac{n_i!}{(n_i - y_{\alpha,i})!} \rho(n,t) \right) \\
&= \sum_{y_\alpha, y_\beta} \frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}} \left(\frac{(n - y_\beta + y_\alpha)!}{(n - y_\beta)!} \rho(n - y_\beta + y_\alpha, t) - \frac{n!}{(n - y_\alpha)!} \rho(n,t) \right),
\end{aligned} \tag{14}
$$

where we get the last equation by making use of the multi-index notation.

Writing Eq. (14) in the form of equation Eq. (3) by making use of the shift operator (see footnote 1), we get

$$
\begin{aligned}
\frac{\partial \rho(n,t)}{\partial t} &= \sum_{y_\alpha, y_\beta} \left(e^{-(y_\beta - y_\alpha)\frac{\partial}{\partial n}} - 1 \right) \frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}} \frac{n!}{(n - y_\alpha)!} \rho(n,t) \\
&\equiv \hat{\underline{H}}_{\mathrm{CRN}} \left(-\frac{\partial}{\partial n}, n \right) \rho(n,t)
\end{aligned}
$$

and obtain the following Hamiltonian operator for CRN

$$
\hat{\underline{H}}_{\mathrm{CRN}} = \sum_{y_\alpha, y_\beta} \left(e^{-(y_\beta - y_\alpha)\frac{\partial}{\partial n}} - 1 \right) \frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}} \frac{n!}{(n - y_\alpha)!}. \tag{15}
$$

Following the derivation in Appendix C 2, we can replace the differential operator $-\partial/\partial n$ with momentum variable p to obtain the Hamiltonian function

$$
\underline{H}_{\mathrm{CRN}}(p,n) = V \sum_{y_\alpha, y_\beta} (e^{(y_\beta - y_\alpha)p} - 1) \frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}} \frac{n!}{(n - y_\alpha)!}.
$$

Finally, using the observation $\frac{(qV)!}{(qV - y)!} \asymp (qV)^y$, we can pass to the Hamiltonian function in the concentration variable $q = n/V$ using Eq. (7) to get

$$
H_{\mathrm{CRN}}(p,q) = \sum_{y_\alpha, y_\beta} (e^{(y_\beta - y_\alpha)p} - 1) k_{y_\alpha \to y_\beta} q^{y_\alpha}. \tag{16}
$$

For the Hessian of the Hamiltonian function in the momentum variable p, we get

$$
\frac{\partial H_{\mathrm{CRN}}}{\partial p^2} = \sum_{y_\alpha, y_\beta} (y_\beta - y_\alpha)(y_\beta - y_\alpha)^T e^{(y_\beta - y_\alpha)p} k_{y_\alpha \to y_\beta} q^{y_\alpha}. \tag{17}
$$

It can be seen that Hessian of the Hamiltonian in p is positive definite as it is a sum of symmetric dyadics, each of which is positive definite. Thus, as commented in the discussion below Eq. (8), the Hamiltonian function for a CRN is indeed convex in the momentum variable.

As noted in the previous subsection, the convexity of the Hamiltonian in the momentum coordinates allows us to find stationary or optimal paths where the variation of the action functional vanishes. The optimal phase-space paths $(q(t), p(t))$ satisfy the Hamilton's equations of motion (EoM)

$$
\begin{aligned}
\frac{dq}{dt} &= \frac{\partial H}{\partial p}, \\
\frac{dp}{dt} &= -\frac{\partial H}{\partial q},
\end{aligned} \tag{18}
$$

where we have dropped the $*$ in denoting the optimal path for simplifying notation. For a derivation, see Appendix C 3, where we also show that the Hamiltonian is constant along the optimal path and the equivalence of Hamilton's EoM with the Hamilton-Jacobi equations. Intuitively, the ODEs of Hamilton's EoM are the characteristic curves of the Hamilton-Jacobi PDE, and thus the union of all solutions to Hamilton's EoM form the complete solution to the Hamilton-Jacobi PDE.

Substituting the CRN Hamiltonian from Eq. (16) into Eq. (18) and omitting the species index for notational clarity, we get

$$
\begin{aligned}
\frac{dq}{dt} &= \sum_{y_\alpha, y_\beta} (y_\beta - y_\alpha)(e^{(y_\beta - y_\alpha)p}) k_{y_\alpha \to y_\beta} q^{y_\alpha}, \\
\frac{dp}{dt} &= -\sum_{y_\alpha, y_\beta} (e^{(y_\beta - y_\alpha)p} - 1) k_{y_\alpha \to y_\beta} \frac{\partial q^{y_\alpha}}{\partial q},
\end{aligned} \tag{19}
$$

which are the set of ODEs satisfied by the optimal paths for a CRN.

Let us begin investigating the optimal paths specified by the above equations. For reasons mentioned in Eq. (13), we only consider the subclass of optimal paths in the $H_{\mathrm{CRN}} = 0$ submanifold. For the remainder of this section, we drop the subscript CRN and refer to the CRN Hamiltonian simply by H.

Notice that for $p = 0$, the Hamiltonian is identically zero, i.e., $H(p = 0, q) = 0$ for all q. Substituting this assignment in Eq. (19), we get

$$
\begin{aligned}
\left. \frac{dp}{dt} \right|_{p=0} &= 0, \\
\left. \frac{dq}{dt} \right|_{p=0} &= \sum_{y_\alpha, y_\beta} (y_\beta - y_\alpha) k_{y_\alpha \to y_\beta} q^{y_\alpha}.
\end{aligned} \tag{20}
$$

The first line ensures that, since the time derivative of p vanishes, $p = 0$ is a consistent assignment everywhere along the optimal path. The second line is nothing but the set of ODEs corresponding to the law of mass-action kinetics for a CRN. Thus, the optimal paths with $p = 0$ correspond to the deterministic trajectories of a CRN, which we term "relaxation trajectories" and denote as $q_{\mathrm{rel}}(t)$.

To understand why the $p = 0$ solution must correspond to the deterministic behavior of CRN in the Hamilton-Jacobi formalism, let us calculate the change in action along the

relaxation trajectory. From Eq. (9),

$$\int_0^T dS[q_{\text{rel}}] = \int_0^T \left[0 \frac{dq_{\text{rel}}}{dt} - H(0, q_{\text{rel}}) \right] dt$$
$$= 0. \tag{21}$$

Recall from Eq. (8) that the exponential of the scaled change in action acted as a conditional probability up to a normalization factor. Since the change in action is identically zero, almost all of the probability measure without any suppression evolves along the relaxation trajectory yielding

$$\rho(q_{\text{rel}}(T), T) = N\rho(q_{\text{rel}}(0), 0).$$

It is quite common for the deterministic dynamics of a CRN, Eq. (20), to exhibit multiple fixed points. Let us denote a fixed point by $\underline{q}$. Then each fixed point must satisfy

$$\frac{d\underline{q}}{dt} = 0 = \sum_{y_\alpha, y_\beta} (y_\beta - y_\alpha) k_{y_\alpha \to y_\beta} \underline{q}^{y_\alpha}$$
$$= \sum_{y_\alpha} y_\alpha \sum_{y_\beta} \left(k_{y_\beta \to y_\alpha} \underline{q}^{y_\beta} - k_{y_\alpha \to y_\beta} \underline{q}^{y_\alpha} \right). \tag{22}$$

The existence of multiple steady states can be ruled out in many cases by certain topological considerations and is studied in deficiency theory. For an introduction and detailed derivation of some results, we refer the reader to Refs. [29–32].

For our purposes in this paper, we restrict our attention to generic CRNs with multiple fixed points. These points can be indexed by the number of repelling directions, which can easily be determined using the eigenvalues of the Jacobian of the flow field in Eq. (20). A fixed point with no repelling or all attracting directions is referred to as a *stable attractor*. In a purely deterministic setting, one would expect any evolving probability distribution to end up concentrated at the multiple stable attractors with ratios determined by the initial conditions. In a stochastic setting however, there can be large fluctuations that take the system out from the vicinity of a stable attractor to a generic point $q \neq \underline{q}$ in its basin of escape. The path along which the system arrives at q with a leading exponential probability will be the optimal path that starts from $\underline{q}$ and ends at q, and the probability of this event is precisely given by the action function along that path using (8). We call the optimal paths out of a fixed point in the $H(p, q) = 0$ submanifold as escape trajectories, denoted by q_{esc}, and from Eqs. (12) and (13) it is precisely the action function along paths that determine the value of the NEP,

$$\mathcal{V}(q) - \mathcal{V}(\underline{q}) = \int_{q_{\text{esc}}} p_{\text{esc}}(q') \frac{dq'}{dt} dt$$
$$= \int_{\underline{q}}^q p_{\text{esc}}(q') \, dq', \tag{23}$$

where $(q_{\text{esc}}, p_{\text{esc}})$ is the optimal path that connects $\underline{q}$ and q. Note that the momentum assignment along the escape path must necessarily be nonzero, except at the fixed points, and the Hamiltonian function must evaluate to zero everywhere along the optimal path:

$$p_{\text{esc}}(q) \equiv p(q_{\text{esc}}) \neq 0,$$
$$H(p_{\text{esc}}, q_{\text{esc}}) = 0. \tag{24}$$

To understand why the momentum assignment along escapes must be nonzero everywhere except the fixed points and to prove that the NEP is a Lyapunov function with respect to the relaxation trajectories (mass-action kinetics), let us investigate the geometry of the $H(p, q) = 0$ submanifold in the $2|\mathcal{S}|$-dimensional phase space. Since we are considering a CRN with fixed points, there must exist shift vectors $y_\beta - y_\alpha$ such that $H \to \infty$ when $|p| \to \infty$ in all directions. Also, recall from Eq. (17) that for a fixed concentration q, the Hamiltonian is a convex function in p. This means the $H(p, q) = 0$ submanifold at a given q must be a closed surface and bounded in the interior with $H < 0$. Next, let us consider the outward normal to the surface given by $\partial H/\partial p$. Recall from Hamilton's equations from Eq. (19) that the outward normal is equal to the rate of change of concentration. Since by the fixed-point condition

$$\frac{\partial H(0, \underline{q})}{\partial p} = 0, \tag{25}$$

this means that the outward normal to the surface vanishes and thus the $H(p, q) = 0$ submanifold must pinch off at the fixed points yielding $p = 0$ as the only solution in the submanifold. Finally, to show that the NEP is a Lyapunov function along the deterministic MAK, let us consider its time derivative along a relaxation trajectory:

$$\frac{d\mathcal{V}(q_{\text{rel}}(t))}{dt} = \frac{\partial \mathcal{V}(q)}{\partial q} \frac{dq}{dt} \bigg|_{\text{rel}}$$
$$= p_{\text{esc}}(q) \frac{dq}{dt} \bigg|_{p=0}$$
$$= (p_{\text{esc}} - 0) \frac{\partial H}{\partial p} \bigg|_{p=0}$$
$$\leqslant 0, \tag{26}$$

where the first line follows from chain rule, the second line uses the Hamilton-Jacobi relation of momentum and NEP from Eq. (23), and the third line follows from Eq. (20). The last line follows from the fact that the gradient of a convex function along a chord in the convex subset (secant) is nonpositive, or equivalently the outward normal always makes more than $90°$ with a secant, shown diagramatically in Fig. 2.

C. Switching dynamics for a multistable stochastic CRN

Let us refer to the set of attractors (fixed points) of a CRN, when taken under MAK, by $\underline{q} \equiv \{\underline{q}_1, \underline{q}_2, \ldots, \underline{q}_N\}$. Now, label the attractors by the number of positive eigenvalues of the Jacobian of the MAK rates or the mixed-Hessian of the Hamiltonian $\partial^2 H(p, q)/\partial p \partial q$ evaluated at $p = 0$. Recall that the positive eigenvalues give the number of repelling directions around the attractor. Let us refer to the fixed points with zero and one repelling directions as *stable* and *saddle* attractors, respectively.

For large scale factor V, the switching dynamics of the stochastic CRN will be generally governed by the stable and saddle attractors. Recall from Sec. II A that the most likely
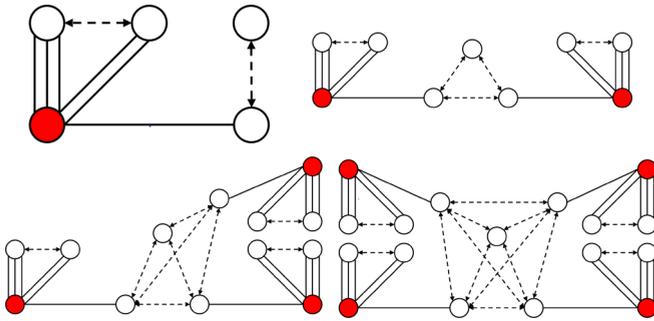
FIG. 1. Diagrammatic representation of 1-Schlogl, 2-Schlogl, 3-Schlögl, and 4-Schlögl model CRNs from top left to bottom right. Species, complexes, stoichiometry, and transitions are represented with solid circles, empty circles, solid lines, and dashed arrows, respectively.

FIG. 2. Diagrammatic proof of the Lyapunov property of the nonequilibrium potential with respect to mass action kinetics.

or optimal paths for a stochastic trajectory lie in the $H = 0$ submanifold. Moreover, in Sec. II B, we showed that there are both trajectories leading into a stable attractor (termed relaxation trajectories) and coming out of a stable attractor (escape paths). In particular then, corresponding to relaxation trajectories that emanate nearby a saddle attractor into a stable attractor, there will also be escape paths emanating from the stable attractor and terminating at the saddle attractor. The data obtained by collecting all the stable and saddle attractors and the relaxation and escape paths joining them is called as the *heteroclinic network* for the CRN. For example, the heteroclinic network obtained for the 2-Schlögl model from Fig. 1 is shown in Fig. 3 (also see Sec. IV B).

Generically, one expects a stochastic CRN to spend a long time near a stable attractor before transitioning to a subsequent one by passing a saddle attractor in between. Recall that the NEP or action along an escape trajectory (see Sec. II A) quantifies the scaled log-improbability of the escape event. Moreover, as we have shown in Sec. II B, the NEP is a Lyapunov function along MAK. This means that the NEP must monotically decrease along the relaxation trajectory, or correspondingly monotonically increase along the escape trajectory. For example, the NEP for the 2-Schlögl model along all the escape heteroclinic orbits is shown in the right panel of Fig. 3. Moreover, in Fig. 4, the NEP is shown for a particular escape and compared against the occupation log-improbability probability obtained by a Gillespie simulation.

We briefly want to remark that since the NEP increases along an escape trajectory, the quantity must not be seen as an *entropy* for a stochastic system. This viewpoint will naturally lead to paradoxical violations of the second law of thermodynamics. However, in Ref. [16], it is shown that for an appropriate definition of entropy for stochastic CRN, its rate of change for any system can be asymptotically estimated by the rate of change of the NEP along a relaxation trajectory, thus yielding a rigorous *second law* for CRN.

III. ACTION FUNCTIONAL GRADIENT DESCENT (AFGD) ALGORITHM

A central property of chemical reaction networks (CRNs) that makes them particularly useful for modeling complex
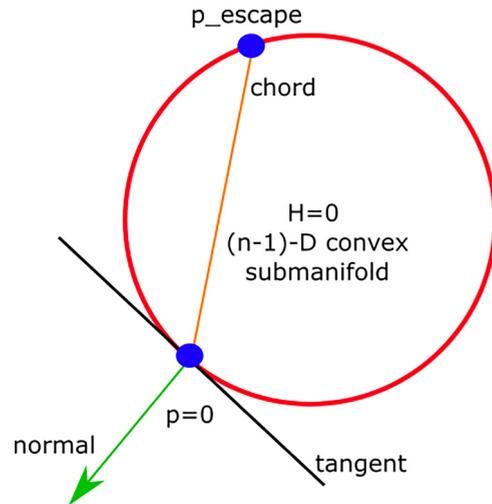
systems is their ability to exhibit multiple attractors when taken under deterministic mass-action kinetics. Moreover, in a stochastic CRN, due to fluctuations in population sizes, it is possible for the system to transition out of a stable attractor into another with some probability. It is of practical value to quantify the transition events, termed as *escapes*, and can help in experiment design as well as network inference [33]. However, the escapes are rare, scaling exponentially in $-V$, compared to fluctuations in population that are polynomial in $1/V$, making stochastic simulation an inefficient approach to estimate the leading consequences of escapes, which dominate basin-switching.

In this section, we provide a deterministic alternative to estimating the escape paths between appropriate start and end points in chemical concentration space using a variational method. Specifically, we present a functional gradient descent algorithm that finds the optimal trajectory, that minimizes the action functional while satisfying the boundary conditions. A mathematical formulation of the optimization problem is posed in Sec. III A and the algorithm is explained in Sec. III B, while relegating the technical details to Appendix D. A discussion of the algorithm's performance costs and its relation to other methods is presented in Sec. III C.

Finally, we will briefly remark on obtaining the probability of system transition from the escape path. First, notice that if the boundary conditions are such that the trajectory emanates from a stable fixed point and terminates in a saddle fixed point, then the resulting escape trajectory does not lie in the solution set of the deterministic mass-action kinetics, and can only be interpreted in the stochastic framework (see Sec. II B). Let the phase-space trajectory of such an escape be denoted by γ^*. For reasons explained below Eq. (8), the likelihood of a system of scale factor V starting from the stable fixed point and escaping to the saddle fixed point can be well-approximated simply by $e^{-V \mathcal{A}[\gamma^*]}$ [see Eq. (12)]. For example, we choose $V = 300$ in Fig. 4, the authors consider values of V as low as 10 in Fig. 4 of Ref. [31], and in both cases the occupation probability from the stochastic simulation data is
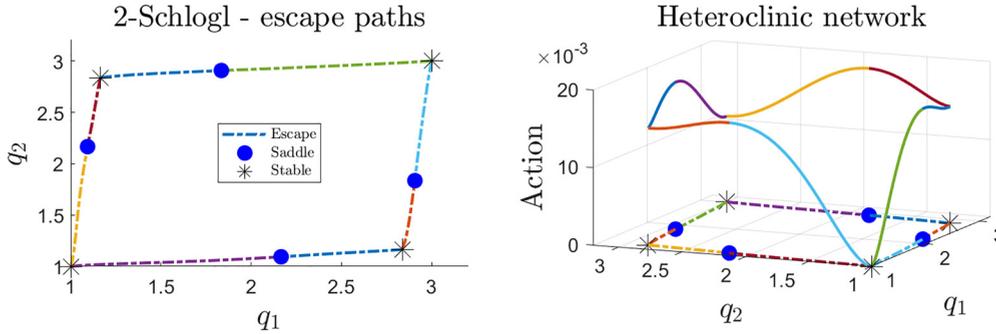
FIG. 3. Escape paths and heteroclinic network for the 2-Schlögl model.

well in agreement with the exponential of the scaled action along the escape trajectory.

A. Formulation as a MinMax problem

We refer to the space of states in which we can find our chemical system as *configuration space*, denoted by Q. Suppose we have $|\mathcal{S}| = D$ species in our network, then the state of a system at a given time is a vector of concentration of each species $\in \mathbb{R}^D_{\geq 0}$. The space of paths that a system can take in configuration space is denoted by $\mathcal{P}_Q$, and we denote a path by $q(s)$. Notice that although the space of parametrized differentiable paths lies in the tangent bundle of Q (see Ref. [34]), the parametrization is arbitrary and we do not yet have a notion of physical time and velocity. To introduce a notion of physical time, we have to introduce a Hamiltonian function [see Eqs. (6) and (16)] defined on the cotangent bundle of Q, commonly referred to as the *phase space*. The coordinates canonically conjugate to configuration coordinates q in the cotangent space are also called as momentum coordinates and denoted by p. Moreover, since the Hamiltonian is convex in momentum, every path in the configuration space has a unique lift in phase space parametrized by time variable t, which we denote by $\gamma \equiv (q(t), p(t))$. To avoid confusion, henceforth we

refer to a path in configuration space in $\mathcal{P}_Q$ as a *curve*, and phase-space path in $\mathcal{P}_{T^*Q}$ as a *trajectory*.

In the notation described in Table I, we can succinctly formulate the problem that we wish to devise an algorithm to solve. We wish to find the optimal phase-space trajectory, i.e., a trajectory that minimizes the action functional, constrained such that its projection to the configuration space curve begins at configuration q_I and ends at q_F, and the phase-space trajectory is along the $H(p, q) = 0$ submanifold [reason for the $H(p, q) = 0$ constraint can be found in discussion around Eq. (13) in Sec. II A]:

$$\text{Find } \gamma^* \equiv (q^*, p^*) \in \mathcal{P}_{T^*Q}$$

$$\text{such that } \mathcal{A}[\gamma^*] = \min_{q(t)} \max_{p(t)} \int_0^T \left(p \frac{dq}{dt} - H(p, q) \right) dt,$$

$$H(p^*(t), q^*(t)) = 0,$$

$$q(0) = q_I, \quad q(T) = q_F. \tag{27}$$

Once the escape path from q_I to q_F is identified, the difference in the nonequilibrium potential between the two configurations is given by value of the action along the
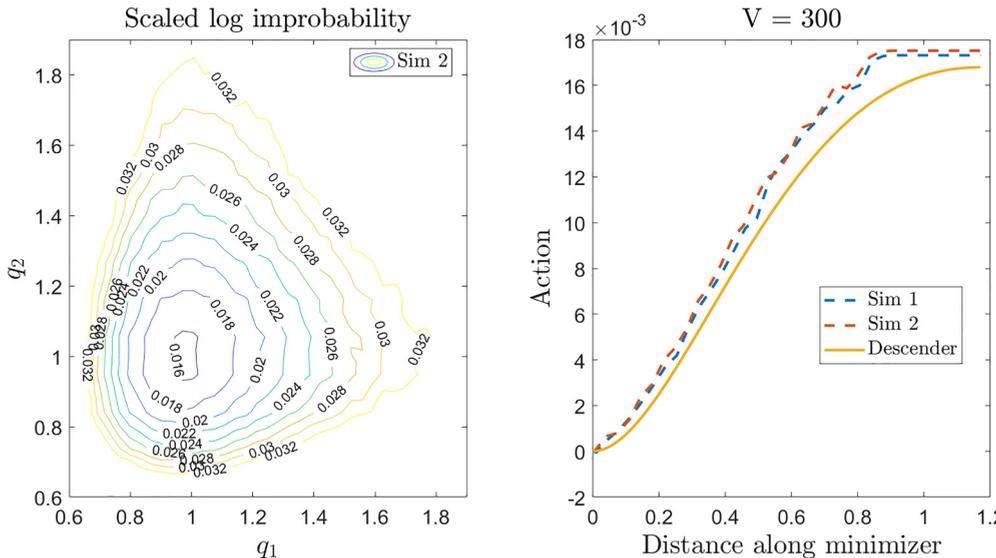

FIG. 4. Comparing the algorithm against Gillespie stochastic simulation for the 2-Schlögl model (Sec. IV B). Scaled log-improbability of the stationary distribution π is defined to be $(-1/V) \log(\pi)$.

TABLE I. Notation used for denoting paths in configuration and phase space.

Object	Notation
Configuration space	$Q \subset \mathbb{R}^D$
Space of curves	$\mathcal{P}_Q = \{q : [0, 1]$
	$\to Q, q(0) = q_I, q(1) = q_F\}$
Curve	$q \in \mathcal{P}_Q$
Phase space	$T^*Q \subset \mathbb{R}^{2D}$
Space of trajectories	$\mathcal{P}_{T^*Q} = \{(q, p) : [0, T]$
	$\to T^*Q, q^*(0) = q_I, q^*(T) = q_F\}$
Trajectory	$(q, p) \equiv \gamma \in \mathcal{P}_{T^*Q}$

path, i.e.,

$$\mathcal{V}(q_F) - \mathcal{V}(q_I) = \mathcal{A}[\gamma^*]. \tag{28}$$

For reasons discussed in Sec. II C, we are only interested in finding escape paths from a stable to a nearby saddle fixed point, thus we will choose the end points to be in the set of fixed points of the mass-action kinetics $\underline{q}$, as defined in Eq. (22). Moreover, from Eq. (25), we know that the $H(p, q) = 0$ phase-space submanifold pinches off and the only solution to $H(p, \underline{q}) = 0$ is $p(\underline{q}) = 0$. Thus, for our purposes here, the momentum end point constraints are $p(q_I) = 0 = p(q_F)$.

B. Algorithm and pseudocode

1. Main idea

For a CRN exhibiting multiple fixed points, we start by picking a stable and a nearby saddle fixed point. The fixed points are determined by finding roots of the polynomial rate equations of mass-action kinetics, and can be found for high-dimensional systems using a numerical algebraic geometry software such as Bertini [35]. Once the end points have been determined, we obtain a configuration space *curve* by joining the end points with a straight line, which serves as an *initial condition* for the algorithm (see Appendix D 2). While in the actual algorithm the initial curve will be uniformly discretized at finitely many points, in this subsection, for the sake of explaining the main idea, consider that the curve is continuous.

Let the continuous configuration space curve at iteration I be denoted by q^I. Then, as explained in Appendix D 3 and shown in Fig. 5, we *lift* the curve to a phase-space *trajectory* by assigning a scalar dt and vector p^I at each point, such that the variation of the action functional in momentum is zero while the phase-space trajectory is constrained to the $H(p, q) = 0$ submanifold,

$$\frac{dq^I}{dt} = \left.\frac{\partial H}{\partial p}\right|_{(q^I, p^I)},$$

$$0 = H(p^I, q^I).$$

Assigning momentum and time coordinates at each point along the discrete curve (see Fig. 5) can be implemented using standard constraint optimization solvers (see Appendix D 3) and is the *costliest step* in the algorithm. Notice that since the Hamiltonian is convex in momentum, the solution exists. However, recall from Eq. (1) that the CRN Hamiltonian, as
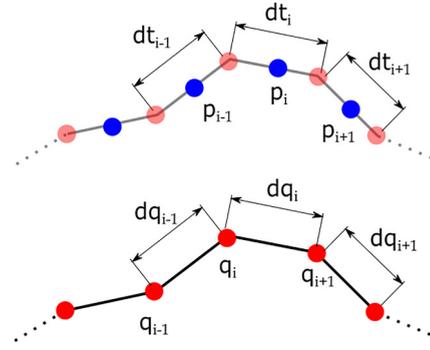

FIG. 5. Lifting a curve in configuration space to a trajectory in phase space by assigning a momentum and time difference to each segment.

opposed to the generic Hamiltonian for mechanical energy, is not quadratic in momentum. Thus, the momentum coordinate generally cannot be solved for analytically and the difficulty of assigning it will depend on both the dimensionality of the system as well as the number of reactions in the network.

From Eq. (C14), the variation of the action functional along a phase-space trajectory (q, p) is

$$\delta\mathcal{A}[q, p] = \int dt \left(\frac{\delta\mathcal{A}}{\delta q}\delta q + \frac{\delta\mathcal{A}}{\delta p}\delta p\right), \quad \text{where}$$

$$\frac{\delta\mathcal{A}}{\delta q} = -\left(\frac{dp}{dt} + \frac{\partial H}{\partial q}\right),$$

$$\frac{\delta\mathcal{A}}{\delta p} = \left(\frac{dq}{dt} - \frac{\partial H}{\partial p}\right). \tag{29}$$

Thus, as explained in Appendix D 4, we update the new configuration curve in the next iteration $q^{I+1} = q^I + \delta q^I$, where a variation δq^I is chosen to be in the steepest descent direction g^I (to first order), where

$$g^I := -\frac{\delta\mathcal{A}}{\delta q} = \left.\left(\frac{dp}{dt} + \frac{\partial H}{\partial q}\right)\right|_{(q^I, p^I)},$$

$$\text{and} \quad \delta q^I = \epsilon\, g^I,$$

with a step-size $\epsilon > 0$ picked by employing a backtracking line search (see Appendix D 6 [36]). This assignment of δq^I ensures that, using Eq. (C14), the variation of the action is negative semidefinite

$$\delta\mathcal{A} = -\epsilon\frac{\delta\mathcal{A}}{\delta q}\frac{\delta\mathcal{A}}{\delta q} \leqslant 0, \tag{30}$$

with equality only at the optimal curve.

The algorithm for finding the solution to Eq. (27) can thus be seen as performing a max by assigning optimal momentum values, followed by moving toward the min by taking a gradient in the descent direction in each iteration. The convergence of the algorithm relies on the convexity of the action functional around the optimal solution, on which we comment in Appendix C 4.

2. Noise and discretization

Since the optimal trajectory is a solution to a variational problem, it must be smooth. A naive implementation along the
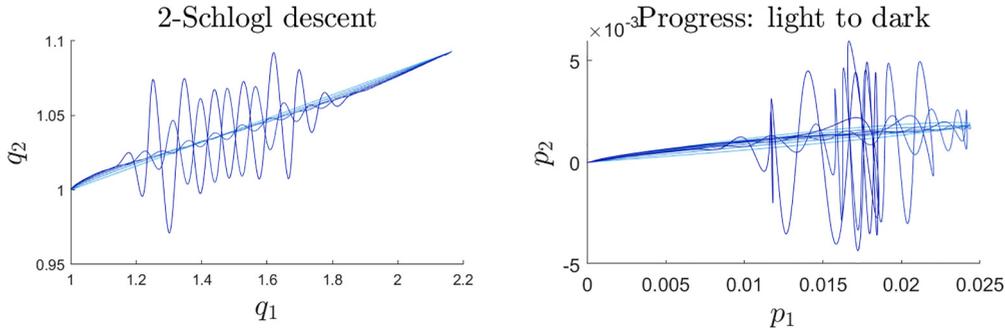
FIG. 6. Functional gradient descent for 2-Schlögl model without filtering. Note that the low-frequency waves first amplify and then concentrate at higher frequencies near a central, badly-performing place on the trajectory.

above account on a finitely sampled or discretized curve will fail on two ends. First, there are two major sources of noise that self-amplify along the iterations, namely due to numerical solving of assigning momentum values (lifting the curve) and discretization. Second, any information in the optimal curve below the length scale of the discretization length will not be able to be captured. We elaborate on this in Appendix D 5, and for an illustration of the results with such a naive implementation, see Fig. 6. Our contribution in this work then is to provide a principled algorithm for controlling these sources of noise and providing a recipe for how, in principle, one could approach the continuum limit in a controlled fashion where the algorithm is guaranteed to converge to the unique solution provided the action functional is convex.

a. Stuck in a local minima. The objective of the algorithm is to converge upon a phase-space trajectory where the value of the action functional along it is minimized. If the difference in the action calculated in two subsequent iterations is below a threshold, or the step size is below a certain threshold, then the descent is considered stuck.

b. Controlling noise. First, to control the noise in the functional gradient g, we employ a low-pass filter and obtain a smoothed gradient g_s, as explained in Appendix D 5. To begin the descent, we choose a small low-pass frequency and allow the algorithm to proceed until it is stuck in a local minimum. The descent gets stuck away from the global minimum whenever the remaining frequencies to be relaxed in the trajectory fall above the pass-band shoulder. Thus, we slightly increase the cut-off frequency of the low-pass filter, allowing more meaningful frequencies to go through and at the same time keeping noise from creeping in, and let the algorithm converge to a new curve (explained in Appendix D 7). Moreover, it must be noted that the algorithm uses a single step size for the complete curve, i.e., $q^{I+1} = q^I + \epsilon g_s$, where ϵ is a scalar. Thus, we find it best to filter the gradient in a space-uniform parametrization of the curve [see Eq. (D16)], which ensures that the magnitude of the gradient affects all the discrete segments in a uniform fashion.

Second, we also low-pass filter the configuration curve q^I in a time-uniform sampling [see Eq. (D17)] each time we update the cutoff frequency. The reason for this is, in practice, although we smooth the gradient, some noise might still accumulate in the descended curve. This step ensures that the descended curve is smooth, and in practice, can often take the curve closer to the optimal curve even in the absence of

a descent step (for instance, see Fig. 19 and the discussion in Appendix D 8).

c. Annealing/Reducing discretization interval. If the algorithm is stuck at a curve where no further descent is possible by increasing the low-pass frequency or filtering the trajectory, then we increase the sampling of the curve by linearly interpolating on a smaller discretization interval. We also refer to this process as *annealing*, the details of which can be found in Appendix D 7 b. One might consider adding $\geqslant 1000$ points in the trajectory, and restarting the algorithm from a new value of the low-pass cutoff frequency at this finely sampled trajectory. A typical reason why the algorithm does not descend further, despite not finding the optimal trajectory, is that the gradient might have meaningful information at length scale smaller than the discretization interval. Increasing the number of sample points then ensures that the meaningful information is not canceled by the low-pass filter. Through examples considered in the next section, we show that this step can take the curve arbitrarily close to the optimal curve, as it theoretically should since we are approaching the continuum limit of the trajectory.

3. Verification protocol

At the optimal trajectory, we know that the functional gradient of the action in both position and momentum must be zero, in other words, the phase-space trajectory must satisfy Hamilton's equations of motion. This gives a canonical verification method for the algorithm, namely numerical integration of Hamilton's equations forward and backward from each point in the phase-space trajectory and looking for how closely they approach the boundary points. The algorithm halts when we have an exact trajectory that starts and ends at the fixed points, and the deviation of the numerical integration from the end points gives us a way to quantify convergence. We explain the process by which one can obtain such a quantity in Appendix D 8. However, it must be noted that symplectic integration in high dimensions is often numerically unstable and we refer readers to Refs. [37,38], for their careful and efficient implementation.

4. Algorithm in a nutshell

Starting with a uniformly discretized straight line between the end points, the action functional gradient descent (AFGD) algorithm consists of three nested subroutines:

Algorithm 1. Pseudocode for Action Functional Gradient Descent (AFGD) algorithm (implementation [13]).

$I \leftarrow 1; \epsilon \leftarrow \epsilon_{IC}; f_c \leftarrow f_{IC}; \Delta^1 \leftarrow \infty$
$q_{IC} \leftarrow$ Initial_condition
$q^1 \leftarrow$ Space_uniform_sampling$[q_{IC}]$
$(q, p, t)^I \leftarrow$ Lift_curve$[q^I]$ ▷ Lift curve to phase-space trajectory
$S^I \leftarrow$ Action$[(q, p)^I]$ ▷ Calculate action
$g^I =$ Functional_gradient$[(q, p, t)^I]$ ▷ Calculate functional gradient
$g_s^I =$ Filter_gradient$[g^1, f_c]$ ▷ Low-pass filter gradient
$\Delta S \leftarrow 0$
while $I \leqslant I_{Max}$ **or** $f_c \leqslant f_{Max}$ **or** $\Delta^I > \Delta_{thresh}$ **do**
  $I \leftarrow I + 1$
  $\epsilon^I \leftarrow$ Step_Size$(q^{I-1}, g_s^{I-1}, \epsilon^{I-1})$ ▷ Pick step size
  $q^I \leftarrow q^{I-1} + \epsilon^I \times g_s^{I-1}$
  $(q, p, t)^I \leftarrow$ Lift_curve$[q^I]$
  $S^I \leftarrow$ Action$[(q, p)^I]$
  $\Delta S \leftarrow S^I - S^{I-1}$
  if $|\Delta S| < \Delta S_{thresh}$ **or** $\epsilon < \epsilon_{thresh}$ **then**
    $f_c \leftarrow f_c + \Delta f; \epsilon^I \leftarrow \epsilon_{IC}$ ▷ Loosen pass-band; refresh step-size
    $q_s^I \leftarrow$ Filter_in_time_uniform$[q^I]$ ▷ Optional: increase sampling/anneal
    $(q, p, t)^I \leftarrow$ Lift_curve$[q_s^I]$
  end if
  $\Delta^I \leftarrow$ Distance_Hamilton's_equations$[(q, p)^I]$ ▷ Verification: Least distance from integrated EoM
  $g^I \leftarrow$ Functional_gradient$[(q, p, t)^I]$ ▷ Calculate functional gradient
  $g_s^I \leftarrow$ Filter_in_space_uniform$[g^I, f_c]$ ▷ Low-pass filter gradient
end while

---

(1) Lift the curve to a trajectory and calculate the functional gradient. Filter the gradient in space-uniform parametrization and calculate a step size to descend. Descend to obtain a new curve.

(2) If the descent is stuck, then filter the curve in time-uniform parametrization and advance the space-uniform filter parameters. Go to routine 1.

(3) If the descent is stuck, then anneal or add more sample points and refresh the filter parameters. Go to routine 1.

Optionally, to quantify a distance from the global minimum, one can also repeatedly employ the verification protocol.

A pseudocode is presented in Algorithm 1 and the algorithm is diagramed in Fig. 7. The mathematical and implementation details of the subroutines can be found in Appendix D and the code can be found in Ref. [13].

C. Performance costs and relation to other methods

From the last subsection, recall that the costliest step of our algorithm is the constrained optimization function needed to lift the curve. Let us refer to the cost of solving the optimization problem as $C_{c.opt.}$. Also, recall that the algorithm consists of three nested subroutines, namely obtaining the gradient and descending, advancing the filter-parameter when stuck, and advancing the number of sample points when stuck in the second routine. Let us refer to the number of points in the curve to be N_P, and the number of times the second and third routines are called as N_2 and N_3, respectively. Moreover, suppose that each time the third routine is performed, ΔN_P points are added to the curve. The cost of running the algorithm then can simply be estimated as

$$\text{Cost of algorithm} \approx \sum_{i=0}^{N_3} (N_P + i\Delta N_P) C_{c.opt.} \times N_2.$$

As commented upon in Appendix D 3, a parallel implementation of calculating the lift of a curve can reduce the cost of computation by a factor proportional to the number of cores.

The proposed AFGD algorithm is a novel functional gradient descent method for solving two-point boundary value problems for a Hamiltonian system. Since it is known that the solutions satisfy the Hamilton equations of motion, another popular method for solving the same type of problem is the *shooting method* [39]. In the shooting method, after selecting the two boundary points, one searches for the trajectory connecting them by integrating a brush of trajectories forward and backward from the starting and end points, respectively, to reach one or more fitting points in between. The technical difficulties associated with the shooting method method are twofold. First, the search cost of the method increases exponentially with system dimension, as the (cross-sectional) brush of trajectories that must be sampled and refined has codimension 1 in the configuration space (starting position) and also codimension 1 in the tangent space (direction). Second, the method as a single algorithm can also fail for CRNs because the Hamiltonian diverges exponentially in the conjugate momentum variables [see Eq. (16)]. Stabilizing numerical integrators with diverging exponentials is challenging [37,40], and if floating-point precision is magnified so that no trajectory can converge to a fixed point, the shooting method must be performed recursively along a path, introducing a further step of error bounding and path updating (for example, see Ref. [32]).
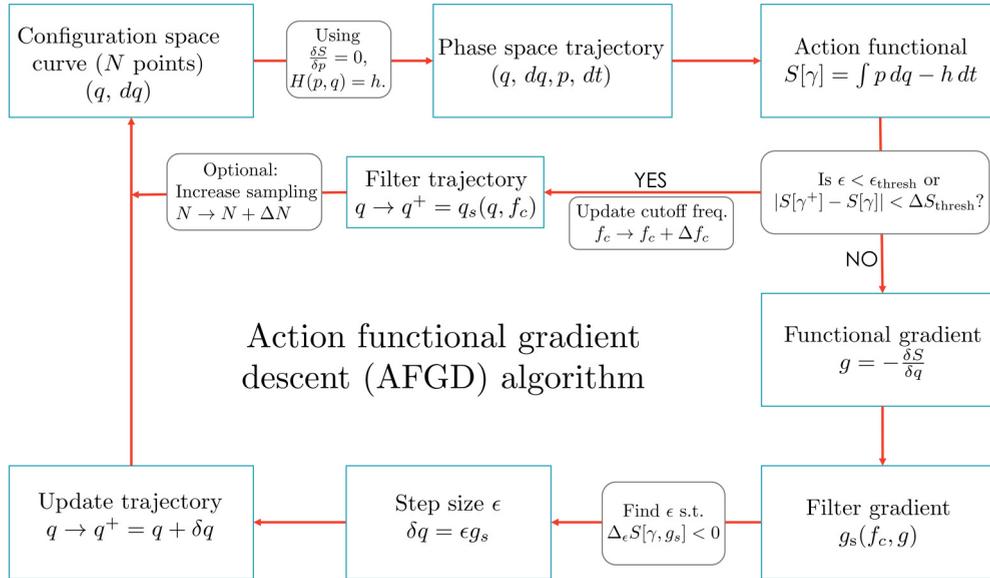
FIG. 7. The AFGD algorithm in a nutshell. For the CRN Hamiltonian, when the end points are fixed points of the system then $h = 0$ must be chosen. For a detailed description, see Sec. III B.

The AFGD algorithm, however, does not suffer from the same issues. First, it relies on computing derivatives, minimizing objective functions, and filtering, all of which are significantly cheaper and more robust than numerical integration. Second, due to the above mentioned reasons, it is also feasible in large dimensions. However, it must be noted that the computation time scales with proportionally to the number of points in the trajectory, and in high dimensions a rather fine sampling of the curve might be needed to get a meaningful escape. The demerit of the AFGD algorithm is that it does not find the exact escape path, as the action can reach its minimum value within tolerance without having completely descended upon the solution to the equations of motion. An example of this can be seen toward the ends of the trajectory near the stable fixed points in Fig. 9, where the trajectory converged upon by the AFGD does not agree completely with that from the shooting method. This suggests that, even in high dimensions, some hybrid of the shooting method and AFGD can be used if the exact escape trajectory is desired. However, if simply a reliable estimate of the action along the escape trajectory is required (to estimate the escape probability, for example), then the AFGD algorithm should be preferred.

IV. APPLICATION AND RESULTS

In this section, we will demonstrate the applicability of the AFGD algorithm using three models with varying features. The first application we consider will be to determine the escape paths and NEP of the Selkov model. The Selkov model, first introduced in 1968 by Sel'kov [41] to model self-oscillations in glycolisis, exhibit relaxation curves that spiral into the stable attractors. Since a spiralling curve, in principle, needs an exponential number of sample points near the fixed point, it poses a difficult challenge for the AFGD algorithm. The other two applications that we consider are higher-dimensional analogs of the Schlögl model. The Schlögl model

was introduced in 1972 to understand nonequilibrium phase transitions in chemical reaction systems [42]. The Schlögl model is an example of a one-dimensional (1D) birth-death process, for which the NEP can be analytically found using Hamilton-Jacobi theory, as we show in Sec. IV B. We then define N-dimensional analogs of the Schlögl model, that we term as N-Schlögl models, for which no analytic results are yet known. We then make use of the AFGD algorithm for the 2-Schlögl and 6-Schlögl model, and verify it against Gillespie simulation and Hamilton's equations of motion, respectively. It must be noted that throughout the examples considered in this section, we tune only the algorithm parameters while keeping the underlying algorithm identical to the pseudocode 1, thus demonstrating that the algorithm is agnostic to the dimensionality of the model. For a MATLAB implementation of the AFGD algorithm, using which we obtain the figures displayed in this paper, see Ref. [13].

A. Selkov model

In this subsection we will consider the well-known Selkov model, which has been analyzed by eikonal methods in Refs. [10,32]. A peculiarity of the model, that makes it both interesting and challenging for analysis, is that both relaxation and escape trajectories exhibit vorticity around the fixed points. In Refs. [10,32], the authors identify the escape trajectories by using the *shooting method* (see Ref. [39]), which amounts to integrating a brush of optimal trajectories in phase space emanating from a point, selecting the one that approaches the desired end point most closely, and repeating. The AFGD algorithm, however, identifies the optimal trajectory using a functional gradient descent, and we summarize the main results of the implementation as well as compare it to the optimal trajectory obtained via the shooting method in Ref. [32] here.
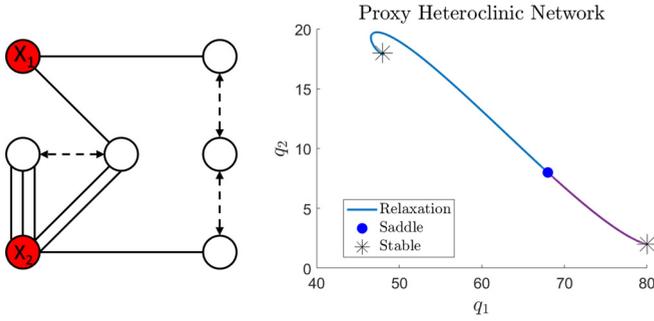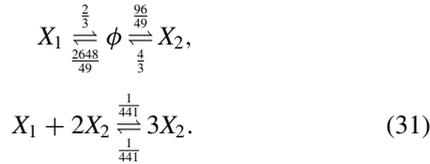
FIG. 8. Diagrammatic representation of the reaction network (left) and a proxy-heteroclinic network using relaxation trajectories for the Selkov model.

To explicitly define the model, our choice of rate constants is identical to Ref. [32]. The concrete model we consider is

$$X_1 \underset{\frac{2648}{49}}{\overset{\frac{2}{3}}{\rightleftharpoons}} \phi \underset{\frac{4}{3}}{\overset{\frac{96}{49}}{\rightleftharpoons}} X_2,$$

$$X_1 + 2X_2 \underset{\frac{1}{441}}{\overset{\frac{1}{441}}{\rightleftharpoons}} 3X_2. \tag{31}$$

A diagrammatic representation of the reaction network can be seen in the left panel of Fig. 8. The resulting fixed points of the model are at $q \in \{(80, 2), (68, 8), (48, 18)\}$, which we first characterize using the number of repelling directions and then find the *relaxation* trajectories that emanate close to a saddle and reach the stable point in the right panel of Fig. 8, which we refer to as as a "proxy-heteroclinic network." Our goal is to find the *escape* trajectories that emanate from each of the two stable points and reach the saddle point, and thus obtain the "true heteroclinic network." For the sake of brevity we will only demonstrate the workings of the algorithm on the top left escape path starting from (48,18) and ending at (68,8); however, the same procedure can be used to obtain the other escape as well.

The first feature that one might notice in the top left relaxation path is that it spirals inwards toward the stable fixed point. This is not uncommon for a CRN and we discuss how to obtain initial conditions for such trajectories in Appendix D 2. We then sample the initial configuration space curve uniformly with 2500 points and run the algorithm, the results of which are summarized in Fig. 20. The top row of the figure displays how the descent progresses across the iterations, and the lower two rows summarize how the action, step size, cutoff frequency and minimum distance from the integrated EoM change along the descent.

Although the action seems to stabilize and the step size is zero roughly beyond iteration 150, the momentum assignment at the converged trajectory (top right panel) clearly looks unreliable. As a rule of thumb, we expect the optimal trajectory to be smooth and containing only low frequency modes. Thus, one can see that despite the AFGD algorithm's effort to smooth the cusp in the momentum initial conditions, it has not yet converged even close to the true solution. To remedy this, we sample the converged curve in 4000 points, a process we call *annealing* (explained in Sec. III and Appendix D 7 b), run the algorithm again starting with this initial condition and
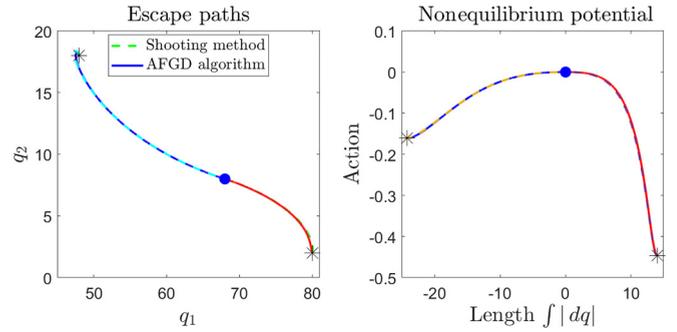

FIG. 9. Comparing escape paths obtained from AFGD algorithm against the shooting method (left panel) and plotting the nonequilibrium potential along the escape (right panel). The data is taken from Ref. [32], and can be compared against Figs. 9 and 15 of the same reference. (For details on the differences between the two methods, see the last paragraph in Sec. III C.)

plot the summary in Fig. 10. This time we indeed find that the converged momentum assignment is without any cusp (right panel of first row), the action has converged to a yet lower value (second row) and the least distance from the integrated equations of motion Δ drops below 0.5 (third and fourth row), guaranteeing convergence near the true solution.

We now proceed by the same method to find the other escape trajectory and plot them against the optimal trajectory found by the shooting method in Ref. [32] in the left panel of Fig. 9. Using Eq. (23), we find the NEP along the escape trajectories and thus obtain the true heteroclinic network for the Selkov model, as displayed in the right panel of Fig. 9.

B. *N*-dimensional Schlögl model

One-dimensional Schlögl model. We begin our discussion by first considering 1D birth-death processes, of which the 1D Schlögl model is an example. A one-species reaction network is called as a birth-death process if the difference of the vectors denoting the target and source complex for each reaction is either positive or negative one, that is

$$|y_\beta - y_\alpha| = 1 \text{ for all reactions in } \mathcal{R},$$

$$\text{where } \mathcal{R} = \{y_\alpha \xrightarrow{k_{y_\alpha \to y_\beta}} y_\beta : k_{y_\alpha \to y_\beta} \geqslant 0\}.$$

The reactions where the stoichiometry of the target complex is one more or less than the source correspond to "birth" or "death" reactions, respectively [27].

Using the form of the CRN Hamiltonian in Eq. (16), we can write the Hamiltonian function of a 1D birth death process as

$$H_{1\text{-b.d.}}(p, q) = (e^p - 1) r_{+1}(q) + (e^{-p} - 1) r_{-1}(q), \tag{32}$$

where $r_{+/-1}$ are polynomials in q with coefficients as the corresponding rate constants appearing in the birth/death reactions. From the Hamiltonian, one can read that the deterministic rate of growth in the concentration of the species is $\dot{q} = r_{+1}(q) - r_{-1}(q)$ and the roots of $\dot{q}$ correspond to fixed points of the system.

We can now proceed to find the NEP for such processes by solving for an escape momentum assignment in the $H_{1\text{-b.d.}} = 0$ submanifold. Since we have only one species and a
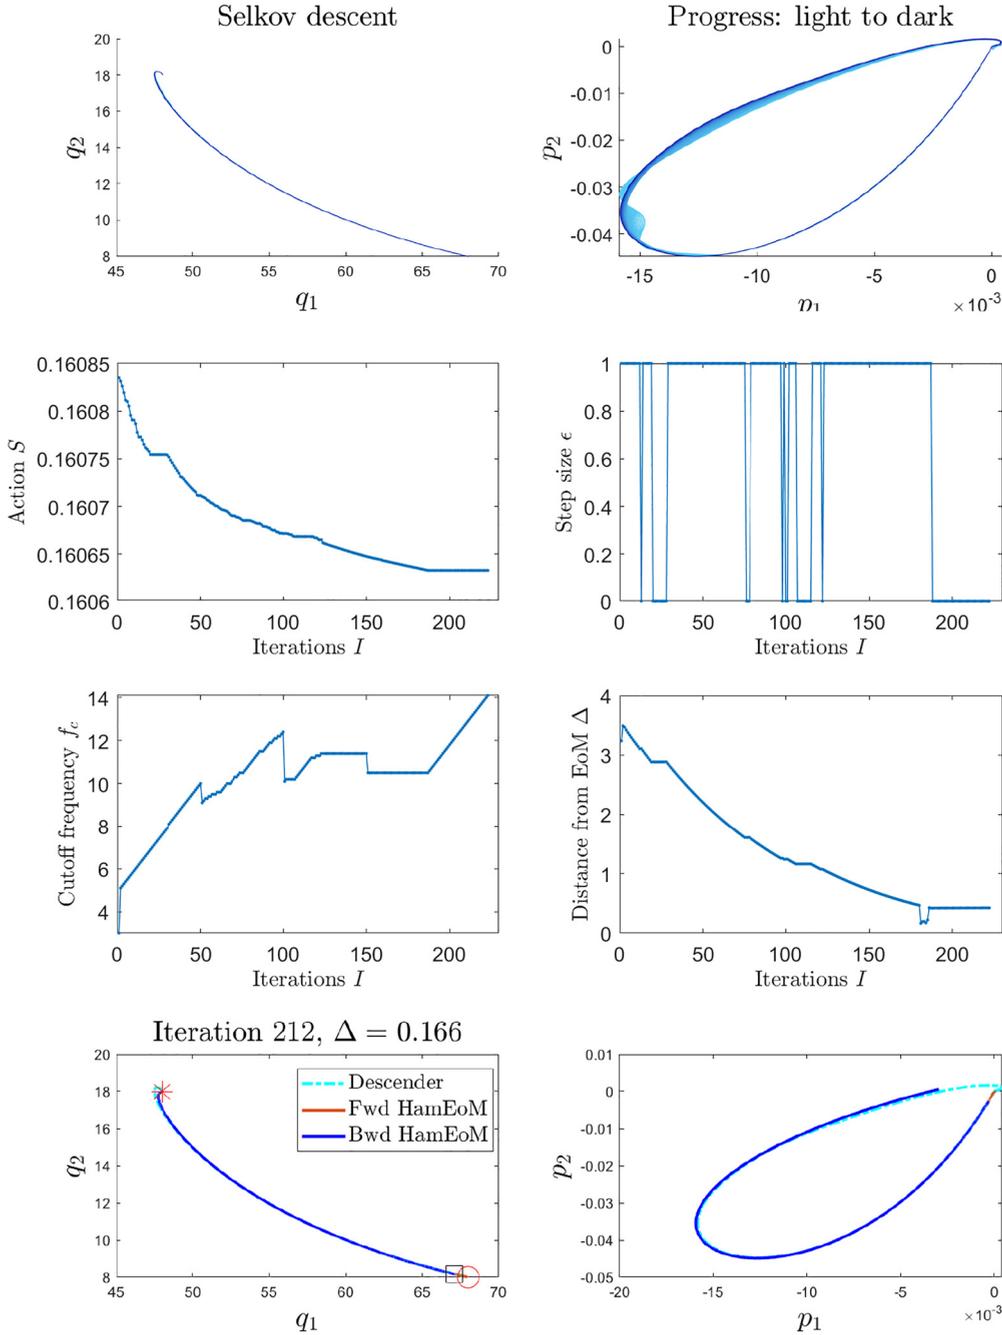
FIG. 10. Descent progress, summary and proof of convergence for the Selkov model after annealing to a trajectory with 4000 points. For initial descent with 2500 points, see Fig. 20.

two-dimensional phase-space manifold, for a birth-death process this constraint uniquely picks a $p_{esc} \neq 0$,

$$H_{\text{1-b.d.}}(p_{esc}, q) = 0,$$

$$p_{esc}(q) = \ln \left[\frac{r_{-1}(q)}{r_{+1}(q)} \right]. \qquad (33)$$

Following the discussion on Hamilton-Jacobi theory in Sec. II B, integrating the escape momentum in Eq. (33) yields the NEP $\mathcal{V}$. Notice that at the fixed points, $p_{esc} = 0$, and correspondingly the NEP is at a local extrema. To ensure that the NEP is always greater than zero, we find a possibly

nonunique fixed point $\underline{q}$, such that

$$\underline{q} = \arg\min_q \int_0^q \ln \left[\frac{r_{-1}(q)}{r_{+1}(q)} \right] dq,$$

using which the NEP is defined as

$$\mathcal{V}_{\text{1-b.d.}}(q) = \int_{\underline{q}}^q \ln \left[\frac{r_{-1}(q)}{r_{+1}(q)} \right] dq. \qquad (34)$$

The ratio of the value of the stationary distribution at points n_2 and n_1 in a stochastic simulation of scaling volume V is then
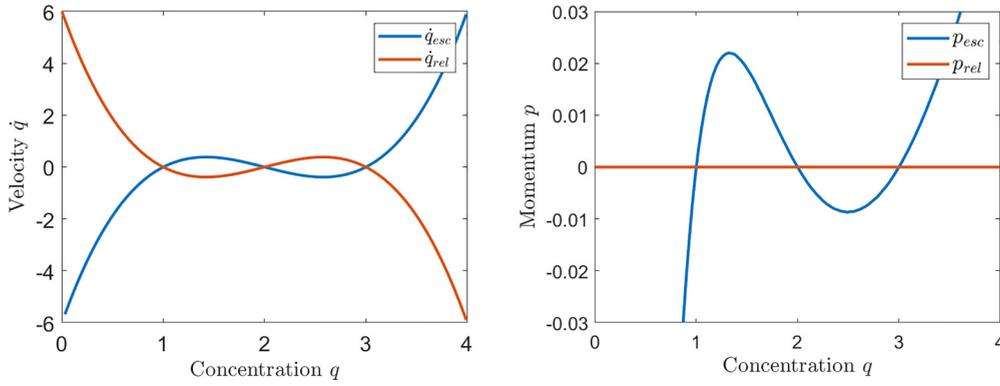
FIG. 11. $H(p, q) = 0$ submanifold in the tangent and cotangent space for the 1D Schlögl model.

given by

$$\frac{\pi(n_2)}{\pi(n_1)} = \exp\left\{-V \int_{q_1}^{q_2} \ln\left[\frac{r_{-1}(q)}{r_{+1}(q)}\right] dq\right\}. \quad (35)$$

The results concerning the NEP of 1-D birth death systems are well known, and an alternative derivation can be found in Ref. [27]. A particular example of a birth-death process relevant to our purposes here is the Schlögl model [42]. For a pedagogical exposition of an application of Hamilton-Jacobi theory to Schlögl model and its generalizations, see Refs. [16,43]. For this section, we use the concrete one-species reaction network

$$\phi \underset{1}{\overset{6}{\rightleftharpoons}} X,$$

$$2X \underset{1}{\overset{6}{\rightleftharpoons}} 3X, \quad (36)$$

where the rate constants are identical to example 10 of Ref. [27]. For a diagrammatic representation of the reaction network, see top left panel in Fig. 1. For the choice of reaction network in Eq. (36), we have

$$r_{-1}(q) = 11q + q^3,$$

$$r_{+1}(q) = 6 + 6q^2,$$

$$\frac{dq_{\rm rel}}{dt} = \frac{\partial H}{\partial p}\bigg|_{p=0} = r_{+1}(q) - r_{-1}(q)$$

$$= -(q-1)(q-2)(q-3). \quad (37)$$

Notice that the relaxation flow field has three roots, which means that the model exhibits three fixed points for the rate constants chosen $q \in \{(1), (2), (3)\}$. We display the velocity and momentum assignment of the variational solution along the $H(p, q) = 0$ submanifold for this model in Fig. 11.

Two-dimensional Schlögl model. We define the two-dimensional (2D) Schlögl model to be a two-species reaction network with 1-Schlögl reactions in each species and diffusion between the two species. For the results in this section and Appendix D, we use the concrete network

$$\phi \underset{11}{\overset{6}{\rightleftharpoons}} X_1 \quad \phi \underset{11}{\overset{6}{\rightleftharpoons}} X_2 \quad X_1 \underset{0.15}{\overset{0.15}{\rightleftharpoons}} X_2,$$

$$2X_1 \underset{1}{\overset{6}{\rightleftharpoons}} 3X_1 \quad 2X_2 \underset{1}{\overset{6}{\rightleftharpoons}} 3X_2, \quad (38)$$
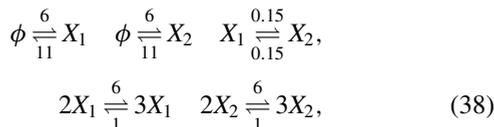
which we represent diagrammatically in the top right panel in Fig. 1. For the choice of rate constants, the relaxation flow field has nine fixed points, out of which four are stable (with no repelling direction), four are saddle (with one repelling direction) and one is unstable (with all repelling directions).

As explained in Sec. II A, to understand the switching dynamics from one stable fixed point to another, we need to only consider the least-improbable path of escape. We know that the least-improbable path between two nearby fixed points is through the saddle point between them, and thus we look for the equation of motion that emanates from a stable attractor and ends at any of the closest saddle points using the AFGD algorithm. Before looking for escape paths however, we find it useful to find the relaxation trajectories that emanate close to a saddle point and reach a nearby stable point. This information yields a "proxy-heteroclinic network" and tells us whether or not the system exhibits any vorticity in the areas of interest. For the position and classification of the fixed points and a "proxy-heteroclinic network" for the 2-Schlögl model, see Fig. 12.

We now proceed to identifying all the relevant escape paths to create a heteroclinic network for the 2-Schlögl model. As an illustrative example of an application of the AFGD algorithm, we first focus our attention to the bottom left escape path from the stable root $q_I = (1, 1)$ to the saddle root $q_F = (2.16, 1.09)$. From the criteria specified in Appendix D 2, we begin the descent from a straight line initial condition as shown in Figs. 16 and 17. We run our algorithm starting from a curve consisting of 500 points, and we summarize the progress of the descender in Fig. 19. The topmost row depicts how the the phase-space trajectory changes in the configuration space (left) and momentum space (right) as the algorithm progresses. The last two rows summarize the progress of the action along the escape, step size, cutoff frequency and least distance from the end points of the integrated EoM during descent.

One can observe that the descent in the action has a step-like behavior which is mirrored in the step size as well as cutoff frequency. What is happening is that the algorithm descends at a given setting of a cutoff frequency as much as it can, by reducing the step size as long as it is above a certain threshold and the value of the action function is strictly decreasing. Once no further step can be taken because the step size has decreased below the minimum threshold
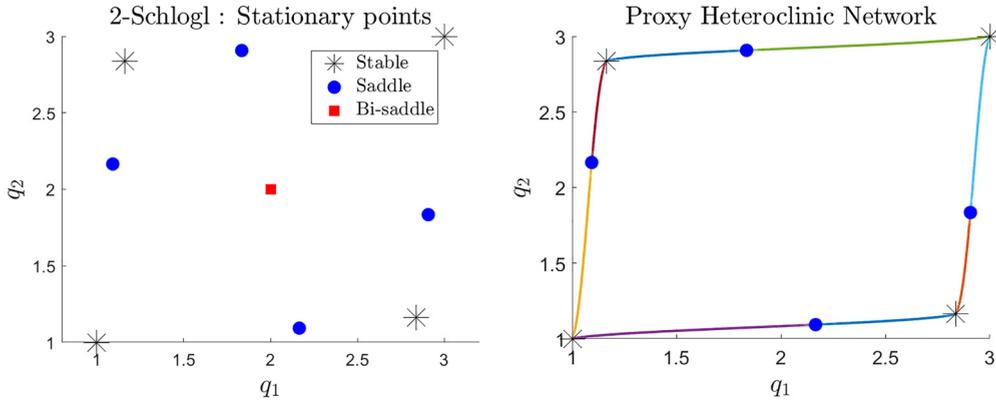
FIG. 12. Fixed points labeled by repelling directions (left) and proxy-heteroclinic network using relaxation trajectories (right) for the 2-Schlögl model.

of ϵ_{thresh}, the cutoff frequency is slightly raised to allow the pass-band to inject more meaningful signal into the gradient. It can be seen in Fig. 13 that the overall magnitude of the gradient decreases and its power spectrum becomes flatter as the descent proceeds through the iterations of the algorithm. For more details, see Appendix D 7.

The bottom right panel in Fig. 19 shows how the minimum distance Δ of the integrated Hamilton's EoM, as explained in Appendix D 8, changes along iterations of the algorithm. Curiously, the distance continues to decrease even when the step size is zero for $I > 55$. The reason for this is that our algorithm also smooths the configuration space curve in a time-uniform sampling, which for the 2-Schlögl demonstrably takes the trajectory toward the optimal solution. The distance, however, ceases to decrease after a point, and it is a signal that the number of points in the trajectory needs to be

increased if further progress is to be made. We show the closest integrated trajectory overlaid on the descender trajectory for a few iterations in Fig. 14. Since, in the previous example, we have already discussed the process of descending further toward the true solution by 'annealing' (increasing the number of sample points in the trajectory), we will not pursue it here.

Validation against Gillespie algorithm. Now that we have verified the correctness of the descended trajectory against Hamilton's EoM, we can ask how well does it perform against a stochastic modeling method such as the Gillespie algorithm [44]? To compare the two, we run two simulations with scaling volume $V = 300$ and find the scaled-log improbability of their stationary distribution π, i.e., we calculate $(-1/V) \log(\pi)$. Since the two simulations give similar results, we only display the contour plot resulting from the second
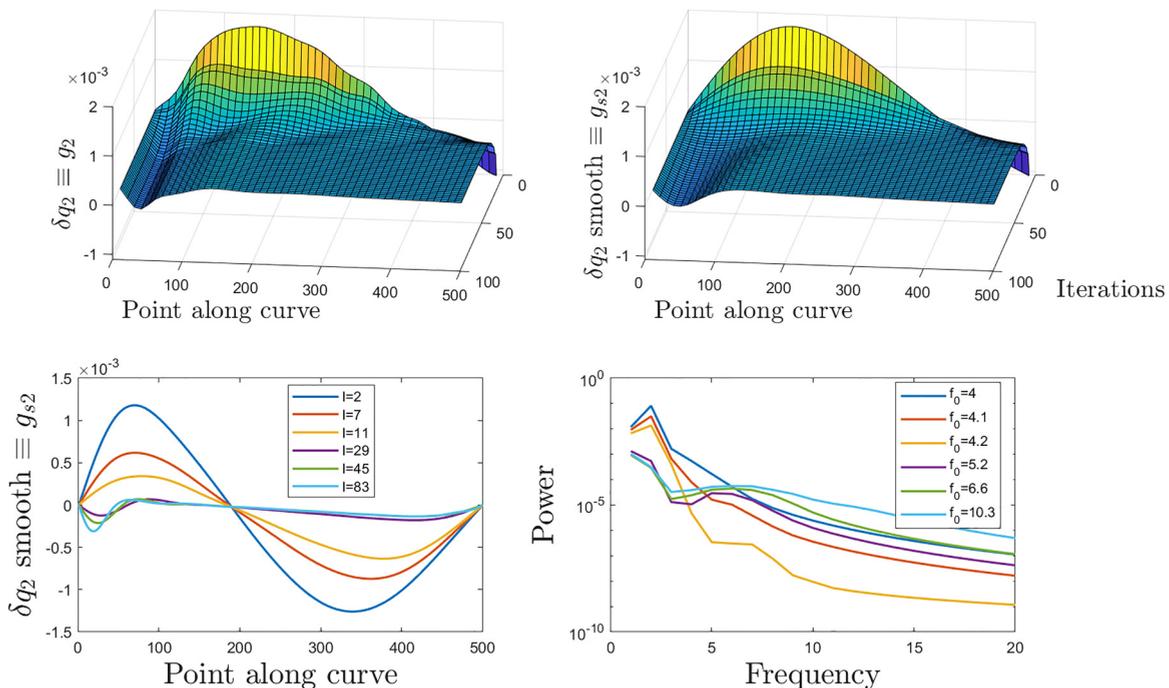

FIG. 13. Extracting functional gradient from phase-space trajectory and smoothing it (top panel). Smoothed gradient and power spectrum at selected iterations (bottom panel) for the 2-Schlögl model.
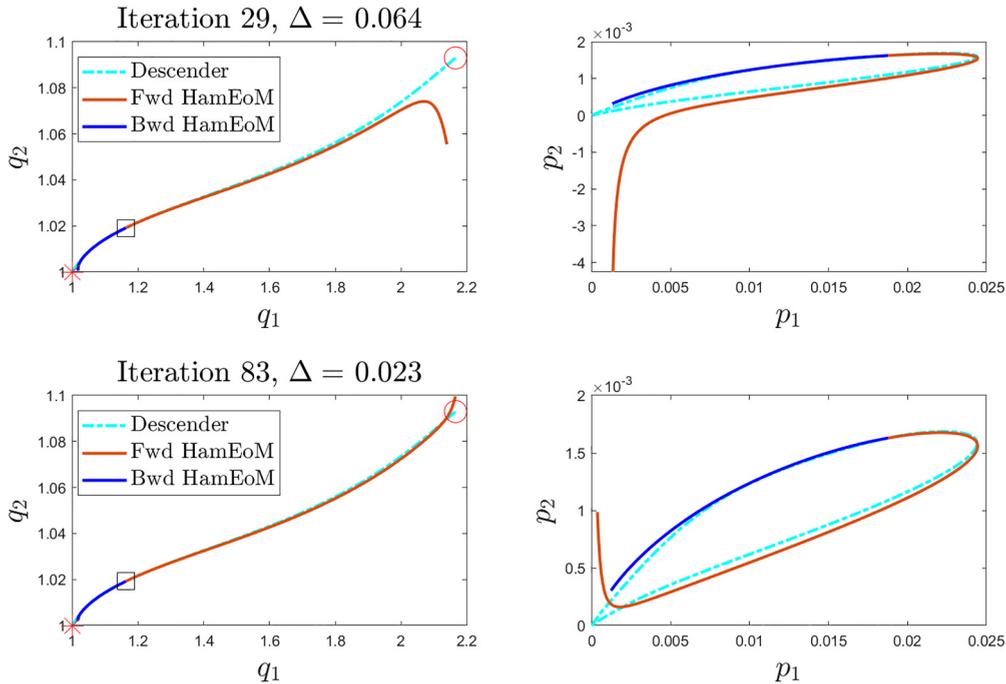
FIG. 14. Progress of integrated Hamilton's equations of motion against the AFGD algorithm at selected iterations for the 2-Schlögl model.

simulation for clarity in the left panel of Fig. 4. We then proceed to find the scaled log-improbability along only the escape trajectory output by the AFGD algorithm and compare it against the action or NEP found using Eq. (23). The result is displayed in the right panel of Fig. 4. We can see that the scaled NEP obtained from the simulation is higher than the log-improbability obtained from the algorithm, as it theoretically should because the variational solution provides a lower bound, becoming exact in the $V \to \infty$ limit. For a similar plot for the 1-Schlögl model, see Fig. 1 in Ref. [27].

Having convinced ourselves of the correctness of the algorithm by two different means, i.e., via integrating Hamilton's EoM and verifying against stochastic simulation, we proceed to find the true heteroclinic network for the 2-Schlögl model. The heteroclinic network consists of all the escape trajectories from every stable attractor, as well as the log improbability or NEP along each escape. We obtain the escape paths by running the algorithm for every pair of nearby stable and saddle points, and display the resulting heteroclinic network in Fig. 3.

N-dimensional Schlögl model. The AFGD algorithm is defined independently of the dimensions of the system, and thus works equally well for dimensions higher than 2. To demonstrate this point, we consider an *N*-dimensional generalization of the Schlögl model. Analogous to the 2-Schlögl model, we define a system with *N* species and include the Schlögl reactions for each species, as well as diffusion among the species.

For $N > 2$ we have to make a choice regarding the underlying diffusion network, and to simplify considerations we choose a fully connected diffusion network, i.e., each species is diffusing with all species. Other choices could also be made, such as the species could form a 1D or a 2D lattice with diffusion only between the nearest neighbors, but this choice

is beside the point of our purposes here. A diagrammatic representation of the 3D and 4D Schlögl model can be seen in the lower line of Fig. 1.

For an application of the AFGD algorithm to higher-dimensional systems, we consider the 6-Schlögl model. To demonstrate the efficacy of the algorithm, the descended phase-space trajectory is plotted against Hamilton's EoM in Fig. 21. The output trajectory is obtained using 2000 sample points and running for under 300 iterations at a relatively low cutoff frequency. Since other details are not particularly illuminating, we omit the summary plots for brevity, but make it available in Ref. [13].

V. DISCUSSION

Chemical reaction networks (CRNs) are essential for modeling a wide range of natural phenomena such as star formation, the origin of life, spatial or ecological patterns in living organisms, and climate [4,7,45,46]. The widespread utility of CRNs stems from their ability to exhibit dynamic equilibria which, unlike the equilibrium at the top or bottom of a potential well where the velocities of objects are zero, are states of a system where the composition remains unchanging although the constituent species are being dynamically exchanged. The state of a star, organism, ecosystem, or climate, when modeled as a stochastic CRN, can undergo transitions from one dynamic equilibrium to another, and it is the probability of such a transition occurring that we give an algorithm to numerically estimate in this work. It must be noted that since the probability of transitions is exponentially suppressed in the number of simulated species, finding cheaper ways of estimating them is of practical importance.

Our main contribution has been to employ the Hamilton-Jacobi formalism to rigorously formulate the problem of

finding transitions between steady states (or fixed points) of a CRN as a MinMax problem, and to construct a principled algorithm to solve it. The functional whose value along its optimal points are the desired transitions is called the *action* functional, due to which we name our algorithm as the *action functional gradient descent* (AFGD) algorithm. Our algorithm only requires computation of derivatives, solving function optimization problems, and basic tools from signal processing. Moreover, while the algorithm itself does not rely on numerical integration, its validity can be readily verified by integrating the equations of motion starting anywhere along the converged output (for more detailed, see Sec. III). Finally, in Sec. IV, we explore applications of the algorithm on several high-dimensional problems and validate them against other methods of obtaining transition paths and probabilities.

While in this work, we only use the algorithm to calculate transitions between two fixed points, in principle it can be used to find a transition between any two points which are guaranteed to have a direct optimal phase-space trajectory connecting them. In particular, rather than escaping from a stable fixed point to a nearby saddle point, one might be interested in finding the probability of an escape to any point within the basin of escape of the stable point in the long time limit. As explained in Sec. II, the optimal trajectories that connect the stable fixed point to the desired point will also be in the $H(p, q) = 0$ submanifold, and can readily be found (by leaving the momentum at the end point unspecified) by the AFGD algorithm. In this way, by connecting arbitrary points to their nearest stable fixed points, in principle, one can assign a transition probability to each point on the state space and recover the occupation probability distribution that one would otherwise obtain by running a stochastic simulation of a CRN for a very long time. We will then leave is as future work to take as input the time-series data or occupation probability distribution obtained from a stochastic simulation and to learn the CRN from which it was generated. Due to the widespread utility of CRN in modeling real-world phenomena, a machine learning algorithm to infer the CRN from its simulated data would have several significant scientific applications.

APPENDIX A: EXPERIMENTAL APPLICATIONS OF AFGD

Consider a dynamical experiment where the individual populations of a list of different types or "species" of objects is tracked through time. Supposing all objects are discrete and objects of the same species are indistinguishable, we can label the different species as $\{S_1, S_2, \ldots\} \equiv \mathcal{S}$ and obtain a population vector $(n_1, n_2, \ldots) \equiv \vec{n}$ which represents the state of the experiment. Furthermore, we will denote the entry in the time series data at time t as $(\vec{n}_t, t)$ or simply $\vec{n}_t$.

In a stochastic experiment, starting from the same initial condition does not guarantee the same observation after some time t has elapsed. In this case, if we are to completely understand the dynamics of the system under investigation, then we have to consider a collection or "ensemble" of experiments starting from some initial condition $(\vec{n}_0, 0)$ and see how the relative frequencies of the different observations in the experiment ensemble changes in time. More precisely, from the data of the experiment ensemble, we get the probability that the experiment is found in $\vec{n}_t$ at time t given that it was in $\vec{n}_0$ at time 0, denoted as $\mathbb{P}[(\vec{n}_t, t) \cap (\vec{n}_0, 0)]$.

Recall from the definition of conditional probability that

$$\mathbb{P}[\vec{n}_t \cap \vec{n}_0] = \mathbb{P}[\vec{n}_t | \vec{n}_0] \mathbb{P}[\vec{n}_0], \tag{A1}$$

where $\mathbb{P}[\vec{n}_t | \vec{n}_0]$ denotes the probability of observing $\vec{n}_t$ conditioned on having observed $\vec{n}_0$. This formula can be useful to us in two ways. First, if the conditional probability $\mathbb{P}[\vec{n}_t | \vec{n}_0]$ is known, given an observation at the time of initialization (detection), we can make a probabilistic prediction (retrodiction) of what will be (was) the state of the system at a future (past) time t. Second, if we have access to both $\mathbb{P}[\vec{n}_t \cap \vec{n}_0]$ and $\mathbb{P}[\vec{n}_0]$, then we can improve our model of the system. In the remainder of the paper we focus on the first of these and comment on the second use in Sec. V.

In Sec. II A [Eq. (8)], we give a detailed exposition of how the Hamilton-Jacobi formalism yields asymptotic estimates to the conditional probability term $\mathbb{P}[\vec{n}_t | \vec{n}_0]$. The basic idea can be summarized as follows. Whenever the dynamics of the ensemble of experiments can be modeled as a Hamiltonian dynamical system satisfying certain conditions, in the limit of a large number of objects in each experiment (making sample fluctuations negligible) the conditional probability $\mathbb{P}[\vec{n}_t | \vec{n}_0]$ can be estimated by finding the optimal path that takes the system from $\vec{n}_0$ to $\vec{n}_t$. The optimality condition can take various interpretations such as most probable, least improbable, least costly, etc. [15], but in all these interpretations there is a variational principle at play which derives from the Hamiltonian structure of the dynamics. Instead of yielding the conditional probability, Hamilton-Jacobi theory gives a descaled conditional log improbability that is referred to as the action function and denoted by

$$S(\vec{q}_t || \vec{q}_0) \equiv -\frac{1}{V} \log \left(\mathbb{P}[V \vec{q}_t | V \vec{q}_0] \right), \tag{A2}$$

where $\vec{q} = \vec{n}/V$. In technical terms, the action function is a solution to the Hamilton-Jacobi PDE, and the optimal paths given by ODEs of Hamilton's equations of motion are its characteristic curves. While historically this formalism has been used to describe motion of celestial or terrestrial objects and is the workhorse of classical mechanics, it can indeed

also be used to quantify population space dynamics such as in chemical reactions, population genetics or economics [7,19,47,48].

A particular type of measurement procedure that is of practical relevance is letting the stochastic experiment evolve under its own rules for an extremely long time, i.e., several orders of magnitude more time than any characteristic time scale in the system. Due to a theorem by Perron and Frobenius [49], the probability distribution of the ensemble of experiments will always approach a stationary distribution independent of the initial conditions, which we denote by $\pi(\vec{n})$. By definition then, $\mathbb{P}[(\vec{n}', \infty) \cap (\vec{n}, 0)] = \pi(\vec{n}')$, and the ratio of the stationary distribution at two population states $\vec{n}$ and $\vec{n}'$, will be given by

$$\frac{\pi(\vec{n}')}{\pi(\vec{n})} = \mathbb{P}[(\vec{n}', \infty)|(\vec{n}, 0)] \equiv \pi(\vec{n}'|\vec{n}). \tag{A3}$$

Thus, for the stationary distribution, the conditional probability between two events $\pi(\vec{n}'|\vec{n})$ is simply equal to the ratio of their stationary probabilities $\pi(\vec{n}')/\pi(\vec{n})$. Using Hamilton-Jacobi theory, we will show in Sec. II A that $\pi(\vec{n}'|\vec{n})$ is asymptotically estimated by finding the optimal path connecting $\vec{n}$ and $\vec{n}'$ in the $H(p, q) = 0$ submanifold of the phase space. Thus, the complete solution to the Hamilton-Jacobi PDE in the $H(p, q) = 0$ manifold, also referred to in literature as nonequilibrium potential (NEP), plays a central role for a certain set of questions. We will return to this point and explain its relevance to our algorithm after we introduce the basic idea behind stochastic chemical reaction network theory.

To explain the type of systems that can be modeled using stochastic chemical reaction networks, we first need an understanding of a "reaction." In a system described by the population of its individual species, a reaction is a rule for a transition that replaces a collection or multiset of objects with a different one. For example, suppose in a system consisting of only two species, a reaction

$$S_1 \rightarrow S_2$$
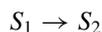
will transition the population vector (n_1, n_2) to $(n_1 - 1, n_2 + 2)$. A "reaction network" is a set of such reactions, each of which are assigned a rate constant determining a proportionality factor in the propensity or how often these reactions are to occur at random. In the particular case of a stochastic 'chemical' reaction network (CRN), the reaction rates are also proportional to the concentrations of the species and model proportional sampling without replacement. Although the dynamics of chemical reaction networks can be seen rigorously as arising from an underlying physical process [28], we will concern ourselves only with the modeling aspects of a stochastic CRN.

Chemical reaction network theory has a long history, and for an excellent review we point the readers to Ref. [8]. The most notable feature of chemical reaction systems is that they can exhibit a wide array of dynamics, with single or multiple attractors, limit cycles, etc. Although there are a host of applications that CRN have, to continue along our practical problem let us pick a particular one. Supposing that our system under investigation can be modeled as a multiattractor CRN, we want to design an experiment to physically observe the different attractors. This amounts to finding how

many experiments our ensemble should consist of to detect a significant amount in each attractor, which, as we will explain, is a type of problem that our algorithm can help get numerical estimates to.

Let us refer to the set of stable attractors in the population space by $\underline{\vec{n}} \equiv \{\underline{\vec{n}}_1, \underline{\vec{n}}_2, \ldots\}$. To estimate the number of experiments our ensemble must consist of, it is sufficient to estimate the ratio of the stationary distribution $\pi(\underline{\vec{n}}_k|\underline{\vec{n}}_j)$ at any two stable attractors $\underline{\vec{n}}_j$ and $\underline{\vec{n}}_k$. From CRN theory and other considerations mentioned in the main text, any optimal path in the $H(p, q) = 0$ submanifold that emanates from the stable attractor $\underline{\vec{n}}_j$ must go through an adjoining saddle point, before going on to the next stable attractor, and so on in the process of reaching some $\underline{\vec{n}}_k$. The optimal paths that take the system out of a stable attractor to an adjoining saddle attractor are called "escape paths," and their collection consists of what we call a "heteroclinic network." The heteroclinic network then determines the asymptotic stationary distribution of the stochastic system at the stable attractors, and the escape paths that constitute the heteroclinic network is precisely what our algorithm is designed to find (see, e.g., Fig. 3).

We conclude this section by explaining how one can also use the algorithm to estimate the ratio $\pi(\vec{n}|\underline{\vec{n}}_1)$ for a general multiattractor CRN at any point $\vec{n}$. First, consider the case where the system exhibits a unique stable attractor $\underline{\vec{n}}_1$. In that case, the ratio $\pi(\vec{n}|\underline{\vec{n}}_1)$ is found simply by determining the optimal path that connects the stable attractor $\underline{\vec{n}}_1$ to the point $\vec{n}$, which is how we recover the Horn-Jackson potential using Hamilton-Jacobi theory in Appendix B. Next, in the case where the system has multiple attractors, we can use the algorithm to determine the stable attractor $\underline{\vec{n}}_k$ from which there is a direct escape path (that does not go through any other stable attractor) to $\vec{n}$. Once the optimal path joining $\underline{\vec{n}}_k$ and $\vec{n}$, or equivalently $\pi(\vec{n}|\underline{\vec{n}}_k)$ is found, we can use the heteroclinic network (as discussed in the preceding paragraph) and

$$\pi(\vec{n}|\underline{\vec{n}}_1) = \pi(\vec{n}|\underline{\vec{n}}_k)\pi(\underline{\vec{n}}_k|\underline{\vec{n}}_1)$$

to determine the desired ratio. In technical terms, since the stationary distribution is asymptotically estimated by the solution of the Hamilton-Jacobi PDE in the $H(p, q) = 0$ submanifold, the conditional probability between any two points $\vec{n}_1$ and $\vec{n}_2$ is obtained by finding the difference in the values of the solution while traversing only along any of the optimal paths or characteristic curves. The geometry of the optimal paths in the $H(p, q) = 0$ submanifold is such that each stable attractor has a region within which each point is connected to the attractor by an escape path (escape basin [32]), and neighboring regions are connected by a saddle point which lies on the boundaries of two adjacent basins (for a detailed exposition, see Ref. [16]). Thus, one can can use our algorithm to find any number of desired characteristic curves of the Hamilton-Jacobi PDE in the $H(p, q) = 0$ submanifold and construct the complete stationary distribution for a chemical reaction network with multiple attractors.

APPENDIX B: REDERIVING PREVIOUS RESULTS IN THE HAMILTON-JACOBI FORMALISM

In this Appendix we illustrate the usage of the formalism developed up to this point by rederiving some well known

results in CRN theory. We start by giving a diagrammatic representation of CRNs, through the help of which we define the complex-balanced condition. Then we find the stationary distribution for such CRN by finding zero-eigenvectors of the corresponding Hamiltonian operators, which recovers the Anderson-Craciun-Kurtz (ACK) theorem [50]. Next, we find the NEP for such a distribution using Hamilton-Jacobi theory which recovers the Horn-Jackson potential and finally show that the negative descaled logarithm of the stationary distribution is indeed the NEP by comparison. We return to the techniques in this section by defining and finding the analytic NEP of another class of models, namely 1D birth-death models in Sec. IV B. There we also consider its generalization to N-dimensional (N-D) birth-death models, and use our algorithm to numerically estimate the corresponding NEPs.

Recall from Sec. II B that a CRN is defined by the triple consisting of the set of species $\mathcal{S}$, set of complexes $\mathcal{C}$, and set of reactions $\mathcal{R}$. For a diagrammatic representation, we represent each species with a solid circle and each complex with an empty circle. Since a complex is a multiset of species, we connect each complex to its constituting species with solid lines, where the number of lines denote the stoichiometry of the complex (denoted by column vector y). Finally, a reaction is a directed edge with a pair of complexes as source (y_α) and target (y_β), which we denote through a dashed line with an arrow indicating direction of reaction. For examples of such a representation, see Figs. 1 and 8 or diagrams in Refs. [32,51].

In this diagrammatic representation, the deterministic flow of mass-action-kinetics can be visualized by assigning weights to the directed edges (dashed lines) equal to $k_{y_\alpha \to y_\beta} q^{y_\alpha}$. An equilibrium is said to be *complex balanced* [52] at $\underline{q}$ if the total flow directed out of each complex equals the total flow inwards or the net flow out of each complex is zero, i.e.,

$$\sum_{y_\beta} \left(k_{y_\beta \to y_\alpha} (\underline{q})^{y_\beta} - k_{y_\alpha \to y_\beta} (\underline{q})^{y_\alpha} \right) = 0 \quad \text{for all } \alpha. \quad \text{(B1)}$$

Given a CRN that exhibits a complex-balanced steady state, we will now find the stationary distribution $\pi(n)$, such that $\hat{H}(-\partial/\partial n, n)\pi(n) = 0$. To identify the stationary distribution, we first need to make the same change of coordinates as we did for deriving the path integral formula in Appendix C 2 Eq. (C9), in which the Hamiltonian operator takes the form

$$\hat{\underline{H}}_{\text{CRN}} \left(z, \frac{\partial}{\partial z} \right)$$

$$= \sum_{y_\alpha, y_\beta} (z^{y_\beta} - z^{y_\alpha}) \frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}} \left(\frac{\partial}{\partial z} \right)^{y_\alpha}$$

$$= \sum_{y_\alpha} z^{y_\alpha} \sum_{y_\beta} \left[\frac{k_{y_\beta \to y_\alpha}}{V^{y_\beta - 1}} \left(\frac{\partial}{\partial z} \right)^{y_\beta} - \frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}} \left(\frac{\partial}{\partial z} \right)^{y_\alpha} \right], \quad \text{(B2)}$$

where in the last line we have switched indices to rewrite the Hamiltonian in a form similar to Eq. (B1), also termed the *complex representation* in Ref. [31]. As mentioned in Appendix C 2, the change of coordinates corresponds to evolving

the $\mathcal{Z}$ transform or the moment-generating function (MGF) rather than the distribution itself. Observe that the coordinate change preserves the commutation relations before and after the transformation,[3]

$$\left[n, -\frac{\partial}{\partial n} \right] = \mathbb{I} = \left[\frac{\partial}{\partial z}, z \right]. \quad \text{(B3)}$$

We will refer to z as the raising operator and $\partial/\partial z$ as the lowering operator. For a representation of the above in abstract linear algebra making use of Fock space operators $a, a^\dagger$ with $[a, a^\dagger] = 1$, see Chapter 4 of Ref. [6].

Next, we find the eigenvectors of the lowering operator $\partial/\partial z$. Let us denote the eigenvector with eigenvalue Vc by $C(z)$. Then we have

$$\frac{\partial C(z)}{\partial z} = VcC(z),$$

$$C(z) = Ne^{cz},$$

where N is a normalization factor (determined in the next line). Recall that the $\mathcal{Z}$ transform or MGF of a probability distribution must be such that $C(z = 1) = 1$, which means that $N = e^{-c}$, yielding

$$C(z) = e^{-Vc + Vcz} = e^{Vc(z-1)}. \quad \text{(B4)}$$

Let us denote the eigenvector of the lowering operator with eigenvalue $V\underline{q}$ by $\underline{Q}(z)$. Then using Eq. (B2),

$$\hat{\underline{H}}_{\text{CRN}} \underline{Q}(z)$$

$$= \sum_{y_\alpha} z^{y_\alpha} \sum_{y_\beta} \left[\frac{k_{y_\beta \to y_\alpha}}{V^{y_\beta - 1}} \left(\frac{\partial}{\partial z} \right)^{y_\beta} - \frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}} \left(\frac{\partial}{\partial z} \right)^{y_\alpha} \right] \underline{Q}(z)$$

$$= V \sum_{y_\alpha} z^{y_\alpha} \underline{Q}(z) \sum_{y_\beta} \left(k_{y_\beta \to y_\alpha} (\underline{q})^{y_\beta} - k_{y_\alpha \to y_\beta} (\underline{q})^{y_\alpha} \right)$$

$$= 0, \quad \text{(B5)}$$

where in the second line we make use of the fact that $\underline{Q}$ is an eigenvector of $\partial/\partial z$ and in the last line we use the complex-balanced condition from Eq. (B1).

Finally, to obtain the distribution $\pi(n)$ from $\underline{Q}(z)$, we can either take the inverse-$\mathcal{Z}$ transform or simply expand the exponential as a summation,

$$\underline{Q}(z) = \sum_{n=0}^{\infty} z^n e^{-V\underline{q}} \frac{(V\underline{q})^n}{n!}$$

$$= \sum_{n=0}^{\infty} z^n \pi(n).$$

---

[3]This transformation is analogous to the Dirac transformation for simple harmonic oscillator in quantum mechanics, where we change from position and momentum to raising and lowering operators. z and $\partial/\partial z$ precisely play the role of the raising and lowering operators, respectively, however, since stochastic dynamics preserves the ℓ_1 norm, the normalization is different. For the relevance of these operators in statistics, see Ref. [19].

By reading the coefficients of the series, we obtain for the stationary distribution

$$\pi(n) = e^{-V\underline{q}}\frac{(V\underline{q})^n}{n!}, \qquad (B6)$$

which is the ACK or multi-Poisson distribution as derived in Refs. [31,50].

Having determined the stationary distribution using the Hamiltonian operator, we proceed to making use of Hamilton-Jacobi theory to find the NEP for complex-balanced system. Recall from Sec. II B that to find the NEP, we need to find a momentum assignment $p_{esc}(q) \neq 0$ at every configuration q, such that $H[p_{esc}(q), q] = 0$ and the Hamilton's equations are satisfied for any initial condition. For finding the momentum assignment, namely escape momentum, along the $H(p, q) = 0$ submanifold it is easiest to recast the CRN Hamiltonian from Eq. (16) in the following form:

$$H(p, q)$$
$$= \sum_{y_\alpha}(e^p)^{y_\alpha}\left[\sum_{y_\beta}\left(k_{y_\beta \to y_\alpha}(e^{-p}q)^{y_\beta} - k_{y_\alpha \to y_\beta}(e^{-p}q)^{y_\alpha}\right)\right].$$

We will now show that $p_{esc} = \ln(q/\underline{q})$ is such an assignment. Observe that

$$H(q, p_{esc})$$
$$= \sum_{y_\alpha}\left(\frac{q}{\underline{q}}\right)^{y_\alpha}\left[\sum_{y_\beta}\left(k_{y_\beta \to y_\alpha}(\underline{q})^{y_\beta} - k_{y_\alpha \to y_\beta}(\underline{q})^{y_\alpha}\right)\right]$$
$$= 0,$$

where the last line follows from the complex-balanced condition. Next, we calculate the total time derivative of p_{esc} and confirm the consistency of such a momentum assignment against the equations of motion,

$$\frac{dp_{esc}}{dt} = \frac{1}{q}\frac{dq}{dt} = \frac{1}{q}\frac{\partial H}{\partial p} = -\frac{\partial H}{\partial q},$$

where the last equality can be verified by a simple calculation.

Thus, the NEP $\mathcal{V}$ for CRN that exhibit a complex-balanced steady state is given by

$$\mathcal{V}(q) = \int_{\underline{q}}^{q} p_{esc}\, dq \qquad (B7)$$
$$= q\ln\left(\frac{q}{\underline{q}}\right) - (q - \underline{q}), \qquad (B8)$$

where an integration constant has been chosen such that $\mathcal{V}(\underline{q}) = 0$, and we recover the Horn-Jackson potential that appears in Ref. [52].

Finally, we show the equivalence of the stationary distribution in Eq. (B6) and NEP in Eq. (B8). Using Eq. (B6) and substituting $n = Vq$, we have

$$-\frac{1}{V}\log(\pi(n)) = -\frac{1}{V}(-V\underline{q} + Vq\log(V\underline{q}) - \log n!)$$
$$\asymp q\log\left(\frac{q}{\underline{q}}\right) - (q - \underline{q}),$$

where the last line is obtained by using Stirling's approximation. We recognize the last equation to be the same as the NEP in Eq. (B8), which completes our expository example of the equivalence of the Hamiltonian operator techniques and Hamilton-Jacobi theory for complex-balanced systems.

APPENDIX C: NONEQUILIBRIUM ACTION, ACTION FUNCTIONAL, AND ITS FIRST AND SECOND VARIATIONAL DERIVATIVES

1. Deriving stochastic and quantum dynamics from the nonequilibrium action (NEA)

Hamiltonian dynamical systems is a mathematical framework for formulating and analyzing dynamics of physical systems defined through a variational principle. The space of states (state space) of a physical system is the space of all possible values one can observe upon measurements at the finest resolution. In physics, the measurable can be position, spin, energy of a particle or a group of particles, while in chemistry or biology, the measurable can be the count, concentration, etc., of a given species of molecule, organelle or organism. An experiment consists of an ensemble of systems, each obeying the same set of rules. These systems can either evolve simultaneously (like particles in a fluid, molecules in a solution, organisms in an ecosystem) or can be spatiotemporally separated (like a series of independent quantum or biological experiments) or both. In all cases, one can assign a distribution over the state space and refer to it as the state of the experiment, which we denote by $|\Psi^R(t)\rangle$, as well as define a sampling protocol that takes a distribution as input and give a scalar quantity, which we denote by $\langle\Psi^L(t)|$.[4]

Following Eyink's construction in Ref. [17], given a Hamiltonian operator $\hat{H}$ we define the nonequilibrium action (NEA) functional Γ to be

$$\Gamma[\Psi^L(t), \Psi^R(t)] = \int \langle\Psi^L(t), (\partial_t - \hat{H})\Psi^R(t)\rangle\, dt,$$

which takes a time-evolving state $|\Psi^R\rangle$ and sampling protocol $\langle\Psi^L|$ as input and yields a scalar quantity. If we expand the NEA around Hilbert-space vectors ϕ^R and ϕ^L,

$$\Psi^R = \phi^R + \delta\Psi^R,$$
$$\Psi^L = \phi^L + \delta\Psi^L,$$

then we get the variation in the NEA to be
$$\Gamma[\Psi^L(t), \Psi^R(t)]$$
$$= \Gamma[\phi^L(t), \phi^R(t)] + \int \langle\delta\Psi^L(t), (\partial_t - \hat{H})\phi^R(t)\rangle\, dt$$
$$+ \int \langle(-\partial_t - \hat{H}^\dagger)\phi^L(t), \delta\Psi^R(t)\rangle\, dt + \mathcal{O}(\delta^2).$$

The stationary condition then yields the following Hamilton's equations of motion on the Hilbert space vectors

$$\partial_t\phi^R(t) = \hat{H}\phi^R(t),$$
$$\partial_t\phi^L(t) = -\hat{H}^\dagger\phi^L(t). \qquad (C1)$$

---

[4] Probability distribution and quantum amplitudes are both distributions, formally defined as vectors in L_1 and L_2 normed Hilbert space, respectively [19].

TABLE II. Correspondence between stochastic and quantum dynamics.

	Stochastic dynamics	Quantum dynamics
Space of states of system	Discrete	Continuous
$q \in Q$	$n \in \mathbb{Z}_{\geqslant 0}^D$	$x \in \mathbb{R}^D$
Space of states of ensemble	Probability distribution	Wave-function
$\mathcal{F} \in L^p$	$\rho^R(n) \in \ell^1$	$\psi(x) \in L^2$
Dual state space	Sampling protocol	State of detector
$\mathcal{F} \in L^{(1-\frac{1}{p})^{-1}}$	$\rho^L(n) \in \ell^\infty$	$\psi(x) \in L^2$
Dual pairing or	$\langle \rho_1^L, \rho_2^R \rangle$	$\langle \psi_1, \psi_2 \rangle$
Inner product	$\sum_n \rho_1^L(n)\rho_2^R(n)$	$\int \psi_1^*(x)\psi_2(x)\,dx$
Hamiltonian operator	ℓ^1 preserving dynamics	L^2 preserving dynamics
	Infinitesimal stochastic	Infinitesimal unitary
	$\sum_n \hat{H}_{nn'} = 0 \quad$ for all n'	$\hat{H} = -\hat{H}^\dagger$
		$\hat{H} = -\frac{i}{\hbar}\hat{M}, \hat{M}^\dagger = \hat{M}$
		($\hat{M}$ is Hermitian)
Forward equation	Master equation	Schrödinger equation
	$\frac{\partial \rho^R(n)}{\partial t} = \hat{H}\rho^R(n)$	$\frac{\partial \psi}{\partial t} = -\frac{i}{\hbar}\hat{M}$
	$\frac{\partial \rho^R(q)}{\partial t} = V\hat{H}\rho^R(q)$	
Adjoint representation	$\mathcal{Z}$ transform	Fourier transform
$\langle p|\Psi\rangle = \tilde{\Psi}(p)$	$\tilde{\rho}(z) = \sum_n z^n \rho(n)$	$\tilde{\psi}(p) = \int dx e^{-\frac{i}{\hbar}px}\psi(x)$
	Laplace transform	
	$\tilde{\rho}_V(p) = \int_0^\infty e^{Vpq}\rho(q)\,dq$	
Momentum operator	$\hat{P} = -\frac{\partial}{\partial n}$	$\hat{P} = \frac{\hbar}{i}\frac{\partial}{\partial x}$
$\langle p|\hat{P}|\Psi\rangle = p\langle p|\Psi\rangle$	$\hat{P} = -\frac{1}{V}\frac{\partial}{\partial q}$	
Scaling limit	$V \to \infty$	$\hbar \to 0$
Large-deviation theory	$\rho^R(q,t) \asymp e^{-VS(q,t)}$	$\psi(x,t) \asymp e^{\frac{i}{\hbar}S(x,t)}$

The norm of the state and the sampling protocol must be preserved during the dynamics, which yields extra conditions on the Hamiltonian operator $\hat{H}$. In particular, when $\hat{H}$ preserves the L_1 norm, then we get the Master equation from Eq. (2).

It is also straightforward to verify that the Hamiltonian operator that preserves the L_2 norm must correspond to an anti-Hermitian operator, implying $\hat{H} = -\hat{H}^\dagger$. Since the dual of an L_2 function is an L_2 function, there is no difference between the evolution of the state or the sampling protocol, thus reducing Eqs. (C1) to one, yielding the Schrödinger equation. We summarize these observations in Table II.

Equations (C1) can be collected in one by defining a density operator $\hat{\phi} = |\phi^R\rangle\langle\phi^L|$ and finding its total time derivative

$$
\frac{d\hat{\phi}}{dt} \equiv \frac{d|\phi^R\rangle\langle\phi^L|}{dt} = \frac{d|\phi^R\rangle}{dt}\langle\phi^L| + |\phi^R\rangle\frac{d\langle\phi^L|}{dt}
$$
$$
= \hat{H}|\phi^R\rangle\langle\phi^L| - |\phi^R\rangle\langle\phi^L|\hat{H}^\dagger
$$
$$
= [\hat{H}, \hat{\phi}]. \tag{C2}
$$

This is also known as the Von-Neumann equation and the operator $[\hat{H}, .]$ is also referred to as the Liouville operator [53]. The statistical observation of the state of an experiment under a sampling protocol is given by $\langle O \rangle_\phi \equiv \mathrm{Tr}\hat{O}\hat{\phi}$. The time evolution of the statistical observable O along the variational solutions of the nonequilibrium action is then given by

$$
\frac{d\langle O \rangle_\phi}{dt} = \mathrm{Tr}[\hat{O}, \hat{H}]\hat{\phi} + \mathrm{Tr}\frac{\partial \hat{O}}{\partial t}\hat{\phi}. \tag{C3}
$$

For an excellent introduction to operator techniques in quantum statistical mechanics and quantum computation, see Refs. [54,55].

2. Derivation of the path integral formula

In this section we will answer the question, given the state of an experiment (PDF) at time $t = 0$ and a stochastic Hamiltonian with which the system evolves following Eqs. (C1) and (2), what is its state at an arbitrary time $t = T$?

Let us revisit the time evolution, from Eq. (3), for a PDF $\rho(n, t)$, where $n \in \mathbb{Z}_{\geqslant 0}^D$ denotes the position in state space and t is the time,

$$
\frac{\partial}{\partial t}\rho(n, t) = \underline{\hat{H}}\left(-\frac{\partial}{\partial n}, n\right)\rho(n, t). \tag{C4}
$$

Since n takes only nonnegative integer values, we can consider the moment-generating function (MGF) or $\mathcal{Z}$ transform of the PDF $\rho(n, t)$, given by $\tilde{\rho}(z, t)$, and its inverse

$$
\tilde{\rho}(z, t) = \sum_n z^n \rho(n, t), \tag{C5}
$$

$$
\rho(n, t) = \frac{1}{2\pi i}\oint dz \frac{1}{z^{n+1}}\tilde{\rho}(z, t). \tag{C6}
$$

We refer to the first line as $\mathcal{Z}$ transform and the second as inverse $\mathcal{Z}$ transform (for an introduction, see Ref. [56]). Here $z \in \mathbb{C}^D$ and the contour integral is done over any contour that encloses $z = 0$.

Since Eq. (C4) in general defines the time evolution of any distribution, in particular we can consider the time evolution
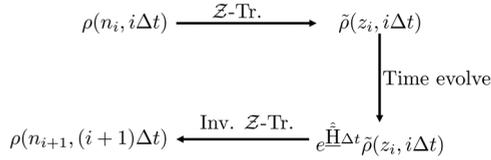
FIG. 15. Evolving distribution by Δt.

of the MGF or $\mathcal{Z}$ transform of a distribution

$$\frac{\partial \tilde{\rho}(z,t)}{\partial t} = \sum_n z^n \hat{H}\left(-\frac{\partial}{\partial n}, n\right)\rho(n,t)$$

$$= \hat{\underline{H}}\left(z, \frac{\partial}{\partial z}\right)\tilde{\rho}(z,t), \qquad (C7)$$

where $\hat{\underline{H}}$ is the time-evolution operator for the $\mathcal{Z}$-transformed distribution, which we will henceforth refer to as the $\mathcal{Z}$-Hamiltonian operator.

Using the observations

$$\sum_n z^n e^{-y\frac{\partial}{\partial n}}\rho(n,t) = z^y \sum_n z^n \rho(n,t),$$

$$\sum_n z^n \frac{n!}{(n-y)!}\rho(n,t) = z^y\left(\frac{\partial}{\partial z}\right)^y \sum_n z^n \rho(n,t),$$

it is easy to see that the $\mathcal{Z}$-Hamiltonian operator for CRN from Eq. (15),

$$\hat{H}_{\text{CRN}}\left(-\frac{\partial}{\partial n}, n\right)$$

$$= \sum_{y_\alpha, y_\beta}\left(e^{-(y_\beta - y_\alpha)\frac{\partial}{\partial n}} - 1\right)\frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}}\frac{n!}{(n-y_\alpha)!}, \qquad (C8)$$

takes the following form (for a more insightful derivation, see Ref. [19]):

$$\hat{\underline{H}}_{\text{CRN}}\left(z, \frac{\partial}{\partial z}\right) = \sum_{y_\alpha, y_\beta}(z^{y_\beta} - z^{y_\alpha})\frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}}\left(\frac{\partial}{\partial z}\right)^{y_\alpha}. \qquad (C9)$$

Returning to the question posed in the beginning of this subsection, given a Hamiltonian operator $\hat{\underline{H}}$ and $\rho(n,0)$, we wish to find $\rho(n,T)$. The approach that we will take is the following (for a pedagogical introduction to these methods, see Refs. [57,58]). First, we discretize time from 0 to T into N intervals of length Δt, such that $N\Delta t = T$. Next, we label the random variable denoting the count of the system at time $i\Delta t$ by n_i. Since, we know that the system is evolving with the given Hamiltonian, there is a relation between $\rho(n_{i+1}, (i+1)\Delta t)$ and $\rho(n_i, i\Delta t)$. To find $\rho(n_{i+1}, (i+1)\Delta t)$, we first compute the $\mathcal{Z}$ transform of $\rho(n_i, i\Delta t)$, evolve it for Δt by Eq. (C7) and compute the inverse $\mathcal{Z}$ transform at n_{i+1}.

Using the above procedure, which we represent in diagrammatic form in Fig. 15, we can read the distribution at

time $(i+1)\Delta t$ to be

$$\rho(n_{i+1}, (i+1)\Delta t)$$

$$= \frac{1}{2\pi i}\oint dz_i \frac{1}{z_i^{n_{i+1}+1}}e^{\Delta t\hat{\underline{H}}(z,\frac{\partial}{\partial z})}\sum_{n_i} z_i^{n_i}\rho(n_i, i\Delta t)$$

$$= \frac{1}{2\pi i}\oint \frac{dz_i}{z_i}\sum_{n_i} z_i^{-n_{i+1}}e^{\Delta t\hat{\underline{H}}(z,\frac{\partial}{\partial z})}z_i^{n_i}\rho(n_i, i\Delta t)$$

$$= \frac{1}{2\pi i}\oint \frac{dz_i}{z_i}\sum_{n_i} z_i^{-n_{i+1}+n_i}e^{\Delta t\underline{\mathcal{H}}(z_i,n_i)}\rho(n_i, i\Delta t), \quad (C10)$$

where in the last line, the function $\underline{\mathcal{H}}(z_i, n_i)$ is obtained by an application of the operator $\hat{\underline{H}}(z, \frac{\partial}{\partial z})$ on $z_i^{n_i}$. In particular, for CRN, using Eq. (C9), we obtain

$$\underline{\mathcal{H}}_{\text{CRN}}(z,n) = \sum_{y_\alpha, y_\beta}(z^{y_\beta - y_\alpha} - 1)\frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}}\frac{n!}{(n-y_\alpha)!}. \qquad (C11)$$

Finally, by repeated application of the above procedure at the initial PDF N times and taking the limit $N \to \infty$ and $\Delta t \to 0$, while their product is held fixed $N\Delta t = T$, we obtain

$$\rho(n_N, N\Delta t)$$

$$= \prod_{i=0}^{N-1}\left(\frac{1}{2\pi i}\oint \frac{dz_i}{z_i}\sum_{n_i} z_i^{-n_{i+1}+n_i}e^{\Delta t\underline{\mathcal{H}}(z_i,n_i)}\right)\rho(n_0, 0),$$

using $z_i = e^{p_i}$

$$= \prod_{i=0}^{N-1}\left(\int \frac{dp_i}{2\pi i}\sum_{n_i} e^{-p_i(n_{i+1}-n_i)}e^{\Delta t\underline{H}(p_i,n_i)}\right)\rho(n_0, 0)$$

$$\xrightarrow[\substack{N\to\infty \\ \Delta t\to 0}]{\lim} \int\left[\frac{dp}{2\pi i}\right]\sum[dn]e^{-\int_0^T dt(p\dot{n}-\underline{H}(p,n)dt)}\rho(n_0, 0),$$

where $[dp]$ and $[dn]$ are the path-integral measures and $\underline{H}(p,n)$. The last line yields us the path-integral formula and we define the action functional as

$$\mathcal{A}[n(t), p(t)] = \int_0^T dt[p\dot{n} - \underline{H}(p,n)dt]. \qquad (C12)$$

For an alternative derivation of the functional integral in the Doi-Peliti formalism using coherent states, see Refs. [6,20,59].

Note that For CRN, using Eq. (C11), the Hamiltonian function takes the particular form

$$\underline{H}_{\text{CRN}}(p,n) = \sum_{y_\alpha, y_\beta}(e^{p(y_\beta - y_\alpha)} - 1)\frac{k_{y_\alpha \to y_\beta}}{V^{y_\alpha - 1}}\frac{n!}{(n-y_\alpha)!}, \quad (C13)$$

which is the same as the functional form of the Hamiltonian operator in Eq. (C8), with the operator $-\frac{\partial}{\partial n}$ replaced by the momentum variable p, thus justifying our notation.

3. Optimality condition: Hamilton's equations of motion and their relation to the Hamilton-Jacobi equation

Analogous to Eq. (C12), in Eq. (6) we defined an action functional in the concentration and momentum coordinates (q, p) to be

$$\mathcal{A}[q(t), p(t)] = \int_0^T \left[p\frac{dq}{dt} - H(p,q)\right]dt,$$

where the relation of the Hamiltonian function $H(p, q)$ to $\underline{H}(p, n)$ was discussed in Eq. (17). As discussed around Eq. (8), the path integral is dominated by the value of the action functional around the optimal or stationary path, i.e., the path for which the first variation of the action is zero.

To calculate the first variation of the action, we consider the difference of the action functional between a path $(q(t), p(t))$ and $(q(t) + \delta q(t), p(t) + \delta p(t))$. In the following equation, we will suppress the time dependence to simplify notation and keep only the first-order terms.

$$
\begin{aligned}
\delta \mathcal{A}[q, p] \\
= \mathcal{A}[q + \delta q, p + \delta p] - \mathcal{A}[q, p] \\
= \int dt \left\{ \delta p \left(\frac{dq}{dt} - \frac{\partial H}{\partial p} \right) + \left(p \frac{d\delta q}{dt} - \frac{\partial H}{\partial q} \delta q \right) \right\} \\
= \int dt \left\{ \delta p \left(\frac{dq}{dt} - \frac{\partial H}{\partial p} \right) - \delta q \left(\frac{dp}{dt} - \frac{\partial H}{\partial q} \delta q \right) + p \delta q \Big|_0^T \right\},
\end{aligned}
\tag{C14}
$$

where the last line is obtained by integrating by-parts. If we fix the end points of the path in configuration space, then we can set $\delta q(0) = \delta q(T) = 0$, getting rid of the last term in the above equation.

Using Eq. (C14) and

$$
\delta \mathcal{A}[q(t), p(t)] = \frac{\delta \mathcal{A}}{\delta q(t)} \delta q(t) + \frac{\delta \mathcal{A}}{\delta p(t)} \delta p(t),
$$

we can read off the variation of the action functional in configuration and momentum around $(q(t), p(t))$ to be

$$
\begin{aligned}
\frac{\delta \mathcal{A}}{\delta q(t)} = -\left(\frac{dp}{dt} + \frac{\partial H}{\partial q} \right), \\
\frac{\delta \mathcal{A}}{\delta p(t)} = \left(\frac{dq}{dt} - \frac{\partial H}{\partial p} \right).
\end{aligned}
\tag{C15}
$$

The optimality condition $\mathcal{A}[q^*, p^*] = 0$ for a trajectory $(q^*(t), p^*(t))$ then yields

$$
\begin{aligned}
\frac{dp^*}{dt} = -\frac{\partial H}{\partial q^*}, \\
\frac{dq^*}{dt} = \frac{\partial H}{\partial p^*}.
\end{aligned}
\tag{C16}
$$

The last two lines are also referred to as Hamilton's equations of motion (EoM) and are the equations that any optimal trajectory must satisfy.

There is a rich and deep mathematical structure underlying these equations, which is the subject of symplectic geometry. For instance, the change in value of an observable function $f(p, q)$ along an optimal trajectory (q^*, p^*, t) can be given in terms of its Poisson-bracket commutator with the Hamiltonian,

$$
\begin{aligned}
\frac{df}{dt} = \frac{\partial f}{\partial t} + \frac{\partial f}{\partial q} \frac{dq}{dt} + \frac{\partial f}{\partial p} \frac{dp}{dt} \\
= \frac{\partial f}{\partial t} + \frac{\partial f}{\partial q} \frac{\partial H}{\partial p} - \frac{\partial f}{\partial p} \frac{\partial H}{\partial q} \\
= \frac{\partial f}{\partial t} + \{f, H\}.
\end{aligned}
\tag{C17}
$$

This immediately shows that the value of a time-independent Hamiltonian as well as any time independent operator that commutes with the Hamiltonian stays constant along equations of motion [notice the resemblance of the above equation with Eq. (C3)]. For a classical introduction to the subject, see Refs. [21,25].

To end this subsection, following Ref. [60], we provide a proof of the equivalence between Hamilton's equations of motion and Hamilton-Jacobi equation on paths where the velocity and the momentum are related by a Legendre transform:

$$
\begin{aligned}
p = \frac{\partial S}{\partial q}, \\
\dot{p} = \frac{d}{dt} \left(\frac{\partial S}{\partial q} \right) \\
= \frac{\partial S}{\partial q} \dot{q} + \frac{\partial S}{\partial q} t \\
= \frac{\partial S}{\partial q} q \frac{\partial H}{\partial p} - \frac{\partial H}{\partial q} - \frac{\partial H}{\partial p} \frac{\partial p}{\partial q} \\
= \frac{\partial p}{\partial q} \frac{\partial H}{\partial p} - \frac{\partial H}{\partial p} \frac{\partial p}{\partial q} - \frac{\partial H}{\partial q} \\
= -\frac{\partial H}{\partial q},
\end{aligned}
$$

where the first and fourth line use the Hamilton-Jacobi equation. The fourth line also makes use of the relation between optimal momentum and velocity or the Legendre-transform condition.

4. Second variational derivative of action functional: Onsager-Machlup action and convexity of instanton

In this subsection, we will first calculate the second variational derivative of the action functional. Next, we provide a derivation of the Onsager-Machlup action, followed by comments on the convexity of the action functional around any optimal trajectory.

To simplify notation we use the following convention. For vectors v, w, we represent their dot product $v^T w$ as vw. For a quadratic form M, we represent $M(v, v) \equiv v^T M v$ simply as Mv^2. Finally, the time derivative and perturbation of a function x is given by

$$
\dot{x} = \frac{dx}{dt},
$$

$$
x^+ = x + \delta x,
$$

respectively.

Recall, in Eq. (6), the action functional was defined to be

$$
\mathcal{A}[q, p] = \int_0^T dt \{ p\dot{q} - H \}.
$$

Thus, the variation in the action functional up to terms second order in the variation is

$$\delta \mathcal{A} = \mathcal{A}[q^+, p^+] - \mathcal{A}[q, p]$$

$$= \int dt \left\{ \delta p \left(\dot{q} - \frac{\partial H}{\partial p} \right) - \delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) \right\} + \delta p \delta \dot{q} - \frac{1}{2} [\delta p \quad \delta q] \begin{bmatrix} \frac{\partial^2 H}{\partial p \partial p} & \frac{\partial^2 H}{\partial p \partial q} \\ \frac{\partial^2 H}{\partial q \partial p} & \frac{\partial^2 H}{\partial q \partial q} \end{bmatrix} \begin{bmatrix} \delta p \\ \delta q \end{bmatrix} + p \delta q \Big|_0^T$$

$$= \int dt \left\{ \delta p \left(\dot{q} - \frac{\partial H}{\partial p} \right) - \delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) - \frac{1}{2} [\delta p \quad \delta q] \begin{bmatrix} \frac{\partial^2 H}{\partial p \partial p} & \left(-d_t + \frac{\partial^2 H}{\partial p \partial q}\right) \\ \left(d_t + \frac{\partial^2 H}{\partial q \partial p}\right) & \frac{\partial^2 H}{\partial q \partial q} \end{bmatrix} \begin{bmatrix} \delta p \\ \delta q \end{bmatrix} \right\} + \left(p \delta q + \frac{1}{2} \delta p \delta q \right) \Big|_0^T$$

$$= \int dt \left\{ \delta p \left(\dot{q} - \frac{\partial H}{\partial p} \right) - \delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) - \frac{1}{2} \frac{\partial^2 H}{\partial p \partial p} \left[\delta p - \left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left(\delta \dot{q} - \frac{\partial^2 H}{\partial p \partial q} \delta q \right) \right]^2 \right.$$

$$\left. + \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left[\left(d_t - \frac{\partial^2 H}{\partial p \partial q} \right) \delta q \right]^2 - \frac{\partial^2 H}{\partial q \partial q} \delta q^2 \right] \right\} + \left(p \delta q + \frac{1}{2} \delta p \delta q \right) \Big|_0^T, \tag{C18}$$

where we have implicitly assumed that the original and perturbed phase-space trajectories lie in the same $H = h$ submanifold.

Now, if we assume that the two trajectories are constrained to the same end points in configuration space, i.e., $\delta q(0) = \delta q(T) = 0$, then the boundary terms vanish. Moreover, if we assume that the momentum assignment before and after the perturbation is optimal, i.e., p and p^+ are assigned to q and q^+ such that the phase-space point is at its Legendre transform before and after the perturbation, then we have

$$\text{for } (q, p): \quad 0 = \left(\dot{q} - \frac{\partial H}{\partial p} \right), \tag{C19}$$

$$\text{for } (q^+, p^+): \quad \dot{q}^+ = \frac{\partial H(p^+, q^+)}{\partial p}, \quad \delta \dot{q} = \frac{\partial^2 H}{\partial p \partial q} \delta q + \frac{\partial^2 H}{\partial p \partial p} \delta p, \quad 0 = \delta p - \left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left(\delta \dot{q} - \frac{\partial^2 H}{\partial p \partial q} \delta q \right). \tag{C20}$$

Substituting Eqs. (C19) and (C20) in Eq. (C18) and making use of the assumptions, we get the variation in the action functional to be

$$\delta \mathcal{A} = \int dt \left\{ -\delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) + \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left[\left(d_t - \frac{\partial^2 H}{\partial p \partial q} \right) \delta q \right]^2 - \frac{\partial^2 H}{\partial q \partial q} \delta q^2 \right] \right\}$$

$$= \int dt \left\{ -\delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) + \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left[d_t \delta q + \frac{\partial H}{\partial p} - \frac{\partial H}{\partial p} - \frac{\partial^2 H}{\partial p \partial q} \delta q \right]^2 - \frac{\partial^2 H}{\partial q \partial q} \delta q^2 \right] \right\}$$

$$= \int dt \left\{ -\delta q \left(\dot{p} + \frac{\partial H}{\partial q} \right) + \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \left[\dot{q}^+ - \dot{q} - \frac{\partial^2 H}{\partial p \partial q} \delta q \right]^2 - \frac{\partial^2 H}{\partial q \partial q} \delta q^2 \right] \right\}, \tag{C21}$$

where we add and subtract the same quantity in the second line and make use of the Legendre transform condition in the third line.

Onsager-Machlup action. $\delta \mathcal{A}$ measures the log-conditional improbability of trajectory q^+ given base trajectory q. If q is Hamilton's equation of motion along $p = 0$, then we have

$$\left(\dot{p} + \frac{\partial H}{\partial q} \right) = 0,$$

$$\frac{\partial^2 H}{\partial q \partial q} \Big|_{p=0} = 0,$$

$$\frac{\partial^2 H}{\partial q \partial p} \Big|_{p=0} = 0, \tag{C22}$$

and we thus recover the Onsager-Machlup action

$$\mathcal{A}_{\text{OM}}[q^+] \equiv \mathcal{A}[(q^+, 0)] - \mathcal{A}[(q, 0)]$$

$$= \int dt \left\{ \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} [\dot{q}^+ - \dot{q}]^2 \right] \right\}. \tag{C23}$$

Recall from the discussion around Eq. (20) that the equations of motion along which $p = 0$ are the deterministic trajectories of the system. Thus, the Onsager-Machlup action measures the log improbability of a trajectory q^+ conditioned on the deterministic or relaxation trajectory q for the stochastic system. For a similar derivation and applications to population biology, see Ref. [47].

Convexity of action functional around an optimal path. If q is an escape curve for a CRN Hamiltoninan, then only the first equation of Eqs. (C22) holds since $p(q_{esc}) \neq 0$ [see Eqs. (24)]. The variation in action around the optimal path then becomes

$$\delta\mathcal{A} = \int dt \left\{ \frac{1}{2} \left[\left(\frac{\partial^2 H}{\partial p \partial p} \right)^{-1} \right. \right.$$
$$\left. \left. \times \left[\dot{q}^+ - \dot{q} - \frac{\partial^2 H}{\partial p \partial q} \delta q \right]^2 - \frac{\partial^2 H}{\partial q \partial q} \delta q^2 \right] \right\}.$$
(C24)

Notice that due to convexity of the Hamiltonian in p, the first term is positive definite. However, nothing can be said *a priori* about the second term since the Hessian of the Hamiltonian in q for multistationary networks is in general nonconvex. For the applications demonstrated in the paper, the variation is indeed positive definite and a global minimum exists in all cases. Generically, however, as can be seen from the equation above, not all optimal curves are global minimizers of the action. For an example of an infinite-dimensional stochastic system exhibiting a saddle optimal solution, we refer the readers to Chapter 8 [61]. The authors are not aware of a similar example for a finite-dimensional chemical reaction network at this point, and leave it as a question for future investigation.

APPENDIX D: DETAILS OF AFGD

1. Notation

As noted in Table I, starting Sec. III A, we use the following notation:

Configuration space : $Q \subset \mathbb{R}^D$

$\mathcal{P}_Q = \{q : [0, 1] \to Q : q(0) = q_I, q(1) = q_F\}$

Configuration space curve : $q \in \mathcal{P}_Q$

Phase space : $T^*Q \subset \mathbb{R}^{2D}$

$\mathcal{P}_{T^*Q} = \{(q, p) : [0, T] \to T^*Q : q^*(0) = q_I, q^*(T) = q_F\}$

Phase space trajectory : $\gamma \in \mathcal{P}_{T^*Q}$

We refer to paths in configuration space as "curves," and paths in phase space as "trajectories." This is an unusual choice intended to avoid confusion and simplify terminology, and will be used consistently through this section. Notice that curves are parametrized on the unit interval [0,1], which is the normalized arc length. However, trajectories are parametrized with time, on the interval $[0, T]$ (for more details, see Sec. III A).

Since computationally we have to work with discrete paths, we now introduce notation for discretization. Say we sample a continuous path $x \in X$ on N points, we denote its discrete counterpart with $\tilde{x} \in \tilde{X}$ where

$$\tilde{x}(n) = x\left(\frac{n}{N}\right), \quad \text{where } n \in [1, N] \subset \mathbb{Z}.$$

We denote the Ith iteration of the algorithm with a superscript I. For instance, the configuration space curve in the first

iteration is denoted by $\tilde{q}^1$, where

$$\tilde{q}^1(n) = q^{IC}\left(\frac{n}{N}\right).$$

2. Initial condition

To begin the descent toward an optimal trajectory, in principle, one can start with an arbitrary curve connecting the end points q_I and q_F that does not pass through or go around another fixed point and remains in q_I's basin of escape. In practice however, the descent might require a large number of sample points and high numerical accuracy to obtain the optimal trajectory. In this subsection we provide a method to classify the optimal trajectory and pick an appropriate initial condition with the desired number of sample points effectively.

We classify the optimal trajectory near a fixed point by linearizing Hamilton's equations of motion (EoM) and analyzing the eigenvalues of the Hessian. If the eigenvalues have a nonzero imaginary part, then the optimal trajectory spirals into or out of the fixed point. An exposition of this technique can be found in Chapter 4 of Ref. [61] (their "activation trajectory" is our "escape trajectory"), but we repeat the relevant construction here.

Let the fixed point be denoted by $\underline{q}$. Recall that the momentum value at the fixed point in the $H(p, q) = 0$ submanifold is identically zero, i.e., $p(\underline{q}) = 0$. Linearizing Hamilton's equations of motion around the fixed point, we get

$$\begin{bmatrix} \frac{dq}{dt} \\ \frac{dp}{dt} \end{bmatrix} = \begin{bmatrix} \frac{\partial H}{\partial p} \\ -\frac{\partial H}{\partial q} \end{bmatrix}_{(\underline{q},0)} + \begin{bmatrix} \frac{\partial H}{\partial p \partial q} & \frac{\partial H}{\partial p^2} \\ -\frac{\partial H}{\partial q^2} & -\frac{\partial H}{\partial q \partial p} \end{bmatrix}_{(\underline{q},0)} \begin{bmatrix} (q-\underline{q}) \\ (p-0) \end{bmatrix}$$
$$= \begin{bmatrix} \frac{\partial H}{\partial p \partial q} & \frac{\partial H}{\partial p^2} \\ 0 & -\frac{\partial H}{\partial q \partial p} \end{bmatrix}_{(\underline{q},0)} \begin{bmatrix} (q - \underline{q}) \\ (p - 0) \end{bmatrix},$$
(D1)

where the second line is obtained by using the fact that $\underline{q}$ is a fixed point and $p(\underline{q}) = 0$. To simplify notation, we denote the column vector $(q - \underline{q}, p - 0)^T$ by $\Delta\gamma$. Then we can rewrite the above equation as

$$\frac{d\Delta\gamma}{dt} = M\Delta\gamma,$$
$$M = \sum_{i=1}^{2n} \lambda_i v_i v_i^T,$$
(D2)

where M is the Hessian with eigenvectors v_i and their corresponding eigenvalues λ_i.

Let the initial condition be denoted by q^{IC}.

a. Systems without voriticity. If all λ_is are real, then we know that the optimal trajectory will not spiral into the saddle or out of the stable fixed point. In this case, we pick the initial condition to be a straight line starting from the stable fixed point toward the saddle fixed point:

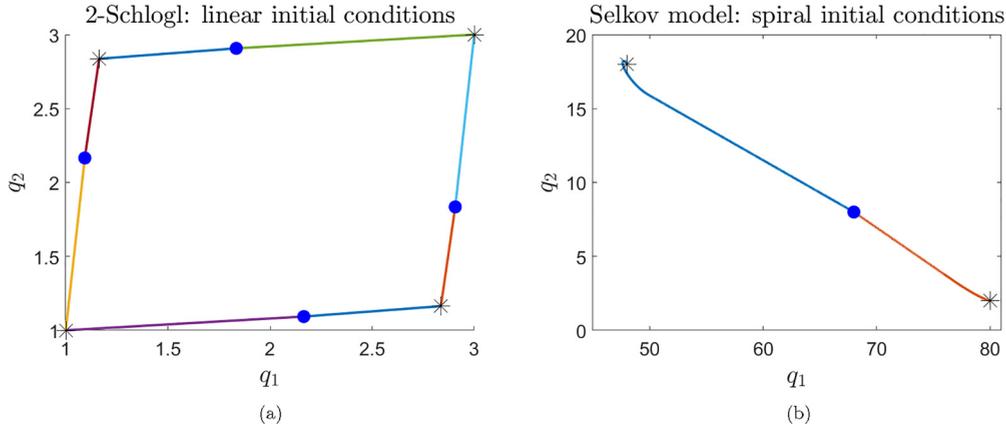$$q^{IC}(s) = q_F + (1 - s)(q_I - q_F).$$
(D3)

FIG. 16. Initial conditions for the AFGD algorithm. (a) 2-Schlögl model and (b) Selkov model.

The discrete initial curve for the algorithm will then be given by

$$\tilde{q}^1(n) = q^{IC}\left(\frac{n}{N}\right). \tag{D4}$$

b. Systems with vorticity. If some λ_is are complex, then the trajectory spirals outwards or inwards around the stable or saddle fixed point, respectively. For the purpose of this subsection, let us assume that the eigenvalues are complex around the stable fixed point q_I and real around the saddle fixed point q_F.

All the trajectories that depart from the fixed point, including the escape trajectory in the $H(p, q) = 0$ submanifold, lie in the D-dimensional submanifold[5] (in the $2D$-dimensional phase space) given by the superposition of the eigenvectors with nonzero eigenvalues. Thus, we can write the escape trajectory near the fixed point as

$$\Delta\gamma(0) = \sum_i \mu_i v_i,$$

$$\Delta\gamma(t) = \sum_i \mu_i e^{\lambda_i t} v_i, \tag{D5}$$

where we must pick μ_i such that

$$\mu_i \in [-1, 1] \subset \mathbb{R} \quad \text{for all } i,$$
$$\mu_i = 0 \quad \text{if } \text{Re}(\lambda_i) < 0,$$
$$\text{and } |\mu_i| = |\mu_j| \quad \text{if } \overline{\lambda_i} = \lambda_j,$$

which ensures that we only consider the repelling subspace and the trajectory has purely real coordinates. There is still some arbitrariness in the selection of the magnitude and sign of the coefficients which must be resolved through experimentation on a case-by-case basis, the goal being to select coefficients such that the curve spirals outwards [$\text{Re}(\lambda_i) > 0$] from q_I toward q_F.

We know that the trajectory obtained by integrating the linearization is only valid in some small distance around the fixed point. To end the curve at the saddle fixed point, we find time t such that the first derivative of the spiral

matches the slope of the joining line, i.e., solve for $t = t^* > 0$ such that

$$\frac{\frac{dq_i}{dt}}{\frac{dq_1}{dt}} = \frac{[q(t) - q_F]_i}{[q(t) - q_F]_1}, \quad \text{for all } i \in [2, n]. \tag{D6}$$

We obtain the complete initial condition by taking the union of the spiral part and straight part,

$$q_{\text{spiral}}(u) = q_I + \Delta\gamma(u) \quad \text{for } u \in [-T, t^*],$$
$$q_{\text{straight}}(u) = q_F + (1 - u)(q_I + \Delta\gamma(t^*) - q_F) \quad \text{for } u \in [0, 1],$$
$$q^{IC} = q_{\text{spiral}} \bigcup q_{\text{straight}}. \tag{D7}$$

We then use the length along the curve to parametrize q^{IC}, and obtain the initial condition

$$\tilde{q}^1(n) = q^{IC}\left(\frac{n}{N}\right). \tag{D8}$$

3. Lift curve to trajectory

In every iteration, we begin with a configuration space curve q^I and lift it to a phase-space trajectory $(q, p, t)^I$, which we then use to obtain a variation curve δq^I, which gives us a curve for the next iteration $q^{I+1} = q^I + \delta q^I$. In this subsection, we define the *lift* map and comment on methods for its implementation.

For a discretized curve $\tilde{q}$, for each segment $\tilde{q}(i)$ to $\tilde{q}(i + 1)$ we assign an optimal momentum value $\mathfrak{p}(i)$ at the center of the segment and a time interval $\Delta t(i)$ denoting the time taken for transition from $q(i)$ to $q(i + 1)$ (see Fig. 5), using the following equations. The process of assigning an optimal momentum value using a convex Hamiltonian is an instance of a Legendre-transformation, for a detailed exposition see Ref. [62]:

$$0 = H\left(\mathfrak{p}(n), \tilde{q}(n) + \frac{\Delta\tilde{q}(n)}{2}\right),$$

$$\frac{\Delta\tilde{q}(n)}{\Delta t} = \left.\frac{\partial H}{\partial p}\right|_{\left(\mathfrak{p}(n), \tilde{q}(n) + \frac{\Delta\tilde{q}(n)}{2}\right)}, \tag{D9}$$

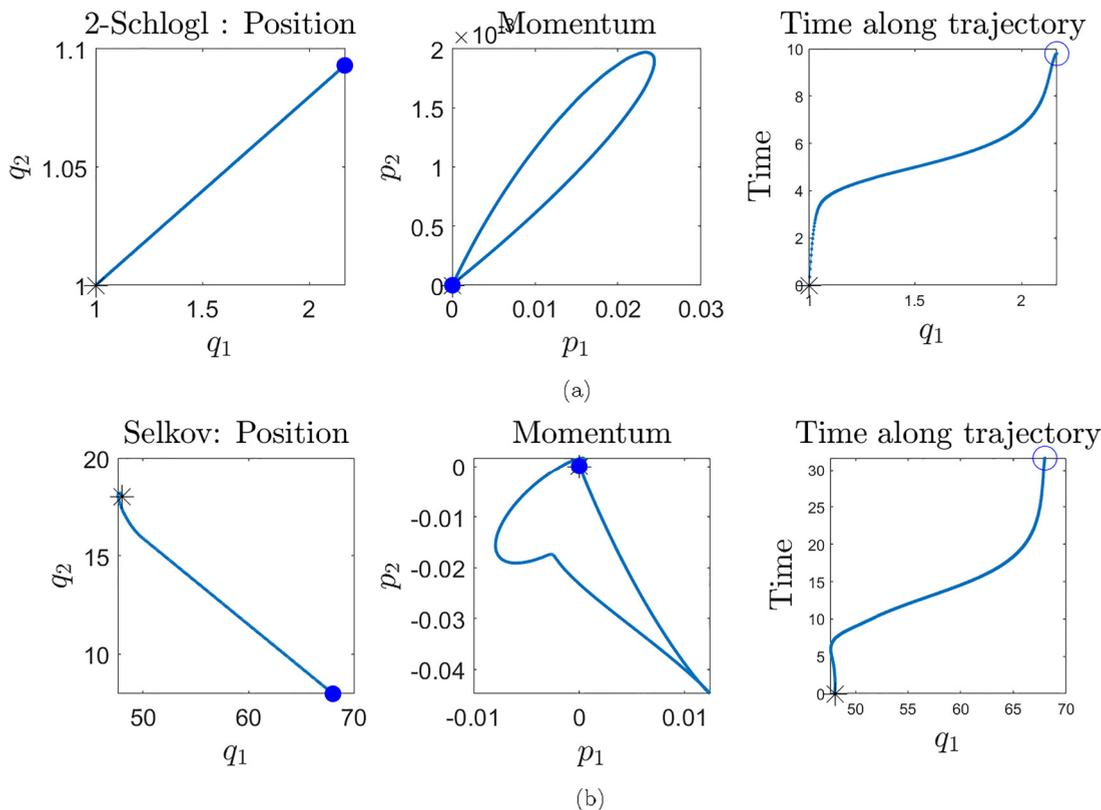[5]Also called the Lagrangian submanifold [61].

FIG. 17. Phase space trajectory obtained by lifting the initial conditions shown in Fig. 16. (a) 2-Schlögl model initial condition phase space trajectory and (b) Selkov model initial condition phase space trajectory.
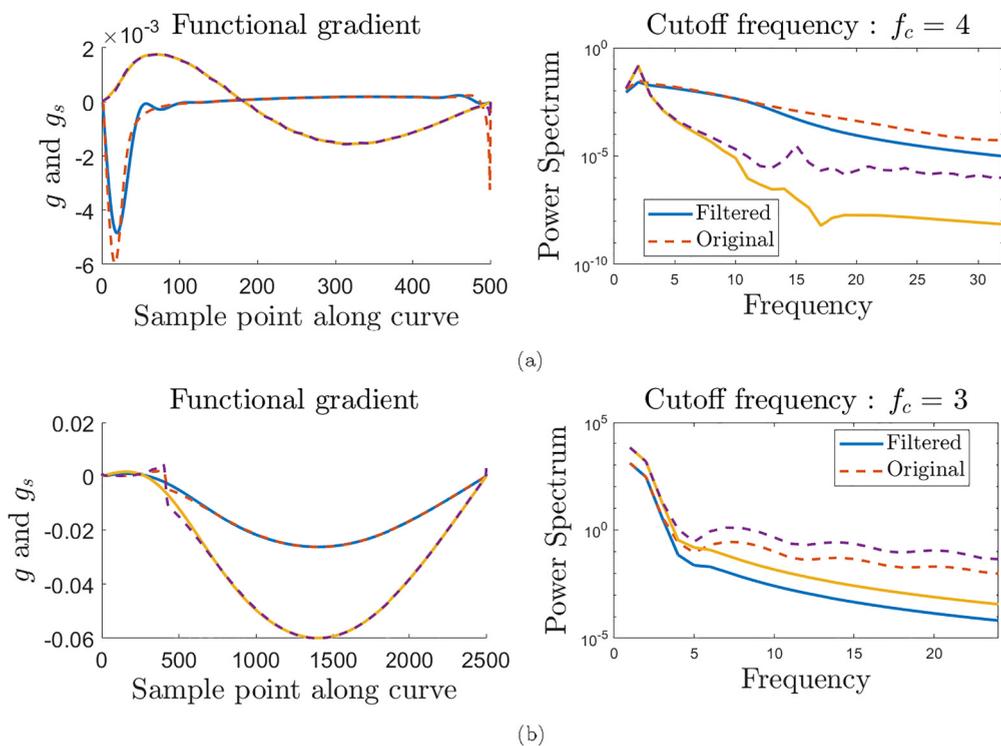

FIG. 18. Extracting functional gradient from phase-space trajectory and smoothing it. In all panels, the unfiltered and filtered curves are represented by dashed and solid lines, respectively. (a) 2-Schlögl model original and filtered functional gradient at the first iteration and (b) Selkov model original and filtered functional gradient at the first iteration.

or equivalently

$$\mathfrak{p}(n), \Delta t(n) = \arg\max_p \min_{\Delta t}$$

$$\left\{ p\Delta\tilde{q}(n) - \left[H\left(p, \tilde{q}(n) + \frac{\Delta\tilde{q}(n)}{2} \right) - 0 \right]\Delta t \right\}. \quad \text{(D10)}$$

The $D + 1$ values are assigned to simultaneously satisfy the $H(p, q) = 0$ constraint (one equation) and the Legendre transform condition (D equations), as shown in Eq. (D9). In fact, the two problems can equivalently be assimilated in a single optimization problem, as shown in Eq. (D10), where Δt plays the role of a Lagrange-multiplier. Since the Hamiltonian is convex in momentum, the objective function is concave in p, so a unique max exists. At the optimal momentum assignment, the gradient of H in p and dq are in the same direction, and the magnitude of the two vectors is made equal via Δt.

Finally, we appropriately assign the momentum and time assignment along the configuration curve as shown below and obtain a discrete phase-space trajectory:

$$t(1) = 0,$$

$$t(n) = \sum_{m=1}^{n-1} \Delta t(m),$$

$$\tilde{p}(1) = 0 = \tilde{p}(N),$$

$$\tilde{p}(n) = \frac{\mathfrak{p}(n-1) + \mathfrak{p}(n)}{2},$$

$$\tilde{\gamma}(n) = (\tilde{q}, \tilde{p}, t)(n). \quad \text{(D11)}$$

We call this procedure as the *lift* of a configuration curve to phase-space trajectory. Denoting the lift map by Λ, we define it as

$$\Lambda : \tilde{\mathcal{P}}_Q \rightarrow \tilde{\mathcal{P}}_{T^*Q},$$

$$\tilde{q} \mapsto \tilde{\gamma} \equiv (\tilde{q}, \tilde{p}, t). \quad \text{(D12)}$$

We denote the inverse or the projection map from phase-space trajectory to configuration space curve by Π, and define it to be

$$\Pi : \tilde{\mathcal{P}}_{T^*Q} \rightarrow \tilde{\mathcal{P}}_Q,$$

$$\tilde{\gamma} \mapsto \tilde{q}. \quad \text{(D13)}$$

Comments on implementation. Equations (D9) can be solved numerically using a nonlinear least-squares problem solver, such as MATLAB's *lsqnonlin*. Equation (D10) can be solved using a constrained optimization solver, such as MATLAB's *fmincon*.

Equation (D9) can also be solved analytically by first doing a change of coordinates, like the one in Eq. (C11), in which the Hamiltonian function and its derivatives are polynomials rather than consisting of exponential terms.

$$p \rightarrow z = e^p$$

$$\mathcal{H}(q, z) = 0$$

$$\frac{\Delta q_i}{\Delta q_1} = \frac{\frac{\partial\mathcal{H}}{\partial z_i} z_1}{\frac{\partial\mathcal{H}}{\partial z_1} z_i} \quad i \in [2, n] \subset \mathbb{Z}.$$

This gives D equations for D components of z. We can then use this to solve for Δt

$$\Delta t = \Delta q_1 \left(z_1 \frac{\partial\mathcal{H}}{\partial z_1} \right)^{-1}.$$

It must be noted that numerically solving for roots of polynomials can get costly in high dimensions. Also, the numerical precision decreases as we go from z to p coordinates, and it is the latter that we need for integrating Hamilton's equations of motion.

Irrespective of the implementation, for calculating momentum and time assignments along the complete trajectory, the problem (D10) or (D9) must be solved individually for each segment. Thus, a parallel implementation over multiple cores is natural and can save a lot of computational time. Also, we recommend smoothing the momentum values along the trajectory using a moving average at this step to partially reduce noise introduced by discretization and even out the effect of numerical errors. For a concrete MATLAB implementation, see Ref. [13].

4. Functional gradient—Obtaining and filtering

In the action functional gradient descent (AFGD) algorithm, we perform a functional gradient descent using negative of the functional gradient of the action functional, derived in Eq. (C15). Denoting the functional gradient by g, we get

$$g = -\frac{\delta\mathcal{A}}{\delta q}$$

$$= -\left(\frac{dp}{dt} + \frac{\partial H}{\partial q} \right).$$

Thus, we define the discrete functional gradient $\tilde{g}$ as

$$\tilde{g}(m) = \left[2\frac{\mathfrak{p}(m) - \mathfrak{p}(m-1)}{\Delta t(m) + \Delta t(m-1)} + \frac{\partial H}{\partial q}\bigg|_{(\tilde{q}(m), \tilde{p}(m))} \right]. \quad \text{(D14)}$$

The result of naively performing a functional gradient descent using the above gradient is displayed in Fig. 6. It can be seen that any numerical inaccuracies in solving for the momentum will be amplified by taking the time derivative, resulting in self-amplifying noise and instability. To smooth the noisy signal thus produced, we employ a filtering routine, that we explain in the next subsection, and obtain a smooth discrete function $\tilde{g}_s$ that we use for updating the algorithm. For a comparison of the unfiltered and filtered gradients, see Fig. 18.

5. Filtering and resampling routines

a. Filtering

To filter a discrete function $\tilde{y}$ which takes value at N points, we employ the following routine.

First, we subtract a straight line joining the end points from the function, to get a new function $\delta\tilde{y}$ which is identically zero at the end point:

$$\tilde{y}_{st}(n) = \tilde{y}(N) + \left(1 - \frac{n-1}{N-1} \right)[\tilde{y}(1) - \tilde{y}(N)],$$

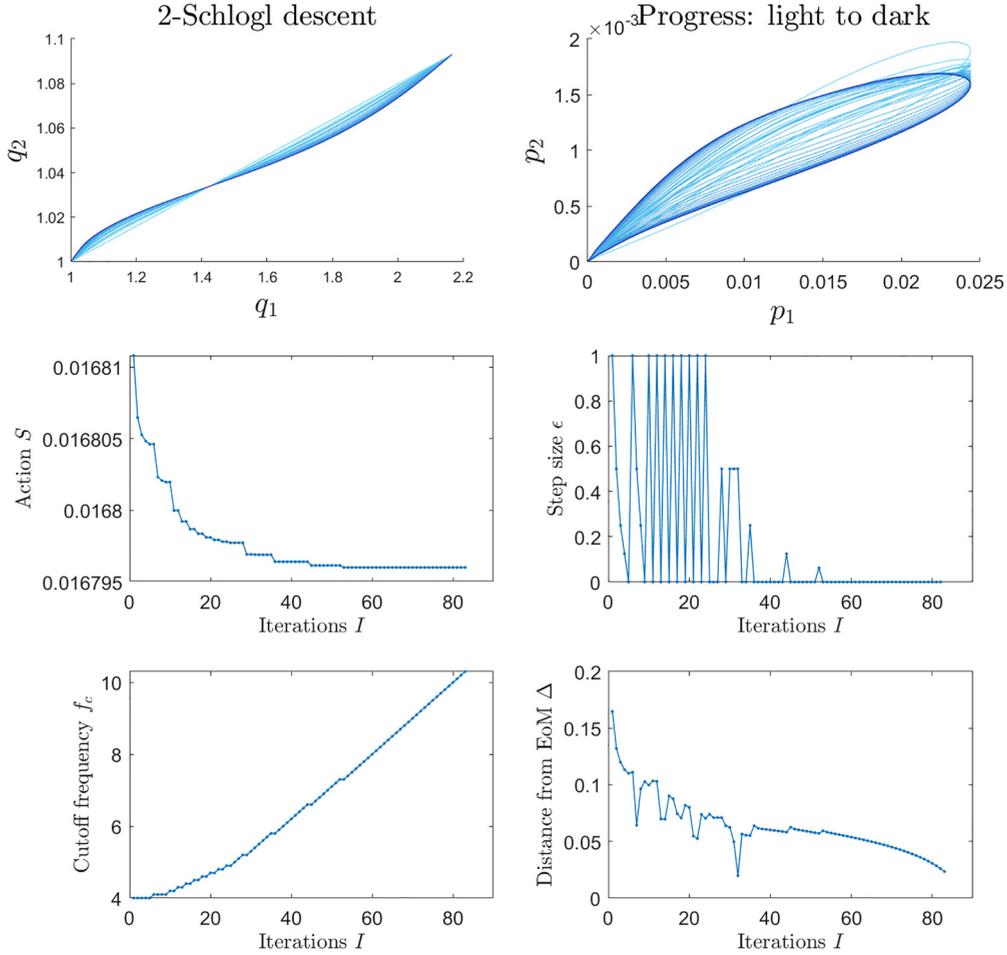$$\delta\tilde{y} = \tilde{y} - \tilde{y}_{st}.$$

FIG. 19. Descent progress and summary for the 2-Schlögl model for a trajectory with 500 points.

Next, we define a concatenated function $\delta\tilde{y}_c$ of size $2N-1$ which is obtained by juxtaposing a flipped copy with a negative sign next to the original signal and by removing the duplicate point at N:

$$\delta\tilde{y}_c(u) = \begin{cases} -\delta\tilde{y}(N-u+1) & \text{if } u \in [1, N-1], \\ \delta\tilde{y}(u-N+1) & \text{if } u \in [N, 2N-1]. \end{cases}$$

Next, we apply a Butterworth low-pass filter $\mathcal{B}(\cdot, f_c)$,[6] with some cutoff frequency f_c, to the concatenated function and obtain a filtered function $\delta\tilde{y}_{cs}$. Notice $\delta\tilde{y}_c$ is an odd function across the mid-point N, and thus $\delta\tilde{y}_{cs}$ will also be of the same form. Concatenation before filtering is a common technique employed in signal processing, without which the end points are not guaranteed to remain at zero after applying the low-pass filter.

$$f_c := \text{Cutoff frequency,}$$
$$\delta\tilde{y}_{cs} = \mathcal{B}(\delta\tilde{y}_c, f_c).$$

Finally, we obtain the desired signal $\tilde{y}_s$ by picking out only the second half of the smoothed concatenated function $\delta\tilde{y}_{cs}$

and adding that back to the straight line from the first step

$$\delta\tilde{y}_s(n) = \delta\tilde{y}_{cs}(n+N-1) \quad \text{for } n \in [1, N],$$
$$\tilde{y}_s = \tilde{y}_{st} + \delta\tilde{y}_s. \tag{D15}$$

b. Resampling

To resample the discrete curve $\tilde{q}$, we first need a parametrization. Consider the continuous curve q obtained by interpolation. Then there are two choices of parametrization canonically available to us, namely

(1) Arc-length parametrization

$$s(x) = \frac{\int_{q_I}^x |dq|}{\int_{q_I}^{q_F} |dq|}, \tag{D16}$$

(2) Time parametrization

$$s(x) = \frac{t(x)}{t(q_F)}. \tag{D17}$$

Now we obtain a new discrete function $\tilde{q}_u$ uniform in a chosen parametrization $s(x)$ by finding a point x where $s(x) = m/N$, i.e.,

$$\tilde{q}_u(m) = \bar{x} \quad \text{such that } s(\bar{x}) = \frac{m}{N} \text{ for } m \in [1, N]. \tag{D18}$$

---

[6]For our implementation we choose Butterworth filter of order 4, but this choice is arbitrary and can be experimented with.
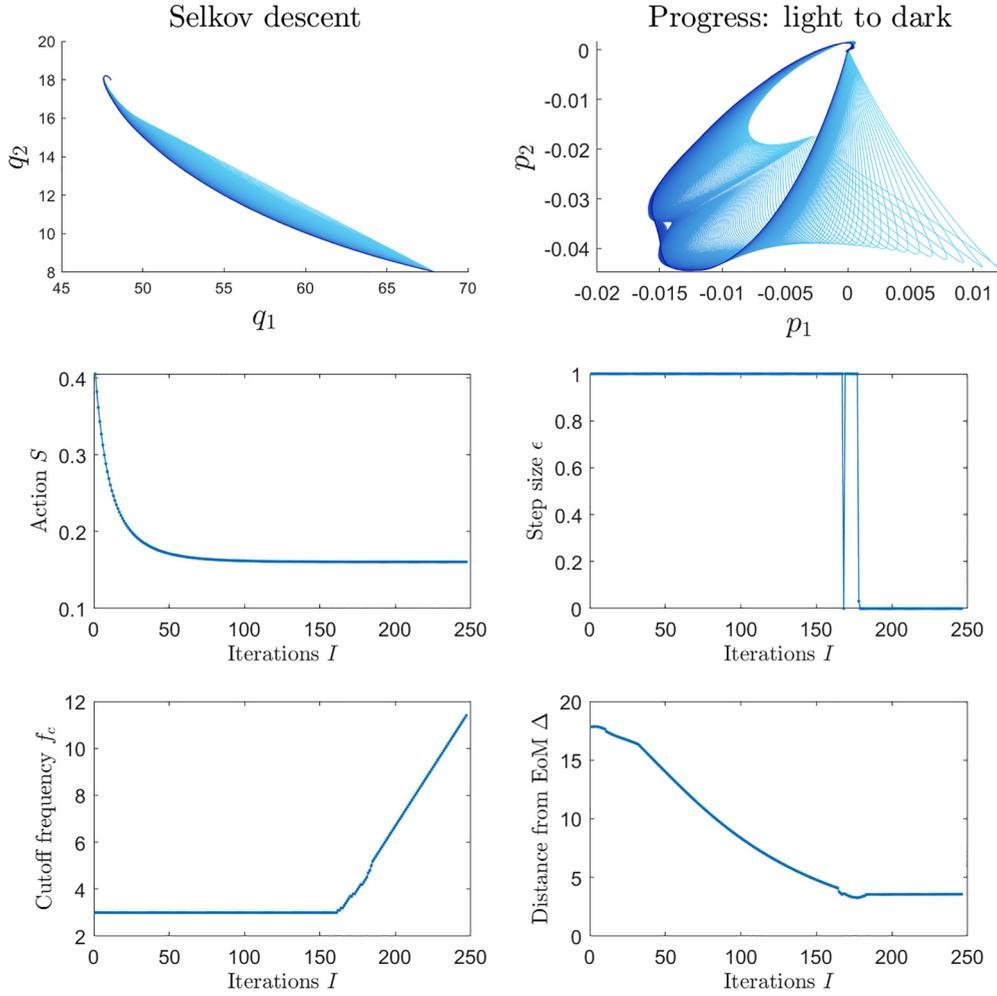
FIG. 20. Descent progress and summary for the Selkov model for a trajectory with 2500 points.

We will call a discrete curve $\tilde{q}_u$ space-uniform or time-uniform sampled if we use the arc-length parametrization in Eq. (D16) or time parametrization in Eq. (D17), respectively.

6. Pick step size

Once we have a descent direction $\tilde{g}_s$ obtained by filtering the functional gradient, we need to pick a step size that ensures that the value of the action functional is strictly decreasing. More precisely, we need to pick $\epsilon > 0$ such that

$$\mathcal{A}[\tilde{\gamma}^\epsilon] < \mathcal{A}[\tilde{\gamma}],$$
$$\text{where} \quad \tilde{\gamma} = \Lambda(\tilde{q})$$
$$\text{and} \quad \tilde{\gamma}^\epsilon = \Lambda(\tilde{q} + \epsilon \tilde{g}_s). \quad \text{(D19)}$$

Ideally, to maximize descent, we want to pick the largest $\epsilon > 0$ such that the above conditions are satisfied. In practice, however, this will require us to solve another optimization problem which can be rather time consuming. Thus, for ease of implementation we employ the *backtracking line search* method in which one starts from a large value for ϵ and keeps making it smaller until the conditions are satisfied. For an exposition of the method and more sophisticated "line search" algorithms for picking a step size, see Ref. [36].

If the step size is below a threshold $\epsilon < \epsilon_{\text{thresh}}$ or change in the value of the action is too small $|\mathcal{A}[\tilde{\gamma}^\epsilon] - \mathcal{A}[\tilde{\gamma}]| < \Delta S_{\text{thresh}}$, then we end the search and assign a step size $\epsilon = 0$. This indicates that the algorithm can not descend further with the given conditions and takes it to the next phase of either updating cutoff frequency or increasing the number of sample points.

As an illustration of convergence, see the fourth panel in Figs. 10, 19, and 20. Notice that in all figures, the step size is identically zero for a few iterations, thus indicating that the algorithm must either be terminated or taken to the next phase.

7. Updating cutoff frequency and increasing sample points during descent

a. Updating cutoff frequency

The low-pass filter on the functional gradient serves the purpose of controlling noise due to discretization and numerical solving. However, it also cuts off meaningful signal in the gradient, especially in the beginning of the descent when the cutoff frequency f_c takes a very small value. To remedy this, at a given cutoff frequency, we let the algorithm converge till it cannot take a further step in the descent direction, and then increase the cutoff frequency to $f_c + \Delta f$. The value of Δf
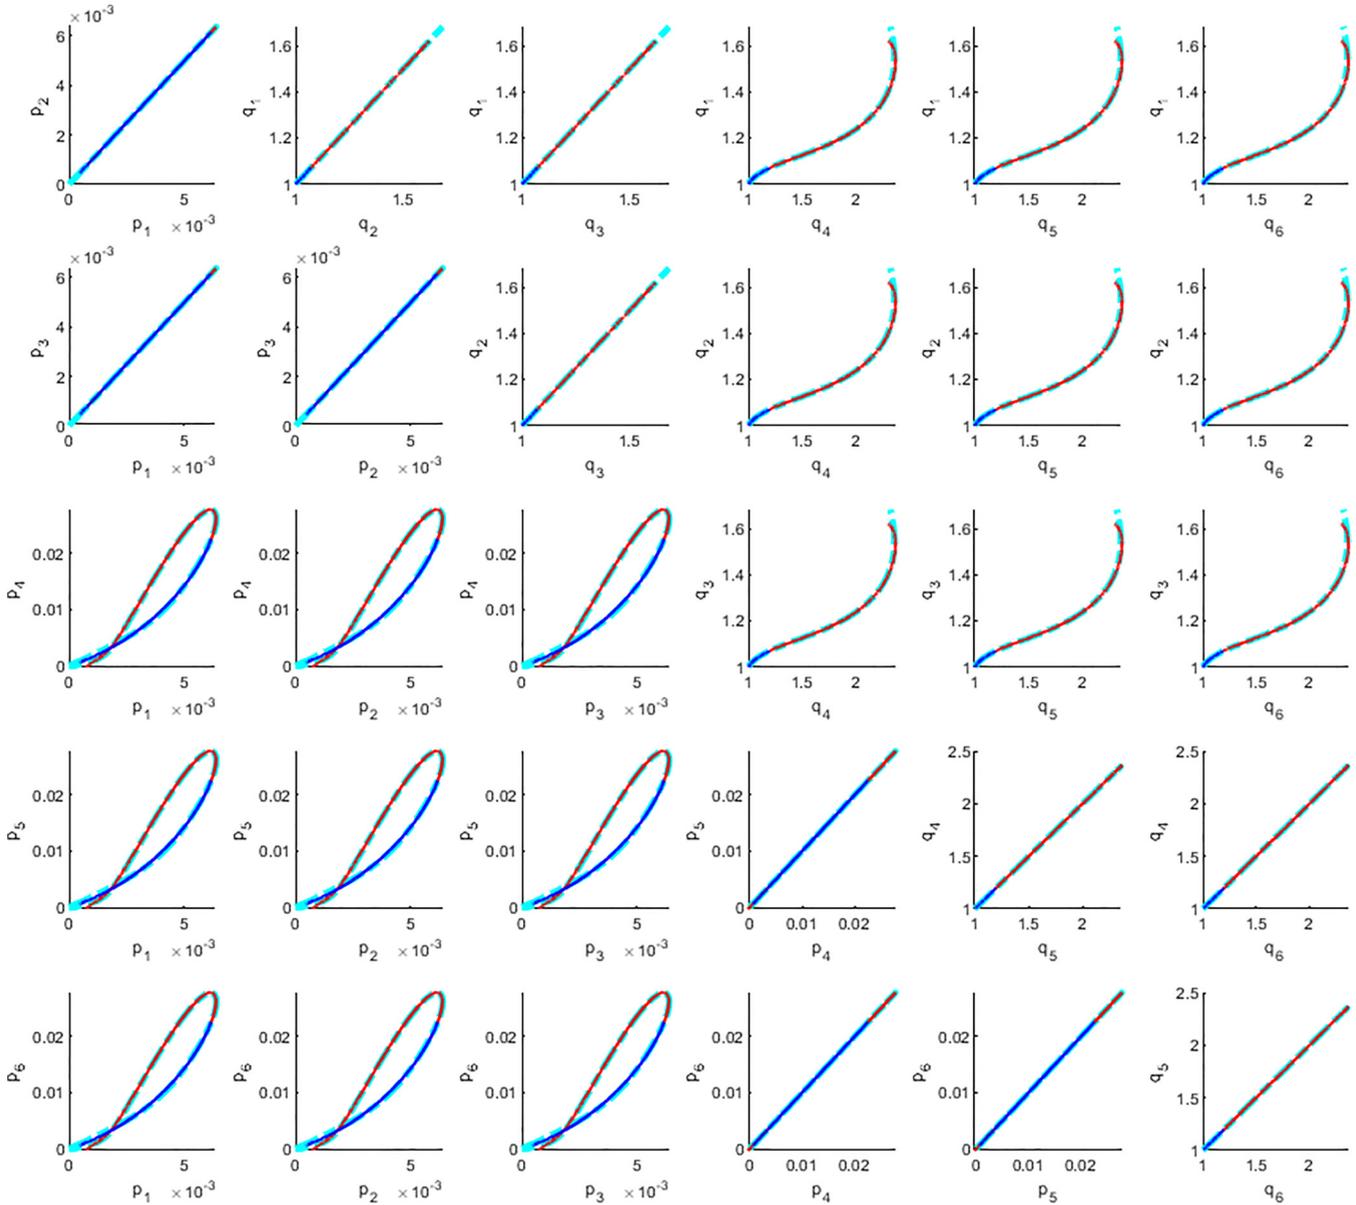
FIG. 21. Integrating Hamilton's equations of motion against AFGD algorithm output for the 6D Schlögl model after 300 iterations. The convention is the same as used in Fig. 14, and is omitted for neatness.

can be chosen by experimentation and we choose it to be 0.1 for our implementation.

We also define a maximum cutoff frequency f_{Max} as a considerable fraction of the Nyquist frequency. However, it must be noted that, in practice, the algorithm will stop descending at a much lower cutoff frequency than f_{Max}. In other words, for no step size will the filtered gradient yield a smaller value of the action functional than its current value. A reason for this might be that the gradient is dominated by noise which is being allowed to go through the pass-band. This is when, for descending further, we employ the *annealing* subroutine.

For an illustration of how the cutoff frequency updates with with iterations, see the fifth panel in Figs. 10, 19, and 20. Notice that in all of these, there are a few iterations where the cutoff frequency remains the same before increasing to a slightly higher value. As explained earlier, during these

iterations the algorithm takes a nonzero step size and the value of the action functional steadily decreases.

b. Annealing or increasing sample points

At a given iteration the curve only has a finite number of points, say N. However, in principle, the optimal trajectory that minimizes the action functional is a continuous function, i.e., we need the limit $N \rightarrow \infty$ to accurately represent it. We get around this problem by first descending with a small value for N until the algorithm converges, and then updating the number of points in the trajectory to $N + \Delta N$.

To illustrate the importance of this subroutine, we consider an application of the AFGD algorithm to the Selkov model. In Fig. 20 it can be seen that the algorithm converges to some trajectory that exhibits nondifferentiability in the momentum coordinates. We know that the true optimal path must be a

smooth function, and thus the algorithm has not converged to the true solution. To remedy this, we take the converged trajectory and resample in the time-uniform sampling, defined in Appendix D 5, with 4000 points. As is evident from Fig. 10 and the discussion in Sec. IV, this indeed takes the algorithm toward the optimal trajectory.

In principle, one must employ the annealing subroutine infinitely many times, since the true optimal trajectory is a continuous function. In practice however, since we can monitor the progress of the algorithm by integrating Hamilton's EoM, as we explain below, after a desired accuracy is reached one can terminate the AFGD algorithm and, if needed, use the shooting-method [10,39].

8. Integrating Hamilton's equations of motion and convergence criteria

The objective of the AFGD algorithm is to converge at the optimal curve constrained at the end points, such that the lifted trajectory is a solution to the Hamilton-Jacobi PDE in Eq. (13). In Appendix C 3, we prove that any solution of Hamilton-Jacobi must also satisfy Hamilton's equations of motion (EoM), thus the optimal lifted trajectory must also satisfy them. Since Hamilton's EoM is a system of coupled ODEs, it is easier to find their solution starting from an initial condition as opposed to solving HJ equations. We will now use this property of the optimal trajectory to define a "distance" from the true solution at a given iteration and a convergence criterion, as explained below.

At a given iteration I, let us denote the discrete phase-space trajectory as $\tilde{\gamma}$ and parametrize it with discrete index $m \in [1, N]$. Now, integrate Hamilton's EoM forward and backward starting from each point m until the configuration space distance from the saddle and stable fixed points, respectively, starts diverging, having passed through its point of closest approach. We take the minimum Euclidean distance in configuration space near the stable and saddle points, and add them to obtain a $\Delta(m)$ for each point m. Finally, we find the minimum over all m to assign a distance of the phase-space trajectory, denoted by $\Delta^I = \min_m \Delta(m)$, and use it as a "measure" for distance from optimality.

Note that since the optimal curve passes through both the saddle and stable fixed point, the distance Δ of the optimal trajectory must be zero by definition. In practice, however, one can define a Δ_{thresh} such that when $\Delta \leqslant \Delta_{\text{thresh}}$, we will declare the algorithm to have converged.

For an illustration of how this measure changes with iterations for the Selkov model, see the bottom right panel in Figs. 20 and 19 or last three panels in Fig. 10. To see how Δ changes during the descent of the 2-Schlögl model, see Fig. 14. In some of these figures, it can be seen that Δ continues to decrease even though the step size is 0. The reason for this is the additional time-uniform filtering on the curve each time the cutoff frequency is updated. Since distance from Hamilton's EoM is a theoretically rigorous measure of convergence, we can use this to convince ourselves of convergence for higher-dimensional models, where alternative verification methods such as the shooting method or Gillespie simulations can be rather expensive. For an example of the application of this measure to a 6D model, see Fig. 21.

[1] C. Hartmann and C. Schütte, Efficient rare event simulation by optimal nonequilibrium forcing, J. Stat. Mech. (2012) P11004.

[2] Y. Cao and J. Liang, Adaptively biased sequential importance sampling for rare events in reaction networks with comparison to exact solutions from finite buffer dCME method, J. Chem. Phys. **139**, 07B605_1 (2013).

[3] G. Biondini, An introduction to rare event simulation and importance sampling, in *Handbook of Statistics* (Elsevier, Amsterdam, 2015), Vol. 33, pp. 29–68.

[4] R. Benzi, G. Parisi, A. Sutera, and A. Vulpiani, A theory of stochastic resonance in climatic change, SIAM J. Appl. Math. **43**, 565 (1983).

[5] F. J. Dyson, A model for the origin of life, J. Mol. Evol. **18**, 344 (1982).

[6] E. Smith and S. Krishnamurthy, *Symmetry and Collective Fluctuations in Evolutionary Games* (IOP Publishing, Bristol, UK, 2015).

[7] E. Smith and H. J. Morowitz, *The Origin and Nature of Life on Earth: The Emergence of the Fourth Geosphere* (Cambridge University Press, Cambridge, UK, 2016).

[8] P. Y. Yu and G. Craciun, Mathematical analysis of chemical reaction systems, Isr. J. Chem. **58**, 733 (2018).

[9] B. J. Berne, G. Ciccotti, and D. F. Coker, *Classical and Quantum Dynamics in Condensed Phase Simulations: Proceedings of the International School of Physics* (World Scientific, Singapore, 1998).

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

[11] E. Weinan, W. Ren, and E. Vanden-Eijnden, Minimum action method for the study of rare events, Commun. Pure Appl. Math. **57**, 637 (2004).

[12] F. Bouchet, J. Rolland, and E. Simonnet, Rare Event Algorithm Links Transitions in Turbulent Flows with Activated Nucleations, Phys. Rev. Lett. **122**, 074502 (2019).

[13] P. Gagrani, AFGD-for-CRN-escapes (2022), https://github.com/praful12/Descender-for-CRN-escapes.

[14] H. Gang, Stationary solution of master equations in the large-system-size limit, Phys. Rev. A **36**, 5782 (1987).

[15] M. Snarski, Hamilton-Jacobi-Bellman equations for chemical reaction networks, arXiv:2108.00514.

[16] E. Smith, Intrinsic and extrinsic thermodynamics for stochastic population processes with multilevel large-deviation structure, Entropy **22**, 1137 (2020).

[17] G. L. Eyink, Action principle in nonequilibrium statistical dynamics, Phys. Rev. E **54**, 3419 (1996).

[18] M. Doi, Second quantization representation for classical many-particle system, J. Phys. A: Math. Gen. **9**, 1465 (1976).

[19] J. C. Baez and J. Biamonte, Quantum techniques for stochastic mechanics, arXiv:1209.3632.

[20] E. Smith, The information geometry of two-field functional integrals, Information Geometry 1 (2022).

[21] V. I. Arnol'd, *Mathematical Methods of Classical Mechanics* (Springer Science & Business Media, Berlin, 2013), Vol. 60.

[22] K. Kirsten and A. J. McKane, Functional determinants by contour integration methods, Ann. Phys. **308**, 502 (2003).

[23] S. Coleman, *Aspects of Symmetry: Selected Erice Lectures* (Cambridge University Press, Cambridge, UK, 1988).

[24] S. K. Manikandan and S. Krishnamurthy, Asymptotics of work distributions in a stochastically driven system, Eur. Phys. J. B **90**, 1 (2017).

[25] H. Goldstein, C. Poole, and J. Safko, Classical mechanics (American Association of Physics Teachers, 2002).

[26] M. Leok and J. Zhang, Connecting information geometry and geometric mechanics, Entropy **19**, 518 (2017).

[27] D. F. Anderson, G. Craciun, M. Gopalkrishnan, and C. Wiuf, Lyapunov functions, stationary distributions, and nonequilibrium potential for reaction networks, Bull. Math. Biol. **77**, 1744 (2015).

[28] D. T. Gillespie, A rigorous derivation of the chemical master equation, Physica A **188**, 404 (1992).

[29] M. Feinberg and F. J. M. Horn, Dynamics of open chemical systems and the algebraic structure of the underlying reaction network, Chem. Eng. Sci. **29**, 775 (1974).

[30] M. Feinberg, *Foundations of Chemical Reaction Network Theory* (Springer, Berlin, 2019).

[31] E. Smith and S. Krishnamurthy, Flows, scaling, and the control of moment hierarchies for stochastic chemical reaction networks, Phys. Rev. E **96**, 062102 (2017).

[32] E. Smith and S. Krishnamurthy, Eikonal solutions for moment hierarchies of chemical reaction networks in the limits of large particle number, J. Phys. A: Math. Theor. **54**, 185002 (2021).

[33] D. Langary and Z. Nikoloski, Inference of chemical reaction networks based on concentration profiles using an optimization framework, Chaos **29**, 113121 (2019).

[34] J. W. Milnor, M. Spivak, R. Wells, and R. Wells, *Morse Theory*, Number 51 (Princeton University Press, Princeton, NJ, 1963).

[35] D. J. Bates, J. D. Hauenstein, A. J. Sommese, and C. W. Wampler, Bertini: Software for numerical algebraic geometry, available at bertini.nd.edu with permanent doi: 10.7274/R0H41PB5.

[36] S. Wright, J. Nocedal *et al.*, in *Numerical Optimization* (Springer Science, Berlin, 1999), Vol. 35, p. 7.

[37] E. Hairer, M. Hochbruck, A. Iserles, and C. Lubich, Geometric numerical integration, Oberwolfach Rep. **3**, 805 (2006).

[38] R. I. McLachlan, G. R. W. Quispel, and P. S. P. Tse, Linearization-preserving self-adjoint and symplectic integrators, BIT Numer. Math. **49**, 177 (2009).

[39] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing* (Cambridge University Press, Cambridge, UK, 2007).

[40] V. Duruisseaux, J. Schmitt, and M. Leok, Adaptive Hamiltonian variational integrators and applications to symplectic accelerated optimization, SIAM J. Sci. Comput. **43**, A2949 (2021).

[41] E. E. Sel'Kov, Self-oscillations in glycolysis 1. A simple kinetic model, Eur. J. Biochem. **4**, 79 (1968).

[42] F. Schlögl, Chemical reaction models for nonequilibrium phase transitions, Z. Phys. **253**, 147 (1972).

[43] A. Lazarescu, T. Cossetto, G. Falasco, and M. Esposito, Large deviations and dynamical phase transitions in stochastic chemical networks, J. Chem. Phys. **151**, 064117 (2019).

[44] D. T. Gillespie *et al.*, Stochastic simulation of chemical kinetics, Annu. Rev. Phys. Chem. **58**, 35 (2007).

[45] L. Smolin, Galactic disks as reaction-diffusion systems, arXiv:astro-ph/9612033.

[46] A. M. Turing, The chemical basis of morphogenesis, Bull. Math. Biol. **52**, 153 (1990).

[47] E. Smith, Beyond fitness: Selection and information flow through the constructive steps in lifecycles, *bioRxiv* (2021).

[48] M. Shubik and E. Smith, *The Guidance of an Enterprise Economy* (MIT Press, Cambridge, MA, 2016).

[49] S. U. Pillai, T. Suel, and S. Cha, The perron-frobenius theorem: some of its applications, IEEE Signal Process. Mag. **22**, 62 (2005).

[50] D. F. Anderson, G. Craciun, and T. G. Kurtz, Product-form stationary distributions for deficiency zero chemical reaction networks, Bull. Math. Biol. **72**, 1947 (2010).

[51] S. Krishnamurthy and E. Smith, Solving moment hierarchies for chemical reaction networks, J. Phys. A: Math. Theor. **50**, 425002 (2017).

[52] F. Horn and R. Jackson, General mass action kinetics, Arch. Ration. Mech. Anal. **47**, 81 (1972).

[53] D. Manzano, A short introduction to the Lindblad master equation, AIP Adv. **10**, 025106 (2020).

[54] L. P. Kadanoff, *Statistical Physics: Statics, Dynamics, and Renormalization* (World Scientific, Singapore, 2000).

[55] M. A. Nielsen and I. Chuang, *Quantum computation and quantum information* (American Association of Physics Teachers, 2002).

[56] A. V. Oppenheim, A. S. Willsky, S. H. Nawab, G. M. Hernández *et al.*, *Signals & Systems* (Pearson Educación, 1997).

[57] A. Altland and B. D. Simons, *Condensed Matter Field Theory* (Cambridge University Press, Cambridge, UK, 2010).

[58] M. E. Peskin, *An Introduction to Quantum Field Theory* (CRC Press, Boca Raton, FL, 2018).

[59] E. Smith, Large-deviation principles, stochastic effective actions, path entropies, and the structure and meaning of thermodynamic descriptions, Rep. Prog. Phys. **74**, 046601 (2011).

[60] R. Courant and D. Hilbert, *Methods of Mathematical Physics: Partial Differential Equations* (John Wiley & Sons, New York, NY, 2008).

[61] A. Kamenev, *Field Theory of Non-equilibrium Systems* (Cambridge University Press, Cambridge, UK, 2011).

[62] H. Touchette, Legendre-Fenchel transforms in a nutshell (2005), http://www.maths.qmul.ac.uk/.