MULTISCALE MODEL. SIMUL.
Vol. 19, No. 3, pp. 1310–1332

APPROXIMATE OPTIMAL CONTROLS VIA INSTANTON EXPANSION FOR LOW TEMPERATURE FREE ENERGY COMPUTATION*

GRÉGOIRE FERRÉ† AND TOBIAS GRAFKE‡

Abstract. The computation of free energies is a common issue in statistical physics. A natural technique to compute such high-dimensional integrals is to resort to Monte Carlo simulations. However, these techniques generally suffer from a high variance in the low temperature regime, because the expectation is often dominated by high values corresponding to rare system trajectories. A standard way to reduce the variance of the estimator is to modify the drift of the dynamics with a control enhancing the probability of rare events, leading to so-called importance sampling estimators. In theory, the optimal control leads to a zero-variance estimator; it is, however, defined implicitly and computing it is of the same difficulty as the original problem. We propose here a general strategy to build approximate optimal controls in the small temperature limit for diffusion processes, with the first goal to reduce the variance of free energy Monte Carlo estimators. Our construction builds upon low noise asymptotics by expanding the optimal control around the instanton, which is the path describing most likely fluctuations at low temperature. This technique not only helps reducing variance, but it is also interesting as a theoretical tool since it differs from usual small temperature expansions (WKB ansatz). As a complementary consequence of our expansion, we provide a perturbative formula for computing the free energy in the small temperature regime, which refines the now standard Freidlin–Wentzell asymptotics. We compute this expansion explicitly for lower orders, and explain how our strategy can be extended to an arbitrary order of accuracy. We support our findings with illustrative numerical examples.

Key words. free energy, low temperature, large deviations, optimal control, Monte Carlo simulation, variance reduction

AMS subject classifications. 60F10, 82M31, 65C05, 49M99

DOI. 10.1137/20M1385809

1. Introduction.

This work is concerned with the computation of free energy-like quantities arising in statistical physics, for diffusion processes in the low temperature and finite time regime [16, 11, 18, 59]. Although such quantities are defined by integrals, the typical high dimensionality of the problem makes numerical integration impossible, so that one generally resorts to Monte Carlo simulation for numerical estimations. However, naive Monte Carlo methods often fail to provide accurate results because of the high variance of standard estimators. This situation typically arises because the observable of interest is dominated by large values along rare trajectories [6].

There are, in general, two ways for reducing the variance of naive Monte Carlo estimators when computing free energies. One is to introduce a bias in the dynamics, so that rare trajectories become more likely under the new dynamics [16, 53, 36, 47]—a strategy sometimes referred to as tilting. We know at a theoretical level that there exists a control, called optimal, which provides a zero-variance estimator. However,

*Received by the editors December 14, 2020; accepted for publication (in revised form) May 10, 2021; published electronically August 20, 2021.
https://doi.org/10.1137/20M1385809
Funding: This work was funded by the Labex Bézout ANR-10-LABX-58-01 as well as EPSRC projects EP/T011866/1 and EP/V013319/1.
†Université Paris-Est, CERMICS (ENPC), Inria (MATHERIALS), F-77455, Marne-la-Vallée, France (gregoire.ferre@ponts.org).
‡Warwick Mathematics Institute, University of Warwick, Coventry, CV4 7AL, UK (T.Grafke@warwick.ac.uk.

for high-dimensional systems, it is hopeless to compute this optimal control to a high degree of accuracy, and poor approximations may deteriorate the quality of the estimator. It is therefore an important and challenging problem to estimate as accurately as possible, and at a reasonable computational cost, the zero-variance control.

Another strategy is to resort to population dynamics [32, 10, 9, 8, 46, 5], another instance of importance sampling. The idea here is to run a series of systems in parallel, and to select the ones that realize the rare event dominating the expectation defining the free energy. There are various possibilities to design a selection mechanism, some provably behaving better than others [55, 1, 13]. However, it is a known fact that, in high-dimension and at low temperature, the number of replicas needed for performing accurate computations becomes very large [49, 50]. Of course, it is also possible to combine the two approaches; see [49] for an example of application to long time large deviations computations.

We focus here on the construction of approximate optimal controls for diffusion processes in the low temperature regime. Since different equivalent expressions are available for the optimal control (for instance, through stochastics, partial differential equations (PDEs), or variational representations), many approximation techniques have been developed, including cross-entropy methods [61], milestoning [36], Isaacs equation [14], martingale based techniques [48], model reduction [37], and forward-backward stochastic differential equations [42], or more recently machine learning based algorithms [34, 15, 35, 52]. We anticipate already here that, when an approximation of the optimal control is available, it is still not obvious that the resulting estimator should actually decrease the variance (see [56, 25, 2, 33] and references therein for more insight on this subtle issue).

The first goal of this paper is to provide a simple way to construct approximate controls that are well suited for variance reduction of free energy Monte Carlo computations in the small temperature regime. We rely for this on low temperature reaction paths (instantons) [16, 17, 27, 30] by building a time-inhomogeneous Taylor expansion around such reaction paths. This is quite different from the standard Wentzell–Kramers–Brillouin (WKB) (or Freidlin–Wentzell–Graham) small temperature expansion [31, 24, 3, 4, 43] where series run in the small temperature parameter and are defined through solutions to PDEs. From a more mathematical viewpoint, we propose an expansion of a finite noise Hamilton–Jacobi–Bellman (HJB) equation around the solution to the characteristic equation of the associated noiseless partial differential equation (PDE) (see [20, Chapter 3.2]), which is not a standard procedure to the best of our knowledge. With our technique, we manage to build offline (i.e., involving only computations that can be done once before starting the sampling) an approximate control that behaves well at low temperature. This should be put in contrast with techniques that build a precise estimate of the control by requiring costly on-the-fly updates [59] or solving a PDE—see, for instance, the interesting Isaacs subsolution approach [14].

A second output of our work is a perturbative formula for the free energy at low temperature. Using the optimal control expansion, we compute correction terms to the Freidlin–Wentzell zero-order asymptotics to the free energy, which can be estimated without resorting to Monte Carlo simulation. In general, we believe the expansion we propose is an interesting object to understand more precisely from a mathematical standpoint, in particular in view of the theory of viscosity solutions for Hamilton–Jacobi equations [7, 20, 23, 53].

This paper is organized as follows. Section 2 presents our problem (section 2.1) and recalls some well-known facts about zero variance estimators (section 2.2) and low temperature reaction paths (section 2.3). We next turn to the main results of the paper, by first presenting our approximation of the optimal control in section 3.1 and then the resulting perturbative formula for the free energy in section 3.2. We conclude in section 4 with some numerical applications illustrating our results. We finish with a short discussion, pointing out limitations of our technique, and indicating directions to address them.

2. Optimal control and low temperature limit. This section presents the overall setting of the work, and recalls some well-known facts about optimal control and low temperature asymptotics of exponential expectations, which we use in section 3 for our approximation procedure.

2.1. Free energy computation. We consider the computation of integrals of exponential quantities for which numerical integration is impossible and Monte Carlo estimators typically have a large variance. Concretely, for fixed time $T > 0$ and initial condition $x_0 \in \mathbb{R}^d$, we consider

$$A_\varepsilon = \mathbb{E}_{x_0}\left[e^{\frac{1}{\varepsilon}f(X_T^\varepsilon)}\right]. \tag{2.1}$$

Here, $f : \mathbb{R}^d \to \mathbb{R}$ is a smooth function and $(X_t^\varepsilon)_{t\geq 0}$ is solution to the following stochastic differential equation in $\mathbb{R}^d$ (with d a positive integer standing for the physical dimension)

$$dX_t^\varepsilon = b(X_t^\varepsilon)\,dt + \sqrt{\varepsilon}\sigma\,dB_t, \tag{2.2}$$

where $(B_t)_{t\geq 0}$ is an m-dimensional Brownian motion, the function $b : \mathbb{R}^d \to \mathbb{R}^d$ is smooth, and $\sigma \in \mathbb{R}^{d\times m}$ is such that the diffusion matrix $D = \sigma\sigma^{\mathrm{T}} \in \mathbb{R}^{d\times d}$ is positive definite (σ^{T} stands for the transpose of the matrix σ). In (2.1), we denote by $\mathbb{E}_{x_0}$ the expectation with respect to all trajectories solving (2.2) and starting at the initial point $x_0 \in \mathbb{R}^d$ at time $t = 0$. Note that we could also consider a time-dependent function b as well as a time-position dependent diffusion matrix σ without additional difficulty, but restrict ourselves to this setting for notational simplicity. The generator of the dynamics (2.2) reads

$$\mathcal{L} = b\cdot\nabla + \varepsilon\frac{\sigma\sigma^{\mathrm{T}}}{2} : \nabla^2, \tag{2.3}$$

where $\cdot$ is the scalar product in $\mathbb{R}^d$. The notation ∇^2 stands for the $\mathbb{R}^{d\times d}$-valued Hessian operator, while for two matrices $A, B \in \mathbb{R}^{d\times d}$ we write $A : B = \mathrm{Tr}(A^{\mathrm{T}}B)$. The differential operators ∇, ∇^2, and $\mathcal{L}$ can be defined on smooth compactly supported functions, and we assume in what follows that the parameters of the model allow one to define (2.1) as a finite quantity for all $\varepsilon > 0$; see, in particular, [24] for technical considerations.

A motivation for studying (2.1) is the computation of the free energy

$$Z_\varepsilon = \varepsilon \log A_\varepsilon \tag{2.4}$$

in the small temperature regime.[1] It is known by large deviations arguments that, under mild assumptions, it holds [24, 12] that

$$Z_\varepsilon \xrightarrow[\varepsilon\to 0]{} Z^0 \tag{2.5}$$

for some finite value Z^0; see section 2.3 below. The numerical computation of Z^0 is one motivation for estimating (2.1) when $\varepsilon \ll 1$. In a large deviations perspective, it is also useful to compute Z^0 for numerically estimating the rate function associated to the path measure of $(X_t^\varepsilon)_{t\in[0,T]}$, which is related to Z^0 through a Legendre–Fenchel transform. We refer to [58, 21] for numerical examples in the related infinite time context.

In the regime of small temperature, the expectation in (2.1) is often dominated by very large values realized over rare trajectories, which leads to large variance Monte Carlo estimators. However, we know that the dynamics (2.2) can be controlled to be turned into a zero-variance estimator of (2.1), as we recall now.

2.2. Optimal tilting on path space. We now present the modification of (2.2) leading to a zero-variance estimator of (2.1). These computations are standard, provided technical conditions are met; see, for instance, [23, 53]. In this procedure, we consider the *tilted process* $(\widetilde{X}_t^\varepsilon)_{t\geq 0}$ solution to

$$d\widetilde{X}_t^\varepsilon = b(\widetilde{X}_t^\varepsilon)\,dt + D\nabla g(t,\widetilde{X}_t^\varepsilon)\,dt + \sqrt{\varepsilon}\sigma\,dB_t, \tag{2.6}$$

where $g : \mathbb{R}_+ \times \mathbb{R}^d \to \mathbb{R}$ is an arbitrary smooth function, and we call ∇g the *control*. We restrict ourselves to gradient controls since, as shown below, the optimal control is indeed gradient.

First, we introduce the Girsanov weight $\alpha : [0,T]\times\mathbb{R}^d \to \mathbb{R}$ associated with g, namely

$$\forall\, t \geq 0, \quad \forall\, x \in \mathbb{R}^d, \quad \alpha(t,x) = \partial_t g(t,x) + \mathcal{L}g(t,x) + \frac{1}{2}|\sigma\nabla g|^2(t,x).$$

Next, we define the function $\psi_\varepsilon : [0,T]\times\mathbb{R}^d \to \mathbb{R}_+$ as

$$\psi_\varepsilon(t,x) = \mathbb{E}_{t,x}\left[e^{\frac{1}{\varepsilon}f(X_T^\varepsilon)}\right] \tag{2.7}$$

and

$$g_\varepsilon(t,x) = \varepsilon \log\, \psi_\varepsilon(t,x). \tag{2.8}$$

[1] The term free energy is often associated with long time problems through the quantity $\lim_{T\to+\infty} \frac{1}{T}\log \mathbb{E}\big[e^{\int_0^T f(X_s^\varepsilon)\,ds}\big]$. Here, we use the terminology associated with small temperature problems, like in [36]. Note that we could also consider expectations involving a random stopping time τ, such as $\varepsilon \log \mathbb{E}_{x_0}\big[e^{\frac{1}{\varepsilon}f(X_\tau^\varepsilon)}\big]$, or finite time-integrated quantities like $\varepsilon \log \mathbb{E}_{x_0}\big[e^{\frac{1}{\varepsilon}\int_0^T f(X_s^\varepsilon)\,ds}\big]$, for a finite integration time $T > 0$, when $\varepsilon \to 0$. These cases can be treated by appropriately modifying the computations performed in Appendix B. In this paper, we present our method by considering (2.1) and leave the modifications needed in other cases to the interested reader. Combining long time and small temperature asymptotics on the other hand is a difficult problem; see, for instance, [51] for interesting insights.

In (2.7), $\mathbb{E}_{t,x}$ refers to the expectation with respect to all realizations of the dynamics (2.2) started at time t from position x. Under technical conditions, g_ε is well-defined as a solution (at least in a weak sense [20]) to the following HJB equation (see Appendix A):

$$\begin{cases} \partial_t g_\varepsilon + \mathcal{L} g_\varepsilon + \dfrac{1}{2} \left| \sigma \nabla g_\varepsilon \right|^2 = 0, \\ \qquad\qquad g_\varepsilon(T,x) = f(x) \quad \forall\, x \in \mathbb{R}^d. \end{cases} \tag{2.9}$$

We assume in what follows that g_ε actually exists as a unique smooth solution of (2.9) with the probabilistic representation (2.7)–(2.8), and refer to section 3.1 for more details on this assumption.

Then, by setting $g = g_\varepsilon$ in (2.6), the estimator

$$A_\varepsilon = e^{g_\varepsilon(0,x_0)} \mathbb{E}_{x_0} \left[\exp\left(\frac{1}{\varepsilon} \big[f(\widetilde{X}_T^\varepsilon) - g_\varepsilon(T, \widetilde{X}_T^\varepsilon) \big] + \frac{1}{\varepsilon} \int_0^T \alpha(t, \widetilde{X}_t^\varepsilon)\, dt \right) \right] \tag{2.10}$$

has zero variance. Namely,

$$A_\varepsilon = \psi_\varepsilon(0, x_0). \tag{2.11}$$

This result is a consequence of the Feynman–Kac formula and the Girsanov theorem; see Appendix A for a more detailed argument. A consequence of (2.10) is that (2.1) can be estimated with a zero-variance (i.e., deterministic) estimator provided (2.8) is known.

In general, the Monte Carlo estimator built on (2.10) by drawing independent trajectories distributed according to (2.6) cannot be used as such for numerical applications, because estimating (2.8) and its gradient for all $t \geq 0$ and $x \in \mathbb{R}^d$ is still more difficult than solving the initial problem of estimating (2.1). However, this result serves as a guide to design approximate controls that are easier to compute while still reducing the variance of Monte Carlo estimators of (2.1). We will present in section 3 an original strategy to build such approximate controls behaving well in the small ε regime from the low temperature asymptotics provided by transition path theory, which is the main contribution of this work. For this, we first need to recall the definition of the transition path in our context, which is the purpose of the next section.

2.3. Low temperature regime and reaction path. Even though the control g_ε is difficult to estimate in practice, we can nevertheless have access to an *instanton*, or reaction or transition path, which stands for the zero temperature most likely path of fluctuation for the dynamics. In the small noise limit, we know by the Freidlin–Wentzell theory [24, section 3] that the trajectories of $(X_t^\varepsilon)_{t\in[0,T]}$ dominating the expectation (2.1) concentrate exponentially fast on this path for the uniform norm under relatively mild conditions on the parameters of the problem. We only recall the most important features of the theory here, and refer to [30] and references therein for more details.

The instanton is a path $(\phi_t)_{t\in[0,T]}$ taking values in $\mathbb{R}^d$, assumed here to be smooth and uniquely defined (we shall discuss more this assumption in Remark 3.3 below). In order to provide an equation for this path, we also consider a conjugate variable $(\theta_t)_{t\in[0,T]}$, which can be thought of as a momentum. The reaction path $(\phi_t, \theta_t)_{t\in[0,T]}$

is then described by the following forward-backward system of equations:

$$\begin{cases} \dot{\phi}_t = b(\phi_t) + D\theta_t, & \phi_0 = x_0, \\ \dot{\theta}_t = -(\nabla b)^{\mathrm{T}}(\phi_t)\theta_t, & \theta_T = \nabla f(\phi_T). \end{cases} \tag{2.12}$$

Note that the initial condition x_0 of ϕ is the same as the one appearing in the definition (2.1) of the free energy. We insist on the fact that the instanton is defined by a deterministic system of equations, and that the reaction path $(\phi_t)_{t\in[0,T]}$ corresponds to a typical path whose final value of f dominates the expectation in (2.1). Moreover, we mention that the set of equations (2.12) is simply the characteristic system describing the noiseless limit of the HJB equation (2.9); see [20, Chapter 3.2].

Finally, the set of equations defining the reaction path provides a representation of the low temperature limit (2.5) of the free energy through

$$Z^0 = \lim_{\varepsilon\to 0} Z_\varepsilon = f(\phi_T) - \frac{1}{2}\int_0^T \theta_t \cdot D\theta_t \, dt. \tag{2.13}$$

Thus, Z^0 can be interpreted in an optimal control sense as the maximal value of f that can be reached under a quadratic penalization of the momentum [36]. Not surprisingly, this kind of asymptotic is obtained via the Girsanov theorem through computations similar to that of section 2.2.

We now have all the tools to present the main contributions of the paper, which are as follows: (i) an approximation of the optimal control g_ε around the reaction path and (ii) a resulting expansion of Z_ε for small values of ε.

3. Low temperature approximation of the optimal bias. We now present our main results. First, we build an approximation of the optimal control around the instanton in section 3.1. We next deduce in section 3.2 a perturbative formula for the free energy Z_ε.

3.1. Expansion around the instanton. In order to present our expansion, we first recall that the zero-variance control is the solution to the HJB equation (2.9) which reads in full form

$$\begin{cases} \partial_t g_\varepsilon + b\cdot\nabla g_\varepsilon + \varepsilon\dfrac{D}{2} : \nabla^2 g_\varepsilon + \dfrac{1}{2}\left|\sigma\nabla g_\varepsilon\right|^2 = 0, \\ \qquad\qquad\qquad\qquad g_\varepsilon(T,x) = f(x) \quad \forall\, x\in\mathbb{R}^d. \end{cases} \tag{3.1}$$

In the zero-temperature limit $\varepsilon\to 0$, the PDE above becomes

$$\begin{cases} \partial_t g^0 + b\cdot\nabla g^0 + \dfrac{1}{2}\left|\sigma\nabla g^0\right|^2 = 0, \\ \qquad\qquad\qquad g^0(T,x) = f(x) \quad \forall\, x\in\mathbb{R}^d. \end{cases} \tag{3.2}$$

We assume in what follows that (3.1) possesses a unique smooth solution on $[0,T)\times\mathbb{R}^d$, which is typically the case under reasonable assumptions by parabolic regularity (see, for instance, [23, section 4, Theorem 4.1]). Moreover, we also assume that (3.2) has a unique smooth solution. This is a more stringent assumption for which it is difficult to provide general conditions of application. However, we know by the method of characteristics that this assertion is valid when the final time T is small enough [20, section 3.2, Theorem 2]. We place ourselves in this setting in this paper, and refer to Remark 3.3 below for further comments on these assumptions. Note also that it typically holds in a weak sense that $\lim_{\varepsilon\to 0} g_\varepsilon = g^0$.

Solving the characteristics system for (3.2) actually relies [20, 30] on plugging the ansatz

$$g^0(t,x) = \theta_t \cdot (x - \phi_t) \tag{3.3}$$

into (3.2), which allows one to derive the couple of equations (2.12) defining the instanton. The definition (3.3) is motivated by Lagrangian considerations in statistical physics [30], but is simply an application of the method of characteristics for first order nonlinear PDEs [20, Chapter 3.2].

The main idea of this paper is to consider (3.3) as the first term of a polynomial Taylor expansion around the instanton $(\phi_t)_{t\in[0,T]}$. This suggests going to next order by looking for a solution of (3.1) in the form

$$g^1(t,x) = \theta_t \cdot (x - \phi_t) + \frac{1}{2}(x - \phi_t) \cdot K_t(x - \phi_t), \tag{3.4}$$

where $(K_t)_{t\in[0,T]}$ is a $\mathbb{R}^{d\times d}$-valued process to be determined. In what follows, we call the ansatz (3.3) the zeroth order approximation and (3.4) the first order one, because the resulting controls ∇g^0 and ∇g^1 are of zeroth and first order in $x - \phi_t$, respectively (see Remark 3.1 below for expansions to arbitrary order). Although this is not an expansion in powers of ε, the temperature appears implicitly through the relation $\widetilde{X}^\varepsilon_t - \phi_t = \mathrm{O}(\sqrt{\varepsilon})$, which holds when the drift is chosen accordingly. Indeed, taking for example $g = g^0$ in (2.6) and assuming that $\widetilde{X}^\varepsilon_t = \Phi_t + \mathrm{O}(\sqrt{\varepsilon})$ for some path $(\Phi_t)_{t\in[0,T]}$, then (2.6) becomes

$$\dot\Phi_t + \mathrm{O}(\sqrt{\varepsilon}) = b(\Phi_t) + D\nabla g^0(t,\Phi_t) + \mathrm{O}(\sqrt{\varepsilon}).$$

Since $\nabla g^0(t,x) = \theta_t$, we observe that indeed $\Phi_t = \phi_t$ is the instanton. As a result, one should think of $x - \phi_t$ as a quantity of order $\sqrt{\varepsilon}$ along a trajectory $(\widetilde{X}^\varepsilon_t)_{t\in[0,T]}$ when the drift g is built from g^0.

We now derive heuristically the equation satisfied by $(K_t)_{t\in[0,T]}$ for g^1 to be an approximation of g_ε. For this, we note that (3.1) rewrites componentwise $\partial_t g_\varepsilon + b_k\partial_k g_\varepsilon + D_{jk}\partial_j g_\varepsilon \partial_k g_\varepsilon/2 + \varepsilon D_{jk}\partial_{jk} g_\varepsilon/2 = 0$, where we use Einstein's notation for summation over repeated indices. Taking the derivative twice with respect to indices $i \in \{1,\dots,d\}$ and $l \in \{1,\dots,d\}$ shows that

$$\begin{aligned}\partial_t\partial^2_{il} g_\varepsilon + \partial^2_{il} b_k \partial_k g_\varepsilon + \partial_i b_k \partial^2_{kl} g_\varepsilon + \partial_l b_k \partial^2_{ik} g_\varepsilon + b_k \partial^3_{ikl} g_\varepsilon \\ + D_{jk}\partial^2_{jl} g_\varepsilon \partial^2_{ik} g_\varepsilon + D_{jk}\partial_j g_\varepsilon \partial^3_{ikl} g_\varepsilon \varepsilon \frac{D_{jk}}{2} \partial_{ijkl} g_\varepsilon = 0.\end{aligned}$$

This can be written in vectorial form as

$$\begin{aligned}\partial_t \nabla^2 g_\varepsilon + \nabla^2 b \cdot \nabla g_\varepsilon + (\nabla b)^{\mathrm{T}} \nabla^2 g_\varepsilon + \nabla^2 g_\varepsilon \nabla b + (\nabla^2 g_\varepsilon)^{\mathrm{T}} D \nabla^2 g_\varepsilon \\ + b\nabla^3 g_\varepsilon + (\nabla^3 g_\varepsilon)^{\mathrm{T}} D \nabla g_\varepsilon + \varepsilon \frac{D}{2} : \nabla^4 g_\varepsilon = 0,\end{aligned}$$

where the equation is evaluated at any (t,x).

Since we look for an evolution equation for $K_t = \nabla^2 g^1(t,\phi_t)$, we compute

$$\begin{aligned}\frac{d}{dt}\nabla^2 g_\varepsilon(t,\phi_t) &= \partial_t \nabla^2 g_\varepsilon + \dot\phi_t \cdot \nabla^3 g_\varepsilon \\ &= -(\nabla^2 b)^{\mathrm{T}} \nabla g_\varepsilon - (\nabla b)^{\mathrm{T}} \nabla^2 g_\varepsilon - \nabla^2 g_\varepsilon \nabla b - (\nabla^2 g_\varepsilon)^{\mathrm{T}} D \nabla^2 g_\varepsilon \\ &\quad - b\nabla^3 g_\varepsilon - (\nabla^3 g_\varepsilon)^{\mathrm{T}} D \nabla g_\varepsilon + \dot\phi_t \cdot \nabla^3 g_\varepsilon - \varepsilon \frac{D}{2} : \nabla^4 g_\varepsilon,\end{aligned} \tag{3.5}$$

where the right-hand side is evaluated at (t, ϕ_t). Considering the ansatz g^1 given by (3.4) to replace g_ε, we have $\nabla g^1(t,\phi_t) = \theta_t$, $\nabla^2 g^1(t,\phi_t) = K_t$, $\nabla^3 g^1(t,\phi_t) = 0$, and $\nabla^4 g^1(t,\phi_t) = 0$, which can be plugged into (3.5) to obtain the equation satisfied by $(K_t)_{t\in[0,T]}$. The final condition on K_T can be derived similarly by differentiating twice the terminal condition in (3.1). As a consequence, in order for (3.4) to approximate (3.1), $(K_t)_{t\in[0,T]}$ should be solution to

$$\begin{cases} \dot{K}_t + (\nabla b)^{\mathrm{T}} K_t + K_t^{\mathrm{T}} \nabla b + \nabla^2 b \cdot \theta_t + K_t^{\mathrm{T}} D K_t = 0, \\ K_T = \nabla^2 f(\phi_T), \end{cases} \tag{3.6}$$

where b and its derivatives are evaluated at ϕ_t. A precise derivation of (3.6) can be found in Appendix B via the Girsanov theorem (see, in particular, (B.3) and (B.5)). Note that a solution to (3.6) is symmetric.

In order to formalize that g^1 is indeed an approximation of g_ε, the solution to (3.1), we thus consider the first order approximation (3.4) where $(\phi_t, \theta_t)_{t\in[0,T]}$ is defined in (2.12) and $(K_t)_{t\in[0,T]}$ satisfies (3.6). We also define the following function of time:

$$\forall\, t \in [0,T], \quad Z_\varepsilon^1(t) = f(\phi_T) - \frac{1}{2}\int_t^T \theta_s \cdot D\theta_s \, ds + \frac{\varepsilon}{2}\int_t^T D : K_s \, ds. \tag{3.7}$$

Then, we show in Appendix B that, in the small ε limit, for any $t \in [0,T]$ and $x \in \mathbb{R}^d$, it holds that

$$g_\varepsilon(t,x) = g^1(t,x) + Z_\varepsilon^1(t) + \mathrm{o}(\varepsilon) + \varepsilon \mathrm{o}(x - \phi_t). \tag{3.8}$$

This formula shows that g^1 approximates the optimal control g_ε at small temperatures and around the instanton. A key ingredient of the proof is that, as noted above, for all time $t \geq 0$, it holds that $\widetilde{X}_t^\varepsilon - \phi_t = \mathrm{O}(\sqrt{\varepsilon})$ (see (B.2) in Appendix B), so the approximation is valid for the process (2.6) tilted by g^1.

Equation (3.6) is an instance of algebraic Riccati equation [45], which is an interesting feature compared to the more standard instanton presented in section 2.3. Riccati equations recurrently appear in optimal control theory [39], so it is no surprise to encounter such an equation in our approximation procedure (here a linearization) of the optimal control. The original feature, we believe, is the fact that all the objects in the approximation are centered around the zero-temperature instanton. Indeed, we insist on the fact that our construction (3.4) *is not a WKB ansatz* since it is not an expansion in the temperature parameter ε (see Remark 3.1 below). Actually, since the tilted process controlled by g^1 satisfies $\widetilde{X}_t^\varepsilon - \phi_t = \mathrm{O}(\sqrt{\varepsilon})$, this dependency in the temperature is hidden in the expansion around the instanton.

From a numerical perspective, we will use g^1 as an ansatz for the optimal control g_ε defined in (2.8). Note that most techniques relying on optimal control strategies strive to estimate $g_\varepsilon(t,x)$ and its gradient for all time t and position x, which is very difficult and computationally costly in practice [59]. Here, we can construct *offline* an approximation of this optimal control, which is a polynomial expansion whose coefficients depend on time only, which drastically reduces the computational cost of the procedure. However, since the construction relies on a small temperature expansion, we expect this approximation to reduce the variance only in a low temperature regime—a fact confirmed by the numerical simulations below.

Remark 3.1 (higher order expansion and relation to WKB ansatz). It is, of course, possible to push our method to an approximation of order $M > 1$ for g_ε

through

$$g^M(t,x) = \sum_{k=1}^{M+1} T_k(t) \odot \big(x - \phi_t\big)^{\otimes k}, \tag{3.9}$$

where $\otimes k$ stands for the kth order tensorization of a d-dimensional vector, $\odot$ the kth order contraction and, for each $k \geq 1$, T_k is a time dependent kth order tensor. Comparatively, an expansion in the temperature parameter ε, sometimes called WKB expansion, would read, up to order M,

$$g(t,x) = \sum_{k=0}^{M} (\sqrt{\varepsilon})^k u_k(t,x),$$

where each u_k is solution to a PDE [22]. Since $\widetilde{X}_t^\varepsilon - \phi_t = \mathrm{O}(\sqrt{\varepsilon})$, the expansion (3.9) looks like a WKB expansion in powers of $\sqrt{\varepsilon}$ around the tilted process although this parameter does not appear explicitly. Moreover, we do not need to solve any PDE since we work with ordinary differential equations at the process level. This allows faster numerical computations and the simple derivation of a perturbative formula for the free energy (2.4), as presented in section 3.2 below. In some sense, the expansion we propose can be thought of as a Taylor, polynomial expansion version of the standard WKB series.

With the notation (3.9), we have seen above that $T_1(t) = \theta_t$ and $T_2(t) = K_t$. In a Lagrangian perspective, we can interpret θ_t as a momentum, and thus the matrix K_t as an acceleration field. Moreover, in dimension $d = 1$, we can show going one order further in the computations of Appendix B that the third term $T_3(t) = Q_t$ is the solution to

$$\begin{cases} \dot{Q}_t + b'''\theta_t + 3b''K_t + 3b'Q_t + 6K_t DQ_t = 0, \\ Q_T = f'''(\phi_T), \end{cases}$$

where the derivatives of b are evaluated at the instanton ϕ_t. The next terms follow similarly by computing the next orders of the Taylor expansion.

We finally note that similar computations appear in [3]. However, the setting of this paper is different since the authors consider a drift b^λ depending on a free parameter λ unrelated to the temperature, and the expansions are realized with respect to this additional degree of freedom. To the best of our knowledge, expanding around the noiseless characteristic equation is a new technique.

3.2. Perturbative formula for the free energy at finite temperature. In the previous section, we focused on constructing an approximate optimal control for Monte Carlo importance sampling estimators. In addition to this result, we now deduce from (3.8) a perturbative formula for $(Z_\varepsilon)_{\varepsilon>0}$ for small values of ε. For this, it suffices to note that

$$Z_\varepsilon = g_\varepsilon(0, x_0).$$

Considering (3.8) for $t = 0$ and $x = x_0$ then leads to

$$Z_\varepsilon = Z_\varepsilon^1(0) + \mathrm{o}(\varepsilon) = f(\phi_T) - \frac{1}{2}\int_0^T \theta_t \cdot D\theta_t \, dt + \frac{\varepsilon}{2}\int_0^T D : K_t \, dt\,. \tag{3.10}$$

For brevity, we will denote $Z_\varepsilon^1(0)$ by Z_ε^1 in what follows. The above formula provides the first order correction to the zeroth order Freidlin–Wentzell asymptotics (2.13).

As mentioned in Remark 3.1, we could continue to construct higher order corrections and obtain a full expansion of the free energy at finite temperature through integrals of solutions to ordinary differential equations. Since we are more interested with the numerical side of this paper, we propose (3.10) as a way to numerically correct the Freidlin–Wentzell asymptotics (2.13) for small temperatures without resorting to Monte Carlo simulation.

Remark 3.2 (relation to prefactor analysis). The correction term in (3.10) reads

$$\int_0^T \mathrm{Tr}(DK_t)\,dt,$$

meaning that the correction to A_ε defined in (2.1) is

$$\exp\left(\frac{1}{2}\int_0^T \mathrm{Tr}(DK_t)\,dt\right).$$

Defining a matrix $G \in \mathbb{R}^{d\times d}$ via

$$\dot{G}_t = G_t DK_t\,,$$

we can apply Liouville's formula[2] to obtain

$$\exp\left(\frac{1}{2}\int_0^T \mathrm{Tr}(DK_t)\,dt\right) = \frac{\sqrt{\det|G_T|}}{\sqrt{\det|G_0|}}.$$

It is common to express the first correction to the Freidlin–Wentzell small temperature asymptotics as the determinant of a Hessian matrix. As a result, the perturbative formula (3.10) can be understood as a prefactor analysis, and the integral of the Riccati matrix as a continuous version of the determinant prefactor that arises for instance in the Eyring–Kramers formula (see, e.g., [4] and references therein, as well as [29, 57]). Following Remark 3.1, our methodology allows one to compute $Z_\varepsilon = Z_\varepsilon^M + \mathrm{o}(\varepsilon^M)$, where Z^M is a power series in ε up to order $M \in \mathbb{N}$, with coefficients defined as integrals of solutions to ordinary differential equations.

Remark 3.3 (multiple instantons). For now it is clear that our strategy relies on the well-definedness of the reaction path. As we said, this is nothing else than the characteristic solution to the noiseless HJB equation (3.2) associated with the optimal control. However, in many cases the characteristic is ill-defined, which provokes shocks and discontinuities in solutions to the HJB problem. This is why a theory of weak solutions has been developed, in order to provide a sense of solution in cases where a classical solution does not exist.

In these more complicated (yet easy to construct) situations [25, 59], it is not clear yet how to adapt our method. Depending on the problem one wishes to solve, it may be possible to content oneself with the "most important" instanton, that is, the one defining the Freidlin–Wentzell asymptotics. Otherwise, one may want to consider several characteristics and glue their resulting expansions together appropriately. We will not address this issue here, and thus our results, as shown, only apply to the situation where the characteristic system is well defined. Understanding how our methodology can be extended to situations where only a weak solution is available is an interesting open problem.

[2]For a given matrix-valued process $A : [0,T] \to \mathbb{R}^{d\times d}$, the matrix $\Psi(t) : [0,T] \to \mathbb{R}^{d\times d}$ solution to $\dot{\Psi}(t) = A(t)\Psi(t)$ satisfies $\det\Psi(T) = \det\Psi(0)\exp\big(\int_0^T \mathrm{Tr}A(t)\,dt\big)$.

Remark 3.4 (error analysis). Provided the problems raised in the above remark are addressed properly, controlling precisely the error terms in (3.8) is another interesting mathematical problem. We believe this can be tackled by more standard error analysis techniques [24, 47]. However, even if such error estimates were available, it is yet another problem to prove that the resulting importance sampling estimator built on (3.4) indeed reduces the variance for estimating (2.1). This is a subtle problem for which we refer to [56, 25, 2, 33] and references therein.

4. Numerical applications. In the following section, we demonstrate the usefulness of our approximation by performing numerical experiments on a number of example systems, comparing first the value of the free energy Z_ε estimated by Monte Carlo sampling to the Freidlin–Wentzell asymptotics Z^0 (which is constant) and the linear approximation Z^1_ε. Further, we compare the performance of a naive (unbiased) Monte Carlo estimator to the one using importance sampling with the approximate optimal biases g^0 and g^1.

For comparing Monte Carlo estimators, we use the *relative error*, which is the ratio of the standard deviation of our estimator over its average for a number of realizations. If, indeed, our approximation to the optimal bias is effective, heuristically we expect smaller relative error for higher order approximations to the optimal bias. Numerically measuring the relative error is therefore an experimental quantification of the variance reduction capabilities of our proposed estimators. Moreover, for all the numerical simulations we discretize the underlying SDE with a standard Euler–Maruyama scheme with time step $\Delta t > 0$, and neglect the error arising from this numerical quadrature [44].

4.1. One-dimensional Ornstein–Uhlenbeck process. The simplest situation is the one-dimensional case where the drift is given by $b(x) = -\gamma x$, with $\gamma > 0$. We further set $\sigma = 1$ and $f(x) = x$. This particular case corresponds to the Ornstein–Uhlenbeck process.

For the above choice, we know that the optimal control is actually equal to the conjugate momentum $(\theta_t)_{t\in[0,T]}$ from (2.12). Therefore, the zeroth order approximation (3.3) actually already provides the zero variance estimator described in section 2.2. Numerically, we therefore expect an estimator with variance close to zero. The first order correction term should not improve the results, so the matrix (scalar in this case) $(K_t)_{t\in[0,T]}$ should be zero for all times here.

The numerical experiment is performed with $x_0 = -1$, $T = 10$, and $\Delta t = 0.01$ for the numerical discretization, performing $N = 10^6$ experiments. The results are shown in Table 1, where we compare the naive unbiased estimator to the estimator biased with the instanton (zeroth order approximation) and the estimator biased to first order. While the relative error of the naive estimator blows up with decreasing ε, the relative error is zero in both the zeroth and first order estimators, implying that the first order estimator is already equivalent to the optimal zero variance bias. We also numerically observe that the Riccati matrix is indeed equal to zero (not shown).

4.2. Two-dimensional nonlinear nonequilibrium process. We now break detailed balance by considering a drift $b(x)$ that is not the gradient of a potential $V(x)$. In that case it is no longer true that the reaction path is merely a reverse relaxation trajectory driven by the potential level sets, and the reaction path itself must be computed by numerically solving the instanton equations (2.12), which is a well-established problem in the literature [17, 38, 26, 28]. Here, we use the algorithm from [30, section III.A]. Moreover, the optimal control is also no longer explicit like

TABLE 1
Relative error for Z_ε for the Ornstein–Uhlenbeck test case, comparing the naive Monte Carlo estimator with the zeroth and first order biased Monte Carlo estimators.

ε	Naive estimator	Zeroth order estimator	First order estimator
1	0.81	$2.73{\cdot}10^{-12}$	$3.15{\cdot}10^{-12}$
0.5	1.31	$6.12{\cdot}10^{-12}$	$1.05{\cdot}10^{-11}$
0.1	9.28	$5.04{\cdot}10^{-12}$	$3.08{\cdot}10^{-13}$
0.05	40.38	$1.03{\cdot}10^{-11}$	$3.74{\cdot}10^{-12}$
0.01	291.82	$8.71{\cdot}10^{-12}$	$1.2{\cdot}10^{-12}$

in the Ornstein–Uhlenbeck case.

As an example, we take the system

$$\begin{cases} dX_t = (Y_t^3 - X_t^3)\,dt + \sqrt{\varepsilon}\,dB_t^X, \\ dY_t = (-X_t^3 - Y_t^3)\,dt + \sqrt{\varepsilon}\,dB_t^Y \end{cases} \tag{4.1}$$

for two independent Brownian motions $(B_t^X, B_t^Y)_{t\geq 0}$, The dynamics experience a nonlinear attractive force towards the unique fixed point $(x, y) = (0, 0)$ with a nonlinear swirl in clockwise direction that becomes stronger away from the origin. As a further complication, we choose a finite time interval T for the transition to happen, and start away from the fixed point.

As observable in (2.1) we take $f(x, y) = x$, biasing the dynamics towards large values of the x-component of the process. We start at $(x_0, y_0) = (-1, -1)$, away from the fixed point, and run the process for $T = 10$, which is long enough so that the instanton is not a straight line, but short enough so that it does not completely relax to the fixed point and then leave it again at a later time (as would be the case in the limit $T \to \infty$). The resulting event is therefore a complicated interplay between the nonlinear dynamics and the conditioning on large x values, and the expected distribution of end-points is far from the invariant measure.

TABLE 2
Relative error for Z_ε for the two-dimensional test problem, comparing the naive Monte Carlo estimator with the zeroth and first order biased Monte Carlo estimators.

ε	Naive estimator	Zeroth order estimator	First order estimator
0.5	0.04	0.02	0.113
0.2	0.089	0.029	0.135
0.1	0.253	0.036	0.09
0.05	1.005	0.048	0.13
0.02	6.793	0.094	0.133
0.01	15.592	0.753	0.155
0.005	25.709	0.631	0.136

The numerical parameters are $\Delta t = 10^{-2}$ and $N = 10^6$ experiments. The results are shown in Table 2, where it can be seen that for $\varepsilon \to 0$ the relative error blows up for the naive estimator, while it is roughly constant for the zeroth and first order estimators, the later being smaller.

Figure 1 compares the different sampling procedures for $\varepsilon = 0.01$. In the naive estimator, most trajectories cluster around the deterministic decay path, swirling in clockwise direction into the origin, and consequently not reaching a large value of x. With the zeroth order estimator, and to a different degree with the first order one, the samples remain closer to the instanton (with different strengths in different regions).

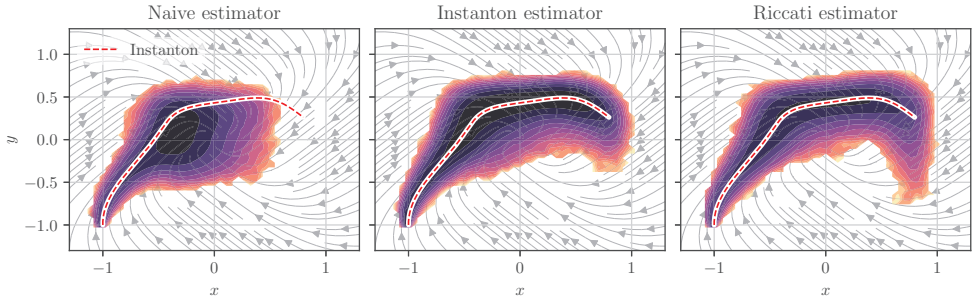

FIG. 1. *Nonlinear two-dimensional process with swirl as defined in* (4.1) *represented by arrows. The red dashed line depicts the instanton. The heat map is a histogram of the trajectories of the Monte Carlo experiment for* $\varepsilon = 0.01$. *For the naive estimator (left), most sample trajectories follow the deterministic decay trajectory by gathering around the fixed point* $(0,0)$, *so that large values for* x *are rarely observed. For the zeroth order control (center) and the first order one (right) the sample trajectories are staying more closely around the instanton, but to a different degree at different locations.*

Figure 2 (left) shows Z_ε, Z^0, and Z^1_ε as a function of ε. While Z^0 captures the constant, ε-independent limiting value of Z_ε, the departure of Z_ε from this constant is captured accurately by the first order approximation Z^1_ε for a prolonged interval in ε. As expected, for larger values of ε, higher order effects come into play, degrading the accuracy of the expansion, which could be improved by considering higher order terms (see Remark 3.2). Finally, Figure 2 (right) shows the evolution of the 2×2 matrix $(K_t)_{t\in[0,T]}$ along the instanton trajectory which is used to compute the approximate optimal bias via (3.4)–(3.6).

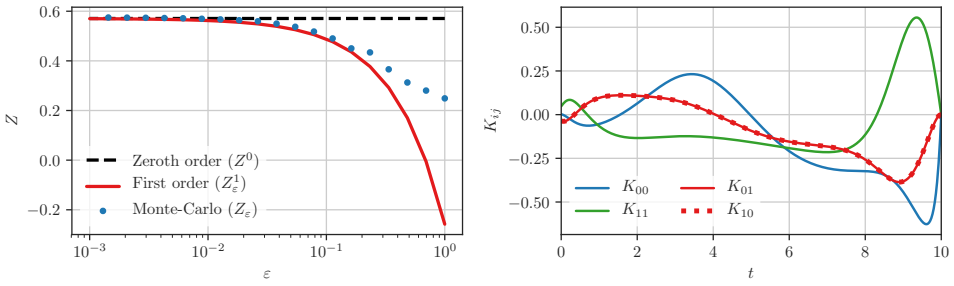

FIG. 2. *Left: Validity of the approximation of* Z_ε *in the two-dimensional problem. In the limit of small* ε, *the Monte Carlo estimator agrees with the constant zeroth order* (2.13). *For larger values of* ε, *the values of* Z_ε *depart from the constant* Z^0. *The first order approximation* Z^1_ε *captures this departure for at least an order of magnitude in* ε. *For still larger values,* Z^1_ε *and* Z_ε *diverge as expected. Right: Evolution of the four components of* K_t *in the two-dimensional problem.*

4.3. Double-well potential. We next consider a one-dimensional double-well potential with $V(x) = \frac{1}{4}(x^2-1)^2$, which has locally stable fixed points at $x = \pm 1$, and set $b(x) = -\nabla V(x)$ and $f(x) = x$. We are starting the process in the left fixed point $x_0 = -1$, so that a typical fluctuation leading to high values of f corresponds to a trajectory crossing to the right well, which becomes a rare event in the low ε limit.

This example is more complicated than the previous ones because two fixed points exist. As a consequence, in the nonconvex regions of the potential, straying from the

globally optimal path is amplified by the dynamics because forward trajectories are spreading. Interestingly, this problem is more or less pronounced depending on the given time interval T: for shorter transition times, the kinetic term $\dot{\phi}^2$ in the Freidlin–Wentzell action dominates, and the dynamics become comparably unimportant. In order to illustrate this phenomenon in our numerical experiments, we choose several values for the final time T.

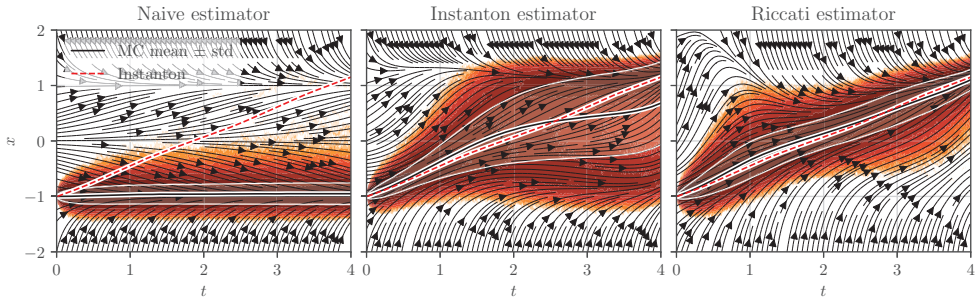

FIG. 3. *One-dimensional double-well process for $\varepsilon = 0.01$. The red dashed line depicts the instanton, while the black solid line and surrounding gray shading displays the (first order biased) Monte Carlo estimator mean and one standard deviation region around it. The heat map is a histogram of all trajectories. The black streamlines display the total drift field, i.e., the sum of the drift $b(x)$ and the bias, respectively equal to zero, $D\nabla g^0(t,x)$, and $D\nabla g^1(t,x)$ in the plots from left to right. For the naive estimator (left), almost no sample trajectory leaves the lower basin, leading to a bad estimate of the expectation. For the instanton estimator (center), many sample trajectories transition to the upper basin, with a wide variance. In the Riccati first order estimator (right), trajectories are staying more closely around the instanton trajectory, which is the optimal one in the small temperature limit.*

Again, we compute the instanton via the algorithm from [30, section III.A]. The sampling procedure, where numerical parameters are set to $\Delta t = 10^{-2}$, $T = 4$, and $N = 10^6$, is depicted in Figure 3. The heat maps display histograms of the sample trajectories. While for the naive estimator only very few manage to transition to the upper basin, many more are driven across the barrier with the instanton drift active. In the first order case, the trajectories are kept in a tube around the instanton, so that a majority of trajectories explore the space around the optimal trajectory at small temperature. We note, however, that the force field has a surprising behavior far from the instanton, which we expect given the estimate (3.8). The corresponding relative errors are listed in Table 3. In particular, while the relative error explodes with $\varepsilon \to 0$ for the naive estimator, it remains roughly constant for the zeroth order estimator and decreases significantly for the first order one. For $T = 8$ (all other parameters being the same), the observation is quite different. In fact, as shown in Table 4, for this time and even longer ones, variance reduction is no longer clearly obtained. This illustrates that the performance of our approximation of the optimal bias obtained here, which is reached under strong assumptions (uniqueness of the instanton in particular; see the beginning of section 3.1 and Remark 3.3), may be deteriorated in nonconvex cases for a large final time T.

However, the fact that the relative error ceases to decrease for $\varepsilon \to 0$ in the first order approximation does not necessarily mean that the approximation Z_ε^1 fails as well. In fact, as shown in Figure 4, Z_ε is well approximated by Z^0 and Z_ε^1 up to $T = 8$. The constant value of the zeroth order term correctly approximates the

limiting value of the numerical experiment, and the departure from that constant is correctly captured by the first order correction.

TABLE 3

Relative error for Z_ε for the double-well test case, $T = 4$, comparing the naive Monte Carlo estimator with the zeroth and first order biased Monte Carlo estimators.

ε	Naive estimator	Zeroth order estimator	First order estimator
1	0.96	1.34	13.6
0.5	2.11	2.55	14.8
0.1	81.21	8.98	25.3
0.05	800.99	6.57	5.58
0.01	343.97	7.21	1.71
0.005	835.89	5.26	0.456

TABLE 4

Relative error for the double-well test case, $T = 8$, comparing the naive Monte Carlo estimator with the zeroth and first order biased Monte Carlo estimators.

ε	Naive estimator	Zeroth Order estimator	First order estimator
1	0.89	1.53	67.9
0.5	1.65	3.14	50.1
0.1	32.66	167.0	47.1
0.05	656.07	42.0	177.0
0.01	541.15	144.	15.6
0.005	721.8	29.1	33.1

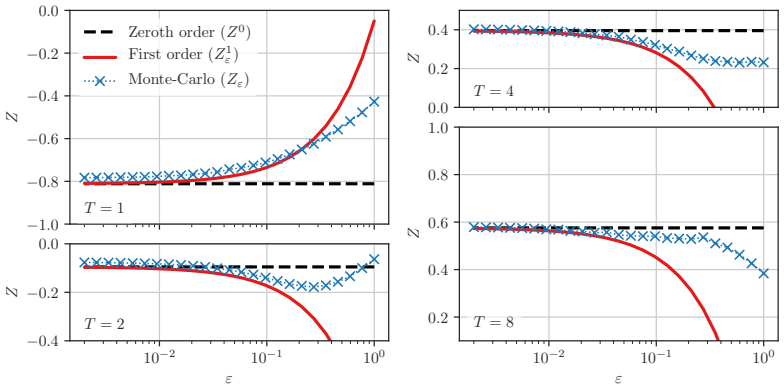

FIG. 4. *One-dimensional double-well process, validity of the approximation of the free energy Z_ε. For $\varepsilon \to 0$, the constant approximation predicts the right limiting value Z^0. Departure from this constant value is correctly approximated by the first order approximation Z^1_ε. The Monte Carlo estimated values are obtained using trajectories biased to first order as above.*

5. Discussion. In this paper, we studied the task of computing a free energy-like quantity as it commonly arises in statistical physics. For this, we consider the optimal control problem associated with finding the optimal bias to reduce the variance of an importance sampling Monte Carlo estimator. We propose a new methodology to approximate the solution to the optimal stochastic control problem. From a physics perspective, it corresponds to designing a nonhomogeneous Taylor expansion around the instanton of the dynamics. From a PDE standpoint, it is an expansion around

the characteristic curve of the zero-temperature limit of the HJB equation associated with the optimal control. Our approach differs from the more standard technique of expanding the solution in the temperature parameter (sometimes called WKB ansatz) since here the temperature appears only implicitly through the distance to the instanton of typical reacting trajectories, and the expansion is defined through solutions to ordinary differential equations instead of PDEs.

With this new tool at our disposal, we achieve two goals. First, we use our approximation to reduce the variance with a standard tilting procedure, replacing the optimal control with our expansion computed offline from simple ordinary differential equations. This approach therefore yields a very efficient method to approximate the optimal control close to the most likely realization given by the instanton, which is exactly the right regime in the low temperature limit. Next, we derive a new formula for expanding the free energy in the small temperature parameter, which refines the Freidlin–Wentzell asymptotics. We explicitly compute the first order term of the series and explain how to pursue the expansion to any order.

Finally, we propose a series of examples to illustrate the validity of our methodology. We demonstrate how the suggested approximate optimal control reduces variance in a set of numerical examples, and how the free energy expansion extends the Freidlin–Wentzell asymptotics. We further show some limitations of our approach, in particular concerning variance reduction in a nonconvex setup and for large final times.

As one can note, the theoretical arguments we use are based on quite stringent conditions: well-posedness and smoothness of the solution to the HJB problem, uniqueness of the instanton, various boundedness assumptions hidden in the expansion analysis, etc. This raises a number of questions on how in its current state the approach may fail, for instance when several solutions exist to the characteristic equation and the HJB equation only has a weak, nonsmooth solution. We know that this situation naturally arises in many contexts, which can be interpreted as a nonconvexity of the rate function from a large deviations perspective, the creation of shocks from a PDE point of view, or caustics from the physicist's viewpoint. On the other hand, from a computational perspective, even though we could quantify the closeness of the expansion to the optimal control, it is yet difficult to assess that the tilted estimator indeed reduces the variance in simulations in general [56, 25, 2, 33]. Our proposed control is therefore only the first step in the direction of rigorously establishing how and when the optimal control can be expanded around the instanton, improving on existing suggestions to use the instanton as an approximation for the optimal tilt as a mere heuristic in importance sampling, for instance in cloning algorithms [60, 30] or in instanton biased importance sampling motivated from path integral techniques [19]. Our results suggest that the approximation is built in such a way that variance is indeed reduced in the small temperature limit for simple systems, but this remains to be proved rigorously, and under which precise conditions. Although studying such issues possibly requires a significant effort, we hope the possible applications in both numerical and theoretical directions will motivate further research in this direction.

Appendix A. Proofs of section 2.2. We first prove an integration by part formula by showing that, given the dynamics (2.2) and (2.6), the expectation in (2.1) rewrites

$$A_\varepsilon = e^{g(0,x_0)}\mathbb{E}_{x_0}\left[\exp\left(\frac{1}{\varepsilon}\left[f(\widetilde{X}_T^\varepsilon) - g(T,\widetilde{X}_T^\varepsilon)\right] + \frac{1}{\varepsilon}\int_0^T \alpha(t,\widetilde{X}_t^\varepsilon)\,dt\right)\right], \tag{A.1}$$

where

$$\alpha(t,x) = \partial_t g(t,x) + \mathcal{L}g(t,x) + \frac{1}{2}|\sigma\nabla g|^2(t,x). \tag{A.2}$$

The proof of this formula relies on the Girsanov theorem and the gradient structure of the drift. We first write the Girsanov formula for the path change of measure [41, 54] between the processes $(X_t^\varepsilon)_{t\geq 0}$ and $(\widetilde{X}_t^\varepsilon)_{t\geq 0}$ (provided technical conditions are met):

$$A_\varepsilon = \mathbb{E}_{x_0}\left[e^{\frac{1}{\varepsilon}f(X_T^\varepsilon)}\right] = \mathbb{E}_{x_0}\left[e^{\frac{1}{\varepsilon}f(\widetilde{X}_T^\varepsilon) - \frac{1}{2\varepsilon}\int_0^T |\sigma\nabla g|^2(t,\widetilde{X}_t^\varepsilon)\,dt - \frac{1}{\sqrt{\varepsilon}}\int_0^T \sigma\nabla g(t,\widetilde{X}_t^\varepsilon)\,dB_t}\right]. \tag{A.3}$$

We now use Itô's formula over a trajectory of $(\widetilde{X}_t^\varepsilon)_{t\geq 0}$ using the generator (2.3):

$$dg(t,\widetilde{X}_t^\varepsilon) = \big(\partial_t g + \mathcal{L}g + \nabla g\cdot D\nabla g\big)(t,\widetilde{X}_t^\varepsilon)\,dt + \sqrt{\varepsilon}\sigma\nabla g(t,\widetilde{X}_t^\varepsilon)\,dB_t.$$

Integrating in time and dividing by ε, the above equation becomes

$$\begin{aligned}-\frac{1}{\sqrt{\varepsilon}}\int_0^T \sigma\nabla g(t,\widetilde{X}_t^\varepsilon)\,dB_t = &-\frac{g(T,\widetilde{X}_T^\varepsilon) - g(0,\widetilde{X}_0)}{\varepsilon}\\ &+\frac{1}{\varepsilon}\int_0^T \big(\partial_t g + \mathcal{L}g + |\sigma\nabla g|^2\big)(t,\widetilde{X}_t^\varepsilon)\,dt.\end{aligned}$$

Inserting this equality into (A.3) leads to (A.1).

We next turn to the derivation of the optimal control (2.8). We first note, using the Feynman–Kac formula [40, Theorem 21.1], that ψ_ε is the solution to the following backward PDE:

$$\begin{cases}\partial_t\psi_\varepsilon + \mathcal{L}\psi_\varepsilon = 0,\\ \quad \psi_\varepsilon(T,x) = e^{\frac{1}{\varepsilon}f(x)} \quad \forall\, x\in\mathbb{R}^d.\end{cases} \tag{A.4}$$

Defining $g_\varepsilon = \varepsilon\log\psi_\varepsilon$, we see that the time derivative of g_ε reads

$$\begin{aligned}\partial_t g_\varepsilon &= \varepsilon\frac{\partial_t\psi_\varepsilon}{\psi_\varepsilon} = \varepsilon\frac{-\mathcal{L}\psi_\varepsilon}{\psi_\varepsilon}\\ &= -\varepsilon e^{-g_\varepsilon/\varepsilon}\mathcal{L}e^{g_\varepsilon/\varepsilon} = \varepsilon\Big(-\varepsilon^{-1}b\cdot\nabla g_\varepsilon - e^{-g_\varepsilon/\varepsilon}\frac{\sigma\sigma^{\mathrm{T}}}{2} : \big(\nabla(e^{g_\varepsilon/\varepsilon}\nabla g_\varepsilon)\big)\Big)\\ &= \Big(-\mathcal{L}g_\varepsilon - \frac{1}{2}\,|\sigma\nabla g_\varepsilon|^2\Big).\end{aligned}$$

Using the terminal condition in (A.4) shows that g_ε is the solution to

$$\begin{cases}\partial_t g_\varepsilon + \mathcal{L}g_\varepsilon + \dfrac{1}{2}\,|\sigma\nabla g_\varepsilon|^2 = 0,\\ \qquad\qquad g_\varepsilon(T,x) = f(x) \quad \forall\, x\in\mathbb{R}^d.\end{cases} \tag{A.5}$$

As a result, (A.5) ensures that $\alpha(t,x) = 0$. Together with the terminal condition, this shows that (2.8) defines a zero-variance control since the estimator is deterministic.

Appendix B. Proof of (3.6)–(3.8). The idea is to rewrite the Feynman–Kac mode ψ_ε defined in (2.7) with the integration by part (A.1) presented in Appendix A in order to exhibit the leading behavior in ε. Consider the dynamics $(\widetilde{X}_t^\varepsilon)_{t\in[0,T]}$ defined

in (2.6) with g^1 given by (3.4). Using the Girsanov theorem like in Appendix A, starting from any $t \geq 0$ and $x \in \mathbb{R}^d$, we have

$$\begin{aligned}\psi_\varepsilon(t,x) &= \mathbb{E}_{t,x}\left[e^{\frac{1}{\varepsilon}(f(\widetilde{X}^\varepsilon_T)-g^1(T,\widetilde{X}^\varepsilon_T)+g^1(t,\widetilde{X}^\varepsilon_t))+\frac{1}{\varepsilon}\int_t^T \alpha(s,\widetilde{X}^\varepsilon_s)\,ds}\right] \\ &= e^{\frac{g^1(t,x)}{\varepsilon}}\mathbb{E}_{t,x}\left[e^{\frac{1}{\varepsilon}(f(\widetilde{X}^\varepsilon_T)-g^1(T,\widetilde{X}^\varepsilon_T))+\frac{1}{\varepsilon}\int_t^T \alpha(s,\widetilde{X}^\varepsilon_s)\,ds}\right],\end{aligned} \tag{B.1}$$

where the function α is defined in (A.2). We now perform an expansion in ε inside the expectation in (B.1). First, the process $(\widetilde{X}^\varepsilon_t)_{t\in[0,T]}$ admits the following expansion:

$$\widetilde{X}^\varepsilon_t = \phi_t + \sqrt{\varepsilon}\zeta_t + \varepsilon\beta_t + \mathrm{O}\big(\varepsilon^{3/2}\big), \tag{B.2}$$

where

$$\begin{cases} d\zeta_t = (\nabla b(\phi_t)\zeta_t + K_t\zeta_t)\,dt + \sigma\,dB_t, \\ d\beta_t = \big(\nabla b(\phi_t)\beta_t + \dfrac{1}{2}\zeta_t\cdot\nabla^2 b(\phi_t)\zeta_t + K_t\beta_t\big)dt. \end{cases}$$

This follows by expanding $(\widetilde{X}^\varepsilon_t)_{t\in[0,T]}$ around the path $(\phi_t)_{t\in[0,T]}$ and identifying the terms of different orders in ε (by Taylor-expanding the drift b). Note that we will actually not need the precise expression for the processes $(\zeta_t)_{t\in[0,T]}$ and $(\beta_t)_{t\in[0,T]}$ in what follows.

We now come back to (B.1) by first considering the terminal terms. Using (B.2), we obtain

$$\begin{aligned} f(\widetilde{X}^\varepsilon_T) &= f(\phi_T) + \sqrt{\varepsilon}\nabla f(\phi_T)\cdot\zeta_T \\ &\quad + \varepsilon\left[\nabla f(\phi_T)\cdot\beta_T + \frac{1}{2}\zeta_T\cdot\nabla^2 f(\phi_T)\zeta_T\right] + \mathrm{O}\big(\varepsilon^{3/2}\big), \\ g^1(T,\widetilde{X}^\varepsilon_T) &= \sqrt{\varepsilon}\theta_T\cdot\zeta_T + \varepsilon\left[\theta_T\cdot\beta_T + \frac{1}{2}\zeta_T\cdot K_T\zeta_T\right] + \mathrm{O}\big(\varepsilon^{3/2}\big). \end{aligned}$$

The terminal conditions for $(\theta_t)_{t\in[0,T]}$ and $(K_t)_{t\in[0,T]}$ lead to

$$f(\widetilde{X}^\varepsilon_T) - g^1(T,\widetilde{X}^\varepsilon_T) = f(\phi_T) + \mathrm{O}\big(\varepsilon^{3/2}\big). \tag{B.3}$$

It remains to study the integral part in (B.1), for which we expand the Girsanov weight α with (B.2). We have

$$\alpha(t,\widetilde{X}^\varepsilon_t) = \partial_t g^1(t,\widetilde{X}^\varepsilon_t) + b(\widetilde{X}^\varepsilon_t)\cdot\nabla g^1(t,\widetilde{X}^\varepsilon_t) + \frac{\varepsilon}{2}D:\nabla^2 g^1(t,\widetilde{X}^\varepsilon_t) + \frac{1}{2}|\sigma\nabla g^1(t,\widetilde{X}^\varepsilon_t)|^2. \tag{B.4}$$

First, we notice that (we omit below the dependency of ϕ_t, θ_t, and K_t on time for concision)

$$\begin{aligned} \partial_t g^1(t,x) &= \dot\theta\cdot(x-\phi) - \theta\cdot\dot\phi + \frac{1}{2}(x-\phi)\cdot\dot K(x-\phi) - (x-\phi)\cdot K\dot\phi, \\ \nabla g^1(t,x) &= \theta + K(x-\phi). \end{aligned}$$

As a result, (B.4) reads

$$\begin{aligned} \alpha(t,\widetilde{X}^\varepsilon_t) &= \dot\theta\cdot(\widetilde{X}^\varepsilon_t-\phi) - \theta\dot\phi + \frac{1}{2}(\widetilde{X}^\varepsilon_t-\phi)\cdot\dot K(\widetilde{X}^\varepsilon_t-\phi) - \frac{1}{2}(\widetilde{X}^\varepsilon_t-\phi)\cdot K\dot\phi \\ &\quad - \frac{1}{2}(\widetilde{X}^\varepsilon_t-\phi)\cdot K^{\mathrm{T}}\dot\phi + \frac{\varepsilon}{2}D:K \\ &\quad + b(\widetilde{X}^\varepsilon_t)\cdot\big(\theta + K(\widetilde{X}^\varepsilon_t-\phi)\big) + \frac{1}{2}\big|\sigma\big(\theta + K(\widetilde{X}^\varepsilon_t-\phi)\big)\big|^2, \end{aligned}$$

which may be reorganized as follows (using (2.12) for estimating the time derivatives of ϕ and θ):

$$\begin{aligned}\alpha(t,\widetilde{X}_t^\varepsilon) = &-|\sigma\theta|^2+\frac{\varepsilon}{2}D:K+\frac{1}{2}|\sigma\theta|^2+\frac{1}{2}\theta\cdot DK(\widetilde{X}_t^\varepsilon-\phi)+\frac{1}{2}\theta\cdot DK^{\mathrm{T}}(\widetilde{X}_t^\varepsilon-\phi)\\&+\frac{1}{2}\big|\sigma K(\widetilde{X}_t^\varepsilon-\phi)\big|^2+\theta\big(b(\widetilde{X}_t^\varepsilon)-b(\phi)\big)\\&-\theta\cdot\nabla b(\phi)(\widetilde{X}_t^\varepsilon-\phi)+\frac{1}{2}(\widetilde{X}_t^\varepsilon-\phi)\cdot\dot{K}(\widetilde{X}_t^\varepsilon-\phi)\\&+\frac{1}{2}\big(b(\widetilde{X}_t^\varepsilon)-b(\phi)\big)\cdot K(\widetilde{X}_t^\varepsilon-\phi)\\&+\frac{1}{2}(\widetilde{X}_t^\varepsilon-\phi)\cdot K^{\mathrm{T}}\big(b(\widetilde{X}_t^\varepsilon)-b(\phi)\big)-\frac{1}{2}\theta\cdot DK(\widetilde{X}_t^\varepsilon-\phi)\\&-\frac{1}{2}\theta\cdot DK^{\mathrm{T}}(\widetilde{X}_t^\varepsilon-\phi)\\=&-\frac{1}{2}|\sigma\theta|^2+\frac{\varepsilon}{2}D:K+\frac{1}{2}\big|\sigma K(\widetilde{X}_t^\varepsilon-\phi)\big|^2+\theta\big(b(\widetilde{X}_t^\varepsilon)-b(\phi)\big)\\&-\theta\cdot\nabla b(\phi)(\widetilde{X}_t^\varepsilon-\phi)\\&+\frac{1}{2}(\widetilde{X}_t^\varepsilon-\phi)\cdot\dot{K}(\widetilde{X}_t^\varepsilon-\phi)+\frac{1}{2}\big(b(\widetilde{X}_t^\varepsilon)-b(\phi)\big)\cdot K(\widetilde{X}_t^\varepsilon-\phi)\\&+\frac{1}{2}(\widetilde{X}_t^\varepsilon-\phi)\cdot K^{\mathrm{T}}\big(b(\widetilde{X}_t^\varepsilon)-b(\phi)\big).\end{aligned}$$

Inserting the expansion (B.2) then leads to

$$\begin{aligned}\alpha(t,\widetilde{X}_t^\varepsilon) = &-\frac{1}{2}|\sigma\theta|^2+\frac{\varepsilon}{2}D:K+\frac{\varepsilon}{2}|\sigma\zeta_tK|^2+\sqrt{\varepsilon}\theta\cdot\nabla b(\phi)\zeta_t+\varepsilon\theta\cdot\nabla b(\phi)\beta_t\\&+\frac{\varepsilon}{2}\zeta_t\cdot\theta\nabla^2b(\phi)\zeta_t\\&-\theta\cdot\nabla b(\phi)(\sqrt{\varepsilon}\zeta_t+\varepsilon\beta_t)+\frac{\varepsilon}{2}\zeta_t\cdot\dot{K}\zeta_t+\frac{1}{2}\varepsilon\zeta_t\cdot\nabla b(\phi)K\zeta_t\\&+\frac{1}{2}\varepsilon\zeta_t\cdot K^{\mathrm{T}}\nabla b(\phi)\zeta_t.\end{aligned}$$

Now, we may identify the terms of various orders in $\sqrt{\varepsilon}$ in the above equation. The leading order is equal to $-|\sigma\theta|^2/2$, as expected. At order $\sqrt{\varepsilon}$ we have

$$\theta\cdot\nabla b(\phi)\zeta_t-\theta\cdot\nabla b(\phi)\zeta_t=0.$$

We now turn to the terms of order ε (excluding the term in $D:K_t$ for now) that are given by

$$\begin{aligned}&\frac{1}{2}\zeta_t\cdot KDK\zeta_t+\theta\cdot\nabla b(\phi)\beta_t-\theta\cdot\nabla b(\phi)\beta_t\\&\qquad+\frac{1}{2}\zeta_t\cdot\dot{K}\zeta_t+\frac{1}{2}\zeta_t\cdot\theta\nabla^2b(\phi)\zeta_t\\&\qquad+\frac{1}{2}\zeta_t\cdot K^{\mathrm{T}}\nabla b(\phi)\zeta_t+\frac{1}{2}\zeta_t\cdot\nabla b(\phi)^{\mathrm{T}}K\zeta_t.\end{aligned}$$

We see that the terms proportional to β_t cancel, while the quadratic product in ζ_t factors out, so we have

$$\dot{K}_t+K_tDK_t+K_t^{\mathrm{T}}\nabla b(\phi)+\nabla b^{\mathrm{T}}(\phi)K+\theta\cdot\nabla^2b(\phi), \tag{B.5}$$

which is equal to 0 since $(K_t)_{t\in[0,T]}$ is the solution to (3.6). Gathering the above results shows that (B.4) becomes, in the small ε limit,

$$\alpha(t,\widetilde{X}_t^\varepsilon) = -\frac{1}{2}|\sigma\theta_t|^2 + \frac{\varepsilon}{2}D : K_t + \mathrm{O}\big(\varepsilon^{3/2}\big). \tag{B.6}$$

In order to use the above estimates in (B.1), we need (B.2) to hold. Such a perturbative formula holds for the dynamics (2.6) provided it starts from the correct initial condition at time t when computing (B.1)—in other words, if x is far from ϕ_t, the error may well be large. One way to solve this problem is to note that (B.2) is actually satisfied at any time when the process is started at time t from the value ϕ_t of the instanton at that time. Introducing the shorthand notation

$$Y_T = e^{\frac{1}{\varepsilon}(f(\widetilde{X}_T^\varepsilon) - g^1(T,\widetilde{X}_T^\varepsilon)) + \frac{1}{\varepsilon}\int_t^T \alpha(s,\widetilde{X}_s^\varepsilon)\,ds},$$

we may thus rewrite (B.1) as

$$\psi_\varepsilon(t,x) = e^{\frac{g^1(t,x)}{\varepsilon}}\mathbb{E}_{t,\phi_t}\left[e^{\frac{1}{\varepsilon}(f(\widetilde{X}_T^\varepsilon) - g^1(T,\widetilde{X}_T^\varepsilon)) + \frac{1}{\varepsilon}\int_t^T \alpha(s,\widetilde{X}_s^\varepsilon)\,ds}\right]\frac{\mathbb{E}_{t,x}[Y_T]}{\mathbb{E}_{t,\phi_t}[Y_T]}. \tag{B.7}$$

In the first expectation starting from (t,ϕ_t) we can then perform the expansions as above.

Plugging the estimates (B.3) and (B.6) into the first expectation in (B.7), we obtain

$$\psi_\varepsilon(t,x) = e^{\frac{g^1(t,x)}{\varepsilon}}\mathbb{E}_{t,\phi_t}\left[e^{\frac{1}{\varepsilon}\left(f(\phi_T) - \frac{1}{2}\int_t^T |\sigma\theta_s|^2\,ds + \frac{\varepsilon}{2}\int_t^T D:K_s\,ds + \mathrm{O}(\varepsilon^{3/2})\right)}\right]\frac{\mathbb{E}_{t,x}[Y_T]}{\mathbb{E}_{t,\phi_t}[Y_T]}.$$

Taking the logarithm and multiplying by ε then leads to

$$\begin{aligned} g_\varepsilon(t,x) = \varepsilon\log\psi_\varepsilon(t,x) &= g^1(t,x) \\ &+ f(\phi_T) - \frac{1}{2}\int_t^T |\sigma\theta_s|^2\,ds + \frac{\varepsilon}{2}\int_t^T D:K_s\,ds \\ &+ \varepsilon\log\mathbb{E}_{t,\phi_t}\left[e^{\frac{1}{\varepsilon}\mathrm{O}(\varepsilon^{3/2})}\right] + \varepsilon\left[\log\mathbb{E}_{t,x}[Y_T] - \log\mathbb{E}_{t,\phi_t}[Y_T]\right]. \end{aligned}$$

Assuming that $(t,x)\to\log\mathbb{E}_{t,x}[Y_T]$ is smooth, we consider the Taylor expansion in x around ϕ_t for the difference of logarithms in addition to the small ε limit. As a result, the optimal control (2.8) admits the following expansion:

$$g_\varepsilon(t,x) = g^1(t,x) + f(\phi_T) - \frac{1}{2}\int_t^T |\sigma\theta_s|^2\,ds + \frac{\varepsilon}{2}\int_t^T D:K_s\,ds + \mathrm{o}(\varepsilon) + \varepsilon\mathrm{o}\big(x-\phi_t\big)$$

in the small ε regime and for x close to ϕ_t, where g^1 is defined in (3.4). This provides the desired result.

Acknowledgments. The authors warmly thank Eric Vanden-Eijnden for his insightful advice on this work, as well as Gabriel Stoltz and Hugo Touchette for interesting discussions. The authors are also grateful to the referees for providing valuable input and mentioning interesting references.

REFERENCES

[1] L. Angeli, S. Grosskinsky, A. M. Johansen, and A. Pizzoferrato, *Rare event simulation for stochastic dynamics in continuous time*, J. Stat. Phys., 176 (2019), pp. 1185–1210.

[2] S. Asmussen and P. W. Glynn, *Stochastic Simulation: Algorithms and Analysis*, Stoch. Model. Appl. Probab. 57, Springer, New York, 2007.

[3] F. Bouchet, K. Gawedzki, and C. Nardini, *Perturbative calculation of quasi-potential in non-equilibrium diffusions: A mean-field example*, J. Stat. Phys., 163 (2016), pp. 1157–1210.

[4] F. Bouchet and J. Reygner, *Generalisation of the Eyring–Kramers transition rate formula to irreversible diffusion processes*, Ann. Henri Poincaré, 17 (2016), pp. 3499–3532.

[5] C.-E. Bréhier and T. Lelièvre, *On a new class of score functions to estimate tail probabilities of some stochastic processes with adaptive multilevel splitting*, Chaos, 29 (2019), 033126.

[6] J. Bucklew, *Introduction to Rare Event Simulation*, Springer Series in Statistics, Springer-Verlag, New York, 2004.

[7] M. G. Crandall and P.-L. Lions, *Viscosity solutions of Hamilton-Jacobi equations*, Trans. Amer. Math. Soc., 277 (1983), pp. 1–42.

[8] T. Dean and P. Dupuis, *Splitting for rare event simulation: A large deviation approach to design and analysis*, Stochastic Process. Appl., 119 (2009), pp. 562–587, https://doi.org/10.1016/j.spa.2008.02.017.

[9] P. Del Moral, *Feynman-Kac Formulae*, Probability and Its Applications, Springer, 2004.

[10] P. Del Moral and L. Miclo, *On the stability of nonlinear Feynman-Kac semigroups*, Ann. Fac. Sci. Toulouse Math. (6), 11 (2002), pp. 135–175.

[11] C. Dellago, P. G. Bolhuis, and P. L. Geissler, *Transition path sampling*, Adv. Chem. Phys., 123 (2003), pp. 1–78, https://doi.org/10.1002/0471231509.ch1.

[12] A. Dembo and O. Zeitouni, *Large Deviations Techniques and Applications*, Stoch. Model. Appl. Probab. 38, Springer-Verlag, Berlin, 2010.

[13] A. Doucet and A. M. Johansen, *A tutorial on particle filtering and smoothing: Fifteen years later*, in The Oxford Handbook of Nonlinear Filtering, Oxford University Press, Oxford, 2011, pp. 656–704.

[14] P. Dupuis and H. Wang, *Subsolutions of an Isaacs equation and efficient schemes for importance sampling*, Math. Oper. Res., 32 (2007), pp. 723–757.

[15] W. E, J. Han, and A. Jentzen, *Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations*, Comm. Math. Stat., 5 (2017), pp. 349–380.

[16] W. E, W. Ren, and E. Vanden-Eijnden, *String method for the study of rare events*, Phys. Rev. B, 66 (2002), 052301, https://doi.org/10.1103/PhysRevB.66.052301.

[17] W. E, W. Ren, and E. Vanden-Eijnden, *Minimum action method for the study of rare events*, Comm. Pure Appl. Math., 57 (2004), pp. 637–656, https://doi.org/10.1002/cpa.20005.

[18] W. E, W. Ren, and E. Vanden-Eijnden, *Finite temperature string method for the study of rare events*, J. Phys. Chem. B, 109 (2005), pp. 6688–6693.

[19] L. Ebener, G. Margazoglou, J. Friedrich, L. Biferale, and R. Grauer, *Instanton based importance sampling for rare events in stochastic PDEs*, Chaos, 29 (2019), 063102, https://doi.org/10.1063/1.5085119.

[20] L. C. Evans, *Partial Differential Equations*, Grad. Stud. Math. 19, American Mathematical Society, Providence, RI, 2010.

[21] G. Ferré and H. Touchette, *Adaptive sampling of large deviations*, J. Stat. Phys., 172 (2018), pp. 1525–1544.

[22] W. H. Fleming and M. R. James, *Asymptotic series and exit time probabilities*, Ann. Probab., 20 (1992), pp. 1369–1384.

[23] W. H. Fleming and H. M. Soner, *Controlled Markov Processes and Viscosity Solutions*, Stoch. Model. Appl. Probab. 25, Springer, New York, 2006.

[24] M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems*, Grundlehren Math. Wiss. 260, Springer, Heidelberg, 1998.

[25] P. Glasserman and Y. Wang, *Counterexamples in importance sampling for large deviations probabilities*, Ann. Appl. Probab., 7 (1997), pp. 731–746, https://doi.org/10.1214/aoap/1034801251.

[26] T. Grafke, R. Grauer, and T. Schäfer, *The instanton method and its numerical implementation in fluid mechanics*, J. Phys. A, 48 (2015), 333001, https://doi.org/10.1088/1751-8113/48/33/333001.

[27] T. Grafke, R. Grauer, T. Schäfer, and E. Vanden-Eijnden, *Arclength parametrized Hamilton's equations for the calculation of instantons*, Multiscale Model. Simul., 12 (2014),

pp. 566–580, https://doi.org/10.1137/130939158.
[28] T. Grafke, T. Schäfer, and E. Vanden-Eijnden, *Long term effects of small random perturbations on dynamical systems: Theoretical and computational tools*, in Recent Progress and Modern Challenges in Applied Mathematics, Modeling and Computational Science, Fields Institute Communications, Springer, New York, 2017, pp. 17–55, https://doi.org/10.1007/978-1-4939-6969-2_2.
[29] T. Grafke, T. Schäfer, and E. Vanden-Eijnden, *Sharp Asymptotic Estimates for Expectations, Probabilities, and Mean First Passage Times in Stochastic Systems with Small Noise*, preprint, https://arxiv.org/abs/2103.04837, 2021.
[30] T. Grafke and E. Vanden-Eijnden, *Numerical computation of rare events via large deviation theory*, Chaos, 29 (2019), 063118, https://doi.org/10.1063/1.5084025.
[31] R. Graham and T. Tél, *On the weak-noise limit of Fokker-Planck models*, J. Stat. Phys., 35 (1984), pp. 729–748.
[32] P. Grassberger, *Go with the winners: A general Monte Carlo strategy*, Comput. Phys. Comm., 147 (2002), pp. 64–70, https://doi.org/10.1016/S0010-4655(02)00205-9.
[33] A. Guyader and H. Touchette, *Efficient large deviation estimation based on importance sampling*, J. Stat. Phys., 181 (2020), pp. 551–586.
[34] J. Han, *Deep Learning Approximation for Stochastic Control Problems*, preprint, https://arxiv.org/abs/1611.07422, 2016.
[35] J. Han, A. Jentzen, and W. E, *Solving high-dimensional partial differential equations using deep learning*, Proc. Natl. Acad. Sci. USA, 115 (2018), pp. 8505–8510.
[36] C. Hartmann and C. Schütte, *Efficient rare event simulation by optimal nonequilibrium forcing*, J. Stat. Mech. Theory Exp., 2012 (2012), 11004.
[37] C. Hartmann, C. Schütte, and W. Zhang, *Model reduction algorithms for optimal control and importance sampling of diffusions*, Nonlinearity, 29 (2016), pp. 2298–2326.
[38] M. Heymann and E. Vanden-Eijnden, *The geometric minimum action method: A least action principle on the space of curves*, Comm. Pure Appl. Math., 61 (2008), pp. 1052–1117, https://doi.org/10.1002/cpa.20238.
[39] V. Jurdjevic, *Geometric Control Theory*, Cambridge Stud. Adv. Math. 52, Cambridge University Press, Cambridge, 1997.
[40] O. Kallenberg, *Foundations of Modern Probability*, Probability and Its Applications, Springer-Verlag, New York, 2002.
[41] I. Karatzas and S. Shreve, *Brownian Motion and Stochastic Calculus*, Grad. Texts Math. 113, Springer-Verlag, New York, 2012.
[42] O. Kebiri, L. Neureither, and C. Hartmann, *Adaptive importance sampling with forward-backward stochastic differential equations*, in International Workshop on Stochastic Dynamics out of Equilibrium, Springer Proc. Math. Stat. 282, Springer, Cham, 2019, pp. 265–281.
[43] Y. I. Kifer, *On the asymptotics of the transition density of processes with small diffusion*, Theory Probab. Appl., 21 (1977), pp. 513–522, https://doi.org/10.1137/1121063.
[44] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*, Stoch. Model. Appl. Probab. 23, Springer, Berlin, 1992.
[45] P. Lancaster and L. Rodman, *Algebraic Riccati Equations*, Clarendon Press, Oxford University Press, New York, 1995.
[46] V. Lecomte and J. Tailleur, *A numerical approach to large deviations in continuous time*, J. Stat. Mech. Theory Exp., 2007 (2007), P03004.
[47] T. Lelièvre and G. Stoltz, *Partial differential equations and stochastic methods in molecular dynamics*, Acta Numer., 25 (2016), pp. 681–880.
[48] H. Lie, C. Schütte, and C. Hartmann, *Martingale-based Gradient Descent Algorithm for Estimating Free Energy Values of Diffusions*, preprint, 2015.
[49] T. Nemoto, F. Bouchet, R. L. Jack, and V. Lecomte, *Population-dynamics method with a multicanonical feedback control*, Phys. Rev. E, 93 (2016), 062123.
[50] T. Nemoto, E. G. Hidalgo, and V. Lecomte, *Finite-time and finite-size scalings in the evaluation of large-deviation functions: Analytical study using a birth-death process*, Phys. Rev. E, 95 (2017), 012102.
[51] D. Nickelsen and H. Touchette, *Anomalous scaling of dynamical large deviations*, Phys. Rev. Lett., 121 (2018), 090602.
[52] N. Nüsken and L. Richter, *Solving High-dimensional Hamilton-Jacobi-Bellman PDEs Using Neural Networks: Perspectives from the Theory of Controlled Diffusions and Measures on Path Space*, preprint, https://arxiv.org/abs/2005.05409, 2020.
[53] H. Pham, *Continuous-Time Stochastic Control and Optimization with Financial Applications*, Stoch. Model. Appl. Probab. 61, Springer-Verlag, Berlin, 2009.

[54] L. Rey-Bellet, *Ergodic properties of Markov processes*, in Open Quantum Systems II, Lecture Notes in Math. 1881, Springer, Berlin, 2006, pp. 1–39.
[55] M. Rousset, *On the control of an interacting particle estimation of Schrödinger ground states*, SIAM J. Math. Anal., 38 (2006), pp. 824–844, https://doi.org/10.1137/050640667.
[56] J. S. Sadowsky and J. A. Bucklew, *On large deviations theory and asymptotically efficient Monte Carlo estimation*, IEEE Trans. Inf. Theory, 36 (1990), pp. 579–588.
[57] T. Schorlepp, T. Grafke, and R. Grauer, *Gel'fand-Yaglom type equations for calculating fluctuations around instantons in stochastic systems*, J. Phys. A: Math. Theor., 54 (2021), 235003, https://doi.org/10.1088/1751-8121/abfb26.
[58] J. Tailleur and V. Lecomte, *Simulation of large deviation functions using population dynamics*, AIP Conf. Proc., 1091 (2009), pp. 212–219.
[59] E. Vanden-Eijnden and J. Weare, *Rare event simulation of small noise diffusions*, Comm. Pure Appl. Math., 65 (2012), pp. 1770–1803, https://doi.org/10.1002/cpa.21428.
[60] J. Wouters and F. Bouchet, *Rare event computation in deterministic chaotic systems using genealogical particle analysis*, J. Phys. A, 49 (2016), 374002, https://doi.org/10.1088/1751-8113/49/37/374002.
[61] W. Zhang, H. Wang, C. Hartmann, M. Weber, and C. Schütte, *Applications of the cross-entropy method to importance sampling and optimal control of diffusions*, SIAM J. Sci. Comput., 36 (2014), pp. A2654–A2672, https://doi.org/10.1137/14096493X.