

STOCHASTIC VISCOSITY APPROXIMATIONS OF HAMILTON–JACOBI EQUATIONS AND VARIANCE REDUCTION

GRÉGOIRE FERRÉ*

Abstract. We consider the computation of free energy-like quantities for diffusions when resorting to Monte Carlo simulation is necessary, for instance in high dimension. Such stochastic computations typically suffer from high variance, in particular in a low noise regime, because the expectation is dominated by rare trajectories for which the observable reaches large values. Although importance sampling, or tilting of trajectories, is now a standard technique for reducing the variance of such estimators, quantitative criteria for proving that a given control reduces variance are scarce, and often do not apply to practical situations. The goal of this work is to provide a quantitative criterion for assessing whether a given bias reduces variance, and at which scale. We rely for this on a recently introduced notion of stochastic solution for Hamilton–Jacobi–Bellman (HJB) equations. Based on this tool, we introduce the notion of k -stochastic viscosity approximation (SVA) of a HJB equation. We next prove that such approximate solutions are associated with estimators having a relative variance of order $k - 1$ at log-scale. In particular, a sampling scheme built from a 1-SVA has bounded variance as noise goes to zero. Finally, in order to show that our definition is relevant, we provide examples of stochastic viscosity approximations of order one and two, with a numerical illustration confirming our theoretical findings.

Mathematics Subject Classification. 49L25, 60H30, 65C05.

Received July 9, 2022. Accepted May 10, 2023.

1. INTRODUCTION

Rare events play a crucial role in many scientific contexts, from networks dynamics to molecular systems and climatic hazards. By rare we mean that the frequency of appearance of such events is low compared to the impact of its realization, relatively to the time scale of observation. If one is interested in sampling these phenomena to estimate probabilities of occurrence or average impact, standard Monte Carlo methods generally lead to a high relative variance, hence poor quality of estimation. This can be understood heuristically, as few samples realize large values of the estimator. Hence, many samples are required to obtain an accurate estimation of the quantity of interest. This is typically not affordable in practice, since one single simulation of a realistic system may already be computationally expensive.

Keywords and phrases. Monte Carlo methods, Variance reduction, Feynman–Kac formula, Hamilton–Jacobi equations, Viscosity solutions.

Capital Fund Management, 23-25 rue de l'Université, 75007 Paris, France.

*Corresponding author: gregoire.ferre@ponts.org

A natural goal is thus to reduce the relative variance of a naive Monte Carlo estimator, and we focus here on the case of diffusions. Many techniques have been designed so far, in particular genetic algorithms relying on a selection mechanism [1, 7], and importance sampling, which modifies the system to enhance the probability of observing the chosen rare event [5]. We consider here the latter technique, which has been applied to various systems of interest, see for instance [9, 20, 25] and references therein.

However, among the possible biases that can be chosen to enhance the probability of rare events, many are not efficient to achieve variance reduction – a wrong bias may even deteriorate the properties of the estimator [19]. In the specific case of diffusion processes we know that, in general, there exists a unique optimal bias leading to a zero-variance estimator [15, 25, 31], and that this control is the solution to a Hamilton–Jacobi–Bellman (HJB) partial differential equation (PDE) depending on the parameters of the system [17]. This PDE also provides the expectation of interest for any initial condition, hence one does not need to resort to Monte Carlo simulation when the HJB problem is solved.

Since a close form of the optimal control is not available in most cases, it needs to be approximated. As for any evolution PDE, finite difference techniques can be used [6, 30], with error estimates available for instance in Lebesgue-like norms. More recently, a series of works also address the approximation of Hamilton–Jacobi problems through a more sampling-oriented point of view [20, 23, 31, 32].

A major issue with all these techniques is that, although they may perform well in some practical cases, one is generally not ensured that the approximation of the optimal control indeed reduces variance. For the case of diffusions, it is proved for the time consuming approach proposed in [31], under a stringent condition of locally uniform convergence of the approximation. More recently, the authors in [21] provide general conditions for an abstract importance sampling estimator to be efficient. Although powerful, the generality of the latter approach makes it difficult to apply to diffusion processes.

The purpose of this paper is to provide a new criterion for assessing the performance of an approximate control in terms of variance reduction. For this we are inspired by the important work [26], which introduces a notion of solution to the Feynman–Kac and Hamilton–Jacobi equations that relies on a stochastic interpretation of the generator. This allows to consider non-smooth problems and solutions – we may call such solutions “stochastic solutions”. This theory uses a finite entropy condition in relation with the Girsanov theorem and extended generators based on martingale techniques. Our main contribution is to adapt this notion of solution to an approximation context: this leads us to introduce the notion of stochastic viscosity approximation of order k (or k -SVA). Such approximate controls solve the Hamilton–Jacobi equation along trajectories tilted by the control itself up to a small process of order k in the noise term coefficient (which we call temperature). We naturally introduce this condition through the Girsanov path change of measure, and then prove that an importance sampling estimator based on such a control is k -log efficient (in particular its relative variance ratio decays exponentially at order $k - 1$), which is our main result.

In order to illustrate that our method can be applied to actual approximation schemes, we consider the approach proposed in [20] to bias the dynamics along the instanton, or reacting path of the dynamics. We prove, under geometric conditions, that the estimator proposed in this work is a stochastic viscosity approximation of order 1, hence it is 1-log efficient. We then turn to the next order approximation presented in [12], showing that under similar conditions it provides a stochastic viscosity approximation of order 2, hence it is 2-log efficient and with logarithmic relative variance decaying linearly to zero. A simple numerical application confirms our theoretical findings. We emphasize that the proposed approximation is not a standard Taylor series expansion in the temperature as in [16], but is a polynomial expansion around a particular path called the instanton, which is quite unusual. We also mention that this expansion for now only applies in the situation where the instanton is unique, which is quite limiting in practical situations. Indeed, realistic potential energy landscapes are strongly non-convex, which typically leads to multiple instantons. We however believe that our results can be generalized to the case of a finite number of instantons, which is the purpose of ongoing work.

Our results are organized as follows. Section 2 presents the general setting of the work with our assumptions, together with reminders on importance sampling. In Section 3, we present our main definition of stochastic viscosity approximation, and the associated variance reduction property. The case of the approximation expanding

around the instanton path is presented in Section 4, with an illustrative numerical application. Finally, we discuss our results and further research directions in Section 5.

2. CONTROLLED DIFFUSIONS AND HAMILTON–JACOBI–BELLMAN EQUATIONS

This section is devoted to generalities on controlled diffusions and importance sampling. We first present in Section 2.1 our setting together with the basic assumptions on the dynamics, before turning in Section 2.2 to the notion of importance sampling. Section 2.3 recalls some elements on the optimal control associated with our problem.

2.1. Dynamics and free energy

This section introduces some notation used throughout. We consider a diffusion process $(X_t^\varepsilon)_{t \geq 0}$ evolving in \mathbb{R}^d with $d \in \mathbb{N}^*$, and satisfying the following stochastic differential equation (SDE):

$$dX_t^\varepsilon = b(X_t^\varepsilon) dt + \sqrt{\varepsilon} \sigma dB_t, \tag{1}$$

where $b : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times m}$ and $(B_t)_{t \geq 0}$ is a m -dimensional Brownian motion for some $m \in \mathbb{N}^*$. In the above equation, $\varepsilon > 0$ is a parameter that may vary. At this stage the initial condition of $(X_t^\varepsilon)_{t \geq 0}$ is a random variable with arbitrary probability distribution. We denote by \cdot the scalar product on \mathbb{R}^d , while, ∇^2 denotes the Hessian operator and for two matrices A, B belonging to $\mathbb{R}^{d \times m}$, we write $A : B = \text{Tr}(A^T B)$. The dynamics (1) induces a semigroup of operators $(P_t^\varepsilon)_{t \geq 0}$ over the space of continuous bounded functions $C_b(\mathbb{R}^d)$ defined by

$$\forall \varphi \in C_b(\mathbb{R}^d), \quad \forall t \geq 0, \quad P_t^\varepsilon \varphi : x \in \mathbb{R}^d \rightarrow \mathbb{E}_x[\varphi(X_t^\varepsilon)], \tag{2}$$

where \mathbb{E}_x denotes expectation over all realizations of (1) when the process is started at $x \in \mathbb{R}^d$. We will also use the notation $\mathbb{E}_{r,x}$ when the process is started at position x at time $r > 0$.

The typical problem we want to address in this paper is the computation of the following exponential-type expectation

$$A^\varepsilon = \mathbb{E}_{x_0} \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right], \tag{3}$$

or the associated free energy

$$Z^\varepsilon = \varepsilon \log A^\varepsilon,$$

where $T > 0$ is a final time and $x_0 \in \mathbb{R}^d$ a fixed initial condition for (1) at time $t = 0$. In particular A^ε serves as a smoothed version of the probability:

$$\mathbb{P}(X_T^\varepsilon \in S),$$

for a Borel set $S \subset \mathbb{R}^d$. This probability, which is interesting for various applications, coincides with (3) when $f(x) = 0$ for $x \in \mathbb{R}^d$ and $f(x) = -\infty$ for $x \notin \mathbb{R}^d$.

For simplifying technical arguments, we use rather strong assumptions gathered below, which still encompass various interesting cases.

Assumption 1. The following regularity requirements hold:

- (1) The matrix $D = \sigma \sigma^T$ is positive definite².
- (2) The field b belongs to $C^\infty(\mathbb{R}^d)^d$ and is globally Lipschitz, *i.e.* there is C_{lip} such that

$$\forall x, y \in \mathbb{R}^d, \quad |b(x) - b(y)| \leq C_{\text{lip}} |x - y|.$$

- (3) The function f belongs to $C^\infty(\mathbb{R}^d)$ and is upper bounded.

²We could also consider multiplicative noise under a uniform ellipticity condition.

Under Assumption 1 it is standard that (1) has a unique strong solution for all times [24]. The semigroup $(P_t^\varepsilon)_{t \geq 0}$ is a strongly continuous semigroup (by continuity of the trajectories) with elliptic regularity. Its generator L^ε has domain $\mathcal{D}(L^\varepsilon)$ and is defined by

$$\forall \varphi \in \mathcal{D}(L^\varepsilon), \quad L^\varepsilon \varphi = \lim_{t \rightarrow 0} \frac{P_t^\varepsilon \varphi - \varphi}{t},$$

where the limit is in supremum norm. When tested over smooth enough functions, Itô calculus shows that L^ε is represented by the differential operator \mathcal{L}^ε defined by

$$\forall \varphi \in C^2(\mathbb{R}^d), \quad \mathcal{L}^\varepsilon \varphi = b \cdot \nabla \varphi + \frac{\varepsilon}{2} D : \nabla^2 \varphi. \tag{4}$$

Finally, Assumption 1 ensures that A^ε is finite for all $\varepsilon > 0$ (because f is upper bounded).

The restrictive conditions of Assumption 1 are imposed in order to avoid non-essential technical details, but can be relaxed depending on the situation of interest. For instance, the second point of the assumption ensures well-posedness of the SDE (1) and the associated PDE, but could be replaced by a local Lipschitz assumption together with Lyapunov condition on the generator [29]. Similarly, this global Lipschitz condition is used in Section 4 to obtain standard error estimates, but could also be replaced there by a local Lipschitz condition together with appropriate growth conditions at infinity on the drift b . Finally, the upper boundedness of f could be weakened under a Lyapunov condition to allow functions that are not bounded from above, see the technique used in [14].

2.2. Estimators and importance sampling

In a small temperature regime and even more when the dimension d is large, it is generally impossible to estimate the integral in (3) by numerical quadrature over a discretized grid. This is why we typically resort to stochastic sampling of the expectation. We know that the estimator of A^ε in (3) based on generating independent samples of the random variable

$$\mathcal{A}^\varepsilon = e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)}$$

is unbiased but typically has high (relative) variance. This is particularly true in the low ε regime, which is the one of interest in physical applications.

A possible way to reduce the variance is to control the process (1) with a function g (for now arbitrary), and to consider the so-called tilted process $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ defined through the stochastic differential equation

$$d\tilde{X}_t^\varepsilon = b(\tilde{X}_t^\varepsilon) dt + D\nabla g(t, \tilde{X}_t^\varepsilon) dt + \sqrt{\varepsilon} \sigma dB_t. \tag{5}$$

We generically denote by $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ a process tilted by a function g to be made precise, and belonging to the class of functions described below.

Assumption 2. A real-valued function g is said to be a biasing or importance sampling function if $g \in C^{1,2}([0, T] \times \mathbb{R}^d)$ and the Novikov condition holds, namely

$$\mathbb{E}_{x_0} \left[e^{\frac{1}{2} \int_0^T |\sigma^T \nabla g|^2(\tilde{X}_t^\varepsilon) dt} \right] < +\infty. \tag{6}$$

Under Assumption 2, the Girsanov change of measure between the paths $(X_t^\varepsilon)_{t \in [0, T]}$ and $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ (see ([24], Ch. 3, Prop. 5.12) or [29]) provides

$$A^\varepsilon = \mathbb{E}_{x_0} \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right] = \mathbb{E}_{x_0} \left[e^{\frac{1}{\varepsilon} f(\tilde{X}_T^\varepsilon) - \frac{1}{2\varepsilon} \int_0^T |\sigma^T \nabla g|^2(t, \tilde{X}_t^\varepsilon) dt - \frac{1}{\sqrt{\varepsilon}} \int_0^T \sigma^T \nabla g(t, \tilde{X}_t^\varepsilon) dB_t} \right]. \tag{7}$$

We next use Itô formula over a trajectory of $(\tilde{X}_t^\varepsilon)_{t \geq 0}$ with the generator (4) (since $g \in C^{1,2}([0, T] \times \mathbb{R}^d)$):

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

Integrating in time and dividing by ε , this equation becomes

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

Inserting the above equation into (7) we finally obtain

$$A^\varepsilon = \mathbb{E}_{x_0} \left[\exp \left(\frac{1}{\varepsilon} \left[f(\tilde{X}_T^\varepsilon) + g(0, x_0) - g(T, \tilde{X}_T^\varepsilon) \right] + \frac{1}{\varepsilon} \int_0^T (\partial_t g + \mathcal{H}^\varepsilon g)(t, \tilde{X}_t^\varepsilon) dt \right) \right], \tag{8}$$

where we introduced the Hamilton–Jacobi–Bellman nonlinear differential operator \mathcal{H}^ε defined by:

$$\forall g \in C^2(\mathbb{R}^d), \quad \mathcal{H}^\varepsilon g = b \cdot \nabla g + \frac{\varepsilon}{2} D : \nabla^2 g + \frac{1}{2} |\sigma^T \nabla g|^2.$$

One can easily check that \mathcal{H}^ε is actually the logarithmic transform of \mathcal{L}^ε , namely

$$\forall g \in C^2(\mathbb{R}^d), \quad \mathcal{H}^\varepsilon g = \varepsilon e^{-g/\varepsilon} \mathcal{L}^\varepsilon e^{g/\varepsilon}. \tag{9}$$

This logarithmic transform is a standard tool for studying partial differential equations related to control problems, see for instance [10, 17, 28].

In what follows, for a function g to be specified and satisfying Assumption 2, we denote by

$$\mathcal{G}^\varepsilon = \exp \left(\frac{1}{\varepsilon} \left[g(0, x_0) - g(T, \tilde{X}_T^\varepsilon) \right] + \frac{1}{\varepsilon} \int_0^T (\partial_t g + \mathcal{H}^\varepsilon g)(t, \tilde{X}_t^\varepsilon) dt \right)$$

the Girsanov weight of the path change of measure. Therefore (8) rewrites:

$$A^\varepsilon = \mathbb{E}_{x_0} \left[e^{\frac{1}{\varepsilon} f(\tilde{X}_T^\varepsilon)} \mathcal{G}^\varepsilon \right].$$

As a consequence of this computation, the estimator of A^ε based on independent realizations of the random variable

$$\tilde{\mathcal{A}}^\varepsilon = e^{\frac{1}{\varepsilon} f(\tilde{X}_T^\varepsilon)} \mathcal{G}^\varepsilon \tag{10}$$

provides an unbiased estimator of A^ε , *i.e.* it holds $\mathbb{E}_{x_0}[\tilde{\mathcal{A}}^\varepsilon] = \mathbb{E}_{x_0}[A^\varepsilon] = A^\varepsilon$.

The goal of importance sampling is to find the best biasing function g to minimize the relative variance of the tilted estimator $\tilde{\mathcal{A}}^\varepsilon$. In order to provide a criterion for the search of the best tilting g , we recall that the relative variance of the tilted estimator (see for instance [31] and references therein) is defined by

$$\frac{\mathbb{E}_{x_0} \left[\left(e^{\frac{1}{\varepsilon} f(\tilde{X}_T^\varepsilon)} \mathcal{G}^\varepsilon \right)^2 \right] - \mathbb{E}_{x_0} \left[e^{\frac{1}{\varepsilon} f(\tilde{X}_T^\varepsilon)} \mathcal{G}^\varepsilon \right]^2}{\mathbb{E}_{x_0} \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right]^2} = \frac{\mathbb{E}_{x_0} \left[e^{\frac{2}{\varepsilon} f(\tilde{X}_T^\varepsilon)} (\mathcal{G}^\varepsilon)^2 \right]}{\mathbb{E}_{x_0} \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right]^2} - 1.$$

We thus consider the ratio:

$$\rho(\varepsilon) = \frac{\mathbb{E}_{x_0} \left[e^{\frac{2}{\varepsilon} f(\tilde{X}_T^\varepsilon)} (\mathcal{G}^\varepsilon)^2 \right]}{\mathbb{E}_{x_0} \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right]^2},$$

which controls the evolution of the relative variance as $\varepsilon \rightarrow 0$. Since ρ typically grows exponentially as $\varepsilon \rightarrow 0$, we are interested in the evolution of ρ at log-scale, and we thus introduce the following standard definition [31].

Definition 1. *The relative log efficiency of an estimator is defined by*

$$R(\varepsilon) = \varepsilon \log \rho(\varepsilon). \tag{11}$$

We say that an estimator is k -log efficient if there is $k \geq 0$ such that

$$R(\varepsilon) = O(\varepsilon^k)$$

when $\varepsilon \rightarrow 0$.

In words, k -log efficiency describes a class of estimators with a small relative variance at exponential scale, in the low temperature regime. Although the notation k suggests that k is an integer (which will be the case in our examples), this does not need to be the case in general. Note also that, when $0 < k < 1$, the relative log efficiency goes to zero but the relative variance might not even be bounded (in the case where $R(\varepsilon)/\varepsilon \rightarrow +\infty$), it is only bounded by a term growing at a subexponential rate. We are thus mostly interested in the case $k \geq 1$. In this situation, there exists a constant $C > 0$ such that

$$0 \leq \rho(\varepsilon) - 1 \leq C\varepsilon^{k-1},$$

hence the relative variance is under control. As a result, from Definition 1, for a given small ε , an estimator with a larger k is likely to have a variance lower by one order of magnitude. This is what we observe in the numerical simulations performed in Section 4.4. Note however that reducing variance by improving log efficiency should not be done at the cost of too much additional computational time.

2.3. The optimal control

There is actually a simple way to find a solution to the problem of minimizing the relative variance of $\tilde{\mathcal{A}}^\varepsilon$. Consider for this the Girsanov formula (8): all the terms are stochastic except $g(0, x_0)$, which is deterministic. It is thus natural to set $g = g^\varepsilon$ solution to

$$\begin{cases} (\partial_t g^\varepsilon + \mathcal{H}^\varepsilon g^\varepsilon)(t, x) = 0, & \forall (t, x) \in [0, T) \times \mathbb{R}^d, \\ g^\varepsilon(T, x) = f(x), & \forall x \in \mathbb{R}^d. \end{cases} \tag{12}$$

If we assume that the solution to the above equation is well-defined, having $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ biased with g^ε indeed makes the random variable $\tilde{\mathcal{A}}^\varepsilon$ deterministic, hence the associated estimator has zero variance. Actually, under Assumption 1, standard parabolic estimates show that there is a unique smooth solution g^ε to (12) for all $\varepsilon > 0$, which solves the problem of variance reduction.

Moreover, inserting g^ε into (8), we see that

$$A^\varepsilon = e^{\frac{g^\varepsilon(0, x_0)}{\varepsilon}},$$

and immediately

$$Z^\varepsilon = g^\varepsilon(0, x_0). \tag{13}$$

In other words, the value of the optimal control at the starting point of the dynamics provides the free energy of the system.

Remark 1 (Stochastic solutions). It is interesting at this point to observe that the condition for g^ε to solve (12) everywhere is too stringent to minimize the variance of $\tilde{\mathcal{A}}^\varepsilon$. It is indeed sufficient that

$$\begin{cases} (\partial_t g^\varepsilon + \mathcal{H}^\varepsilon g^\varepsilon)(t, \tilde{X}_t^\varepsilon) = 0, \\ g^\varepsilon(T, \tilde{X}_T^\varepsilon) = f(\tilde{X}_T^\varepsilon), \end{cases} \tag{14}$$

for almost all $t \in [0, T]$ and almost surely with respect to the process $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ defined in (5) when the bias is g^ε itself. This notion of solution based on a stochastic representation is actually presented in details in the recent paper [26]. While ([26], Thm. 5.9) is motivated by considering non-smooth coefficients and solutions for the HJB equation through extended generators, it motivated the current study for its relation with variance reduction. In other words, it suggests that the optimal control only needs to be well approximated around appropriately titled trajectories (and not everywhere in space) for relative variance to be reduced. This is the basis of the notion of stochastic viscosity approximation that we propose in Section 3.1.

3. STOCHASTIC VISCOSITY APPROXIMATIONS

We now present the main results of the paper: the definition of stochastic viscosity approximations in Section 3.1, and the associated variance reduction property in Section 3.2.

3.1. Definition

The goal of this work is to find criteria for a function g to reduce the variance of the estimator (10). For this, we saw in Section 2.3 that the optimal solution to the variance reduction problem is the solution to a HJB equation, whose associated estimator has zero variance. Since the solution to this equation is in general unknown, it is natural to try to approximate it.

There exist a series of works dedicated to approximating Hamilton–Jacobi PDEs, typically by considering decentered finite difference schemes and assessing convergence through a Lebesgue or Sobolev distance to the exact solution [2, 3, 6, 27, 30]. Such convergence criteria rely on non-local quantities, and they are a priori not related to the problem of variance reduction. In other word, there is no one-to-one correspondance between the approximation of the HJB equation and relative variance of the associated estimator. Moreover, we are often interested in situations where the dimension d of the system is large, so that finite difference schemes are typically not applicable. In this context of stochastic optimal control in high dimension, a series of more specific techniques have also been designed, for instance [11, 23, 31, 32]; yet the same issue of proving that variance reduction is achieved for a given approximation remains – it may even deteriorate the associated estimator [19].

We now propose a new notion of approximation to the HJB problem (12). Such approximate solutions will be relevant to consider the variance reduction properties of \tilde{A}^ε in the low temperature regime. This is the main contribution of this paper. For this, we follow the stochastic solution approach discussed in Remark 1 (see in particular [26] and references therein), which suggests for g to be an approximate solution of (12) around tilted trajectories.

Definition 1 (Stochastic viscosity approximation). Let Assumption 1 hold, g^ε be the solution to (12), and $x_0 \in \mathbb{R}^d$ be the initial condition defining A^ε in (3). Let \hat{g}^ε be a function satisfying Assumption 2 with $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ the process tilted by \hat{g}^ε through (5). We say that \hat{g}^ε is a k -stochastic viscosity approximation of (12) if there exist $k, \varepsilon_0 > 0$, a family of \mathbb{R}^d -valued adapted processes $(Q_t^\varepsilon)_{t \in [0, T]}$ for $\varepsilon \in (0, \varepsilon_0]$ (residual process), and a family of real-valued random variables Q_f^ε for $\varepsilon \in (0, \varepsilon_0]$ (residual terminal condition) such that the following conditions hold:

- Approximate HJB equation along tilted paths: for any $\varepsilon \in (0, \varepsilon_0]$,

$$\begin{cases} (\partial_t \hat{g}^\varepsilon + \mathcal{H}^\varepsilon \hat{g}^\varepsilon)(t, \tilde{X}_t^\varepsilon) = \varepsilon^k Q_t^\varepsilon, \\ \hat{g}^\varepsilon(T, \tilde{X}_T^\varepsilon) = f(\tilde{X}_T^\varepsilon) + \varepsilon^k Q_f^\varepsilon. \end{cases} \tag{15}$$

for almost all $t \in [0, T]$ and almost surely with respect to $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$.

- Bound: there exists a real number $C_Q \geq 0$ independent of ε such that the family of random variables

$$Q^\varepsilon = Q_f^\varepsilon + \int_0^T Q_t^\varepsilon dt \tag{16}$$

satisfies

$$\forall \varepsilon \in (0, \varepsilon_0], \quad \left| \log \mathbb{E} \left[e^{2Q^\varepsilon} \right] \right| \leq 2C_Q. \tag{17}$$

Let us provide some comments on Definition 1. The main condition (15) is for \hat{g}^ε to be an approximate solution of (12) (including the boundary conditions) along trajectories tilted by \hat{g}^ε itself. By approximate we mean that evaluating the Hamilton–Jacobi operator along such stochastic trajectories is an appropriately bounded stochastic process (which we may call residual process) scaled by ε^k . Since the error term vanishes in the small temperature limit, there is a notion of viscosity involved, as it is presented for instance in Chapter 2, Theorem 3.1 from [18]. We thus understand the naming stochastic viscosity approximation, which we abbreviate as SVA.

Remark 2. – We ask for \hat{g}^ε to be of class $C^{1,2}$ for computing derivatives in the classical sense, but using the notion of stochastic derivative ([26], Sect. 2) we could consider less regular approximations (typically almost everywhere differentiable functions, a situation that may arise for piecewise defined functions). This is a nice feature of SVAs.

- We also emphasize that Definition 1 applies way beyond the scope of Assumption 1, typically as soon as the stochastic process (1) and its associated Fokker–Planck equation make sense.
- We claim by no mean that the bound (17) is optimal in any sense: it is simply convenient for the proofs below, and seems reasonable to check at least in simple situations.
- Finally, we assume that (15) holds almost surely. However, it is likely that a weaker probabilistic sense is sufficient, for instance the equality could be satisfied in expectation only.

3.2. Relation to variance reduction

As explained in Section 2.3, we have defined our notion of stochastic viscosity approximation precisely in order to make the tilted estimator $\tilde{\mathcal{A}}^\varepsilon$ close to deterministic by reducing the random terms in the Girsanov weight to a small-noise factor. We first show that Definition 1 ensures a small temperature consistency of the estimated quantity.

Lemma 1. *Assume that \hat{g}^ε is a k -SVA of g^ε for some $k \geq 1$. Then it holds*

$$\hat{g}^\varepsilon(0, x_0) = g^\varepsilon(0, x_0) + O(\varepsilon^k). \tag{18}$$

Recalling (13), it holds $Z^\varepsilon = g^\varepsilon(0, x_0)$. Lemma 1 thus shows that a stochastic viscosity approximation provides an estimate of the quantity of interest through its initial condition according to

$$Z^\varepsilon = \hat{g}^\varepsilon(0, x_0) + O(\varepsilon^k). \tag{19}$$

The proof is as follows.

Proof of Lemma 1. We use the Girsanov theorem (8) to write:

$$\begin{aligned} e^{\frac{1}{\varepsilon}g^\varepsilon(0,x_0)} &= \mathbb{E}_{x_0} \left[e^{\frac{1}{\varepsilon}f(X_T^\varepsilon)} \right] = \mathbb{E}_{x_0} \left[\exp \left(\frac{1}{\varepsilon} \left[f(\tilde{X}_T^\varepsilon) + \hat{g}^\varepsilon(0, x_0) - \hat{g}^\varepsilon(T, \tilde{X}_T^\varepsilon) \right] \right. \right. \\ &\quad \left. \left. + \frac{1}{\varepsilon} \int_0^T (\partial_t \hat{g}^\varepsilon + \mathcal{H}^\varepsilon \hat{g}^\varepsilon)(t, \tilde{X}_t^\varepsilon) dt \right) \right], \end{aligned}$$

where $(\tilde{X}_t)_{t \in [0, T]}$ is the process (5) for the bias function \hat{g}^ε . Using (15) we obtain

$$e^{\frac{1}{\varepsilon}g^\varepsilon(0,x_0)} = e^{\frac{1}{\varepsilon}\hat{g}^\varepsilon(0,x_0)} \mathbb{E}_{x_0} \left[e^{\varepsilon^{k-1}(Q_F^\varepsilon + \int_0^T Q_t^\varepsilon dt)} \right] = e^{\frac{1}{\varepsilon}\hat{g}^\varepsilon(0,x_0)} \mathbb{E}_{x_0} \left[\left(e^{2Q^\varepsilon} \right)^{\frac{\varepsilon^{k-1}}{2}} \right].$$

Given that $k \geq 1$, the function $x \in \mathbb{R}_+ \mapsto x^{\varepsilon^{k-1}/2}$ is concave for $\varepsilon \leq 1$ (we may replace ε_0 by $\min(\varepsilon_0, 1)$ if necessary). As a result, we can use Jensen’s inequality combined with (17) to reach

$$|g^\varepsilon(0, x_0) - \hat{g}^\varepsilon(t, x_0)| \leq C_Q \varepsilon^k,$$

which is the desired conclusion. □

Using Definition 1 of log-efficiency, we can now state the variance reduction property of a k -stochastic viscosity approximation.

Theorem 1. *Assume that \hat{g}^ε is a k -SVA of g^ε for some $k \geq 1$. Then the associated estimator $\tilde{\mathcal{A}}^\varepsilon$ is k -log efficient.*

This is our main result associated to the definition of stochastic viscosity approximation. Since Lemma 1 already provides the estimate (19) of the free energy, Theorem 1 ensures that the correction terms can be estimated by Monte Carlo simulation with a bounded or vanishing variance. Contrarily to Lebesgue-norm criteria that may not ensure good variance reduction properties, we here draw a very clear connection between SVA and log-efficiency of the associated estimator. Let us present the proof, which is rather straightforward.

Proof of Theorem 1. For this proof, we write

$$\hat{\mathcal{G}}^\varepsilon = \exp\left(\frac{1}{\varepsilon} \left[\hat{g}^\varepsilon(0, x_0) - \hat{g}^\varepsilon(T, \tilde{X}_T^\varepsilon) \right] + \frac{1}{\varepsilon} \int_0^T (\partial_t \hat{g}^\varepsilon + \mathcal{H}^\varepsilon \hat{g}^\varepsilon)(t, \tilde{X}_t^\varepsilon) dt\right)$$

for the Girsanov weight associated with the function \hat{g}^ε . Let us consider first the quantity

$$\mathbb{E}_{x_0} \left[e^{\frac{2}{\varepsilon} f(\tilde{X}_T^\varepsilon)} (\hat{\mathcal{G}}^\varepsilon)^2 \right] = e^{\frac{2}{\varepsilon} \hat{g}^\varepsilon(0, x_0)} \mathbb{E}_{x_0} \left[e^{\frac{2}{\varepsilon} (f(\tilde{X}_T^\varepsilon) - \hat{g}^\varepsilon(T, \tilde{X}_T^\varepsilon) + \int_0^T (\partial_t \hat{g}^\varepsilon + \mathcal{H}^\varepsilon \hat{g}^\varepsilon)(t, \tilde{X}_t^\varepsilon) dt)} \right].$$

The approximate HJB equation (15) (including the terminal condition) in Definition 1 implies again that

$$\mathbb{E}_{x_0} \left[e^{\frac{2}{\varepsilon} (f(\tilde{X}_T^\varepsilon) - \hat{g}^\varepsilon(T, \tilde{X}_T^\varepsilon) + \int_0^T (\partial_t \hat{g}^\varepsilon + \mathcal{H}^\varepsilon \hat{g}^\varepsilon)(t, \tilde{X}_t^\varepsilon) dt)} \right] = \mathbb{E}_{x_0} \left[e^{\frac{2}{\varepsilon} (\varepsilon^k Q_F^\varepsilon + \varepsilon^k \int_0^T Q_t^\varepsilon dt)} \right].$$

Using the consistency relation (18) proved in Lemma 1 and recalling the expression (16) of the residual Q^ε , we obtain

$$\mathbb{E}_{x_0} \left[e^{\frac{2}{\varepsilon} f(\tilde{X}_T^\varepsilon)} (\hat{\mathcal{G}}^\varepsilon)^2 \right] = e^{\frac{2}{\varepsilon} (g^\varepsilon(0, x_0) + O(\varepsilon^k))} \mathbb{E}_{x_0} \left[e^{2\varepsilon^{k-1} Q^\varepsilon} \right].$$

Since $k \geq 1$, we use again the concavity of the function $x \in \mathbb{R}_+ \mapsto x^{\varepsilon^{k-1}}$ for $\varepsilon \leq 1$. Together with Jensen’s inequality and (17) we obtain, for $\varepsilon \in (0, \varepsilon_0]$:

$$\mathbb{E}_{x_0} \left[e^{2\varepsilon^{k-1} Q^\varepsilon} \right] \leq \mathbb{E}_{x_0} \left[e^{2Q^\varepsilon} \right]^{\varepsilon^{k-1}} \leq e^{2\varepsilon^{k-1} C_Q}.$$

Finally, by (13), it holds $\mathbb{E} \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right]^2 = e^{\frac{2}{\varepsilon} g^\varepsilon(0, x_0)}$, so that we obtain the bound

$$\rho(\varepsilon) \leq \frac{e^{\frac{2}{\varepsilon} (g^\varepsilon(0, x_0) + O(\varepsilon^k))} e^{2\varepsilon^{k-1} C_Q}}{e^{\frac{2}{\varepsilon} g^\varepsilon(0, x_0)}} = e^{2C_Q \varepsilon^{k-1} + O(\varepsilon^{k-1})}.$$

This implies that

$$R(\varepsilon) = \varepsilon \log \rho(\varepsilon) = O(\varepsilon^k)$$

as $\varepsilon \rightarrow 0$, which is the desired result. □

Note that we cover the case $k \geq 1$ for simplicity and because this is the main situation of interest (since relative variance is a priori not even bounded for $k < 1$), but one can prove a similar result for $k \in [0, 1)$ under stronger assumptions, for instance by assuming that Q^ε is bounded. As discussed after Definition 1, in the particular case where $k \geq 1$, the relative variance $\rho(\varepsilon) - 1$ is bounded, hence we ensure that the variance of numerical simulations remains finite even in the vanishing temperature regime, which is what we are looking for in practice.

4. INSTANTON AND STOCHASTIC VISCOSITY APPROXIMATIONS

In this section, we propose to construct a 1-SVA and a 2-SVA from the theory of low temperature reaction paths called instantons [12, 20] and we provide illustrative numerical results. We first present and motivate the instanton dynamics in Section 4.1. We then show in Section 4.2 that the earlier proposed sampling technique [20] relying on these equations provides a 1-SVA of the optimal control under appropriate assumptions. It is thus associated with a 1-log efficient estimator. This provides a mathematical variance reduction property for this importance sampling scheme under geometric conditions, which had not been proved so far. In Section 4.3 we consider the next order expansion proposed in [12], which is a 2-SVA of the optimal control, hence improving on variance reduction. We finally propose in Section 4.4 a simple numerical application to illustrate the validity of our predictions.

4.1. Low temperature asymptotics

We start by presenting low temperature asymptotics concerning the dynamics that we will use to practically build SVAs. Recalling that $\mathcal{H}^0\varphi = b \cdot \nabla\varphi + |\sigma^T \nabla\varphi|^2/2$, the zero temperature HJB equation reads

$$\begin{cases} (\partial_t g^0 + \mathcal{H}^0 g^0)(t, x) = 0, & \forall (t, x) \in [0, T] \times \mathbb{R}^d, \\ g^0(T, x) = f(x), & \forall x \in \mathbb{R}^d. \end{cases} \quad (20)$$

This is a first order nonlinear PDE for which a smooth solution often does not exist, because of the lack of diffusive regularization. However, a viscosity solution can most often be described through the relation ([18], Ch. 2, Thm. 3.1):

$$\forall (t, x) \in [0, T] \times \mathbb{R}^d, \quad g^0(t, x) = \lim_{\varepsilon \rightarrow 0} g^\varepsilon(t, x).$$

In some cases, mostly when the function g^0 is smooth, we can describe the solution through its characteristics system. The same methodology applied in a large deviations context leads to the instanton dynamics, which is described by a forward-backward system [12, 20] with \mathbb{R}^d -valued position $(\phi_t)_{t \in [0, T]}$ and momentum $(\theta_t)_{t \in [0, T]}$ solution to

$$\begin{cases} \dot{\phi}_t = b(\phi_t) + D\theta_t, & \phi_0 = x_0, \\ \dot{\theta}_t = -(\nabla b)(\phi_t) \cdot \theta_t, & \theta_T = \nabla f(\phi_T). \end{cases} \quad (21)$$

The above formula is a definition of the instanton, which can be derived by minimizing the Freidlin–Wentzell rate function under constraints (see [20] and references therein). This dynamics describes the path realizing fluctuations in the low temperature limit.

There are various ways to introduce the instanton dynamics (21). Let us provide a maybe unusual, PDE-oriented motivation for this object, by introducing the function

$$\forall (t, x) \in [0, T] \times \mathbb{R}^d, \quad \hat{g}_1^\varepsilon(t, x) = f(\phi_T) - \frac{1}{2} \int_t^T \theta_s \cdot D\theta_s \, ds + \theta_t \cdot (x - \phi_t). \quad (22)$$

Note that we keep the ε dependency on \hat{g}_1^ε although it does not depend on ε , because such an approximation may depend on ε in general. This function is made of three parts: a constant term, a time-dependent only term, and a linear term in the variable $(x - \phi_t)$. We may thus call \hat{g}_1^ε a time-inhomogeneous first order polynomial in the variable $(x - \phi_t)$.

To understand better this function, we first see that $\hat{g}_1^\varepsilon(T, x) = f(\phi_T) + \nabla f(\phi_T) \cdot (x - \phi_T)$, which is a linearization of f around the instanton's terminal point ϕ_T . Let us now compute the value of the HJB operator $\partial_t + \mathcal{H}^\varepsilon$ over \hat{g}_1^ε . Noting that $\nabla \hat{g}_1^\varepsilon(t, x) = \theta_t$ and $\theta_t \cdot D\theta_t = |\sigma^T \theta_t|^2$ we get, for any $(t, x) \in [0, T] \times \mathbb{R}^d$:

$$\begin{aligned} (\partial_t \hat{g}_1^\varepsilon + \mathcal{H}^\varepsilon \hat{g}_1^\varepsilon)(t, x) &= \frac{1}{2} |\sigma^T \theta_t|^2 + \dot{\theta}_t \cdot (x - \phi_t) - \theta_t \cdot \dot{\phi}_t + b(x) \cdot \theta_t + \frac{1}{2} |\sigma^T \theta_t|^2 \\ &= |\sigma^T \theta_t|^2 - \nabla b(\phi_t) \cdot (x - \phi_t) - \theta_t \cdot (b(\phi) + D\theta_t) + b(x) \cdot \theta_t \\ &= \theta_t \cdot (b(x) - b(\phi_t) - \nabla b(\phi_t) \cdot (x - \phi_t)) = O\left((x - \phi_t)^2\right), \end{aligned}$$

where we used (21) for the second line, while the last inequality is derived by Taylor expansion. We actually obtain from the above computation that

$$\begin{cases} (\partial_t \hat{g}_1^\varepsilon + \mathcal{H}^\varepsilon \hat{g}_1^\varepsilon)(t, \phi_t) = 0, & \forall t \in [0, T], \\ \hat{g}_1^\varepsilon(T, \phi_T) = f(\phi_T). \end{cases} \tag{23}$$

In other words, \hat{g}_1^ε solves the HJB problem (12) along the instanton.

Since our discussion in Remark 1 suggests that importance sampling only requires approximation of HJB around some trajectories and $(\phi_t)_{t \in [0, T]}$ describes the most likely fluctuation path as $\varepsilon \rightarrow 0$, we understand that the ansatz (22) can serve as a first guess for approximating the solution g^ε of the finite temperature equation (12) along most likely fluctuation paths, and hence reduce variance.

4.2. Instanton bias

We consider the approximation (22) under the following assumption (see Rmk. 3 for further discussion).

Assumption 3. The dynamics (21) admits a unique solution $(\phi_t, \theta_t)_{t \in [0, T]}$, which is of class $C^1([0, T])^d \times C^1([0, T])^d$.

Under Assumption 3, the function \hat{g}_1^ε is well-defined and we denote by $(\tilde{X}_t^{\varepsilon, 1})_{t \in [0, T]}$ the dynamics biased by \hat{g}_1^ε as in (5). Compared to the control proposed in the previous work [12], we add a constant and a time-only dependent components (the first two terms in (22)). This is not crucial from a practical perspective since the biasing force $\nabla \hat{g}_1^\varepsilon$ does not depend on these terms, but it makes the entire approximation argument more elegant.

Before proving that \hat{g}_1^ε is a 1-stochastic viscosity approximation of g^ε , we first show that a trajectory tilted by \hat{g}_1^ε is a perturbation of the instanton – very much in the spirit of the expansions that can be found in Chapter 2 from [18].

Lemma 2. *Under Assumptions 1 and 3, there exist an adapted process $(Y_t^\varepsilon)_{t \in [0, T]}$ and a constant $C_T > 0$ such that*

$$\forall t \in [0, T], \quad \tilde{X}_t^{\varepsilon, 1} = \phi_t + \sqrt{\varepsilon} Y_t^\varepsilon, \tag{24}$$

and for all $t \in [0, T]$ the following bound holds

$$\forall \varepsilon > 0, \quad |Y_t^\varepsilon| \leq C_T \sup_{0 \leq s \leq T} |\sigma^T B_s|, \tag{25}$$

almost surely with respect to realizations of the Brownian motion.

The proof follows a standard Gronwall argument.

Proof of Lemma 2. Let $\varepsilon > 0$ be arbitrary. Since the solution to (21) is unique, we can define for all $t \in [0, T]$ the adapted process³

$$Y_t^\varepsilon = \frac{\tilde{X}_t^{\varepsilon,1} - \phi_t}{\sqrt{\varepsilon}}. \tag{26}$$

Since $\nabla g(t, x) = \theta_t$ for all $(t, x) \in [0, T] \times \mathbb{R}^d$, Itô formula shows that

$$\begin{aligned} \sqrt{\varepsilon} dY_t^\varepsilon &= b\left(\tilde{X}_t^{\varepsilon,1}\right) + D\theta_t dt + \sqrt{\varepsilon}\sigma^T dB_t - \dot{\phi}_t dt \\ &= \left(b\left(\tilde{X}_t^{\varepsilon,1}\right) - b(\phi_t)\right) + \sqrt{\varepsilon}\sigma^T dB_t, \end{aligned}$$

where we used (21). Integrating now in time and taking the absolute value leads to

$$\forall t \in [0, T], \quad \sqrt{\varepsilon}|Y_t^\varepsilon - Y_0^\varepsilon| \leq \int_0^t \left|b\left(\tilde{X}_s^{\varepsilon,1}\right) - b(\phi_s)\right| ds + \sqrt{\varepsilon}|\sigma^T B_t|.$$

Recalling that $Y_0^\varepsilon = 0$ and that b is globally Lipschitz with constant C_{lip} (by Asm. 1), we obtain

$$\forall t \in [0, T], \quad \sqrt{\varepsilon}|Y_t^\varepsilon| \leq \int_0^t C_{\text{lip}} \left|\tilde{X}_s^{\varepsilon,1} - \phi_s\right| ds + \sqrt{\varepsilon}|\sigma^T B_t|,$$

which rewrites

$$\forall t \in [0, T], \quad |Y_t^\varepsilon| \leq C_{\text{lip}} \int_0^t |Y_s^\varepsilon| ds + |\sigma^T B_t|.$$

Gronwall lemma then implies that

$$\forall t \in [0, T], \quad |Y_t^\varepsilon| \leq \sup_{0 \leq s \leq T} |\sigma^T B_s| e^{TC_{\text{lip}}},$$

which holds for any continuous realization of the Brownian motion, hence almost surely, and for any $\varepsilon > 0$. We thus have the bound (25), while (24) follows from the definition (26) of Y_t^ε . \square

Proving that \hat{g}_1^ε is a stochastic viscosity approximation of g^ε is now a matter of Taylor expansion under some technical conditions.

Theorem 2. *Let λ be the largest eigenvalue of D . Let moreover Assumptions 1 and 3 hold and suppose that there exist $\delta_f, \delta_b \geq 0$ such that*

$$\delta_f + \delta_b T < \frac{1}{\lambda TC_T^2}, \tag{27}$$

and:

– The function f satisfies

$$\forall y \in \mathbb{R}^d, \quad \sup_{x \in \mathbb{R}^d} y \cdot \nabla^2 f(x)y \leq \delta_f |y|^2. \tag{28}$$

– The field b satisfies

$$\forall y \in \mathbb{R}^d, \quad \forall t \in [0, T], \quad \sup_{x \in \mathbb{R}^d} y \cdot (\theta_t \cdot \nabla^2 b(x))y \leq \delta_b |y|^2. \tag{29}$$

Then the function \hat{g}_1^ε defined in (22) satisfies Assumption 2 and is a 1-SVA of the HJB problem. As a result, it defines a 1-log efficient estimator.

³We mention that in the case of multiple instantons solution to (21), one could consider defining several approximations centered on each instanton with their residual processes.

The proof of Theorem 2 is presented just below. We may interpret the geometric conditions (28) and (29) as complementary to Assumption 3. First, (28) asks for f to be “mostly concave”. Typically, we may expect problems to occur when f has several pronounced maxima, so that the maximum value of f can be reached in several ways. A similar condition (29) holds for the drift b : if we further assume that ∇f has nonnegative coordinates, then θ_t has nonnegative coordinates at least for small times, and so (29) is a concavity condition on b . Obviously, these conditions are quite drastic and could be relaxed if one proves additional stability properties of the process or has other information on the drift b (for instance bounded second derivative).

Proof of Theorem 2. We check the approximate HJB equation (15) with bound (17) by relying on Lemma 2. Consider the terminal condition first. Using the terminal value of θ_T in (21), we have

$$\hat{g}_1^\varepsilon(T, \tilde{X}_T^{\varepsilon,1}) = f(\phi_T) + \nabla f(\phi_T) \cdot (\tilde{X}_T^{\varepsilon,1} - \phi_T).$$

Taylor’s integral theorem applied at first order to f around ϕ_T shows, using Lemma 2, that we have almost surely

$$\begin{aligned} Q_f^\varepsilon &= \varepsilon^{-1} \left(f(\tilde{X}_T^{\varepsilon,1}) - \hat{g}_1^\varepsilon(T, \tilde{X}_T^{\varepsilon,1}) \right) = \varepsilon^{-1} \left(f(\tilde{X}_T^{\varepsilon,1}) - f(\phi_T) - \nabla f(\phi_T) \cdot (\tilde{X}_T^{\varepsilon,1} - \phi_T) \right) \\ &= \frac{1}{2} \int_0^1 (1-s) Y_T^\varepsilon \cdot \nabla^2 f(\phi_T + s\sqrt{\varepsilon}Y_T^\varepsilon) Y_T^\varepsilon ds. \end{aligned}$$

According to (28), it holds

$$Q_f^\varepsilon \leq \delta_f \frac{|Y_T^\varepsilon|^2}{4},$$

which we will use to control the residual process⁴.

Let us now check the dynamical part in (15). First, we easily compute

$$\begin{aligned} \partial_t \hat{g}_1^\varepsilon(t, x) &= \frac{1}{2} \theta_t \cdot D\theta_t + \dot{\theta}_t \cdot (x - \phi_t) - \theta_t \cdot \dot{\phi}_t, \\ \nabla \hat{g}_1^\varepsilon(t, x) &= \theta_t, \\ \nabla^2 \hat{g}_1^\varepsilon(t, x) &= 0. \end{aligned}$$

Therefore, using Lemma 2 and equation (21) we obtain

$$\begin{aligned} (\partial_t \hat{g}_1^\varepsilon + \mathcal{H}^\varepsilon \hat{g}_1^\varepsilon)(t, \tilde{X}_t^{\varepsilon,1}) &= \frac{1}{2} \theta_t \cdot D\theta_t - \sqrt{\varepsilon} \theta_t \cdot \nabla b(\phi_t) Y_t^\varepsilon - (b(\phi_t) + D\theta_t) \cdot \theta_t + b(\tilde{X}_t^{\varepsilon,1}) \cdot \theta_t + \frac{1}{2} \theta_t \cdot D\theta_t \\ &= -\sqrt{\varepsilon} \theta_t \cdot \nabla b(\phi_t) Y_t^\varepsilon + \left(b(\tilde{X}_t^{\varepsilon,1}) - b(\phi_t) \right) \cdot \theta_t \\ &= \theta_t \cdot \left(b(\tilde{X}_t^{\varepsilon,1}) - b(\phi_t) - \nabla b(\phi_t) (\tilde{X}_t^{\varepsilon,1} - \phi_t) \right). \end{aligned}$$

Finally, by defining the process

$$Q_t^\varepsilon = \frac{\theta_t}{\varepsilon} \cdot \left(b(\tilde{X}_t^{\varepsilon,1}) - b(\phi_t) - \nabla b(\phi_t) (\tilde{X}_t^{\varepsilon,1} - \phi_t) \right),$$

we have

$$(\partial_t \hat{g}_1^\varepsilon + \mathcal{H}^\varepsilon \hat{g}_1^\varepsilon)(t, \tilde{X}_t^{\varepsilon,1}) = \varepsilon Q_t^\varepsilon.$$

⁴Although we impose a constraint on $\nabla^2 f$ over the whole space, we see that it is essentially important to control this quantity around ϕ_T . One difficulty however comes from the fact that ϕ_T itself depends on the function f . This makes a general condition for f to satisfy the bound hard to state at this level of generality. The same is true for the condition on $\nabla^2 b$.

In order to obtain (17) and conclude, it remains to bound Q_t^ε . We apply again Taylor’s theorem at first order between $\tilde{X}_t^{\varepsilon,1}$ and ϕ_t for all time $t \geq 0$ and any continuous realization of $(\tilde{X}_t^\varepsilon)_{t \in [0,T]}$, now to the \mathbb{R}^d -valued function b , to obtain

$$\forall t \in [0, T], \quad Q_t^\varepsilon = Y_t^\varepsilon \cdot \left(\frac{1}{2} \int_0^1 (1-s) \theta_t \cdot \nabla^2 b(\phi_t + s\sqrt{\varepsilon} Y_t^\varepsilon) \, ds \right) Y_t^\varepsilon.$$

Finally, the condition (29) ensures that

$$Q_t^\varepsilon \leq \frac{\delta_b}{4} |Y_t^\varepsilon|^2,$$

and the total residual Q^ε defined in (16) satisfies:

$$Q^\varepsilon \leq \frac{\delta_f}{4} |Y_T^\varepsilon|^2 + \frac{\delta_b}{4} \int_0^T |Y_t^\varepsilon|^2 \, dt,$$

Using (25) in Lemma 2 then leads to

$$Q^\varepsilon \leq C_T^2 \sup_{0 \leq s \leq T} |\sigma^T B_s|^2 \left(\frac{\delta_f}{4} + \frac{\delta_b}{4} T \right).$$

In order to find conditions for the integration condition (17) to hold, we recall that $\lambda > 0$ is the largest eigenvalue of $D = \sigma \sigma^T$, and introduce

$$c = 2\lambda C_T^2 \left(\frac{\delta_f}{4} + \frac{\delta_b}{4} T \right).$$

Since $|\sigma^T B_s|^2 = B_s^T D B_s$, we have

$$2Q^\varepsilon \leq c \sup_{0 \leq s \leq T} |B_s|^2.$$

It is simple to show that

$$\sup_{0 \leq s \leq T} |B_s|^2 = \left(\sup_{0 \leq s \leq T} |B_s| \right)^2.$$

On the other hand, the law of $\sup_{0 \leq s \leq T} |B_s|$ is described by Formula 1.1.8 on page 250 from [4]:

$$\forall y, x \geq 0, \quad \mathbb{P}_x \left(\sup_{0 \leq s \leq T} |B_s| \geq y \right) = \sum_{k=-\infty}^{+\infty} (-1)^k \text{sign}(x + 2(k+1)y) \text{Erfc} \left(\frac{|x + (2k+1)y|}{\sqrt{2}T} \right), \quad (30)$$

where

$$\forall x \in \mathbb{R}, \quad \text{Erfc}(x) = \frac{2}{\sqrt{\pi}} \int_x^{+\infty} e^{-u^2} \, du.$$

Using the case $x = 0$ in (30), we conclude that

$$\mathbb{E} \left[e^{2Q^\varepsilon} \right] \leq \mathbb{E} \left[e^{c(\sup_{0 \leq s \leq T} |B_s|)^2} \right]$$

is finite independently of $\varepsilon > 0$ when $c < 1/(2T)$, which is equivalent to (27). As a result, (17) holds provided δ_f and δ_b are small enough for (27) to hold.

Since $\nabla \hat{g}_1^\varepsilon$ does not depend on the position variable, Novikov’s condition (6) holds and so does Assumption 2. As a result, all the conditions of Definition 1 are satisfied and \hat{g}_1^ε is a 1-SVA of the HJB problem (12). According to Theorem 1, the associated estimator is 1-log efficient, which concludes the proof. □

Remark 3 (About Assumption 3). While Assumptions 1 and 2 are of technical nature and could easily be relaxed, Assumption 3 is of geometrical nature, and has a crucial importance. It is easy to build examples for which all quantities are smooth, but where this assumption does not hold [19], typically for symmetry or convexity reasons. An example is to set $d = 1, b = 0, \sigma = 1, f(x) = \tanh(|x|)$ and $x_0 = 0$: in this case the instanton may maximize the final value of f either by going to the right or to the left for the same kinetic cost, so there is no uniqueness of the instanton. This is a situation similar to the one appearing when solving first order PDEs with the method of characteristics.

In this case, it is natural to divide the space \mathbb{R}^d into different regions where one instanton would dominate. An idea in this direction is to resort to the so-called *regions of strong regularity* defined in [16]. The idea is to split the space in N regions $(\Gamma_i)_{i=1}^N$ such that $[0, T] \times \mathbb{R}^d = \cup_{i=1}^N \Gamma_i$ with $\forall i \neq j, |\Gamma_i \cap \Gamma_j| = 0$, and such that the problem can be treated independently in each region. We would then define N instantons and their approximations $(\hat{g}^{\varepsilon,i})_{i=1}^N$, and a natural candidate approximation is

$$\hat{g}^\varepsilon(t, x) = \hat{g}^{\varepsilon,i}(t, x) \quad \text{if } (t, x) \in \Gamma_i \quad \text{and} \quad x \notin \Gamma_j, \quad \forall j \neq i.$$

Since we assume the boundaries of these sets to have zero Lebesgue measure, this function is well-defined almost everywhere. Of course, using a finite polynomial in each region, \hat{g}^ε could have singularities on the boundaries of the sets, which might be a side effect of going through a low temperature expansion.

4.3. Riccati bias

One motivation for the current work was to better understand the approximation proposed in [12]. The point of the authors was to consider \hat{g}_1^ε as the first term of a time-dependent Taylor expansion in the argument $(x - \phi_t)$, which naturally suggests to continue the series with higher orders. Based on this earlier work ([12], eq. (21)) we introduce the following quadratic approximation, for a $R^{d \times d}$ -valued dynamics $(K_t)_{t \in [0, T]}$,

$$\hat{g}_2^\varepsilon(t, x) = f(\phi_T) - \frac{1}{2} \int_t^T \theta_s \cdot D\theta_s \, ds + \frac{\varepsilon}{2} \int_t^T D : K_s \, ds + \theta_t \cdot (x - \phi_t) + \frac{1}{2}(x - \phi_t) \cdot K_t(x - \phi_t). \tag{31}$$

Like for the expansion performed in the proof of Theorem 2, we can consider the above ansatz and identify the equations associated with each term to obtain a 2-SVA. This is exactly the computation performed in Section 3 and Appendix B from [12], which we don't reproduce here for conciseness. The obtained result is that we shall take $(K_t)_{t \in [0, T]}$ solution to the following Riccati system:

$$\begin{cases} \dot{K}_t + (\nabla b)^T K_t + K_t^T \nabla b + \nabla^2 b \cdot \theta_t + K_t^T D K_t &= 0, \\ K_T = \nabla^2 f(\phi_T). \end{cases} \tag{32}$$

By imposing $(\phi_t, \theta_t)_{t \in [0, T]}$ to be solution to (21) and $(K_t)_{t \in [0, T]}$ to be solution to (32) (assuming well-posedness of these systems, which is true at least for small times) while assuming appropriate bounds on $\nabla^3 f$ and $\nabla^3 b$, we can thus show that \hat{g}_2^ε is a 2-SVA of g^ε . As mentioned above, the computation performed in Section 3 and Appendix B from [12] follows the same Taylor expansion as for Theorem 2 but one order beyond, which naturally leads to a Riccati equation. Since \hat{g}_2^ε defines a 2-SVA, the associated estimator is 2-log efficient. This entails that the second order approximation proposed in [12] indeed improves on variance reduction compared to the first order one earlier proposed in [20]. For a given ε , we expect to find an error one order of magnitude smaller for a 2-SVA compared to a 1-SVA, which is a major improvement.

4.4. Numerical application

We now illustrate our theoretical findings by proposing a simple situation in which all our conditions are satisfied and our results apply. We then perform a numerical simulation that confirms our predictions.

The system is a simple one dimensional Ornstein-Uhlenbeck process, *i.e.* $d = 1, \sigma = 1$ and $b(x) = -x$. In order to propose a smoothed version of a probability like $\mathbb{P}(X_T^\varepsilon > 1)$, we set $f(x) = -(x - 2)^4/4$. This weight




FIGURE 1. *Left*: instanton, momentum and Riccati matrix (here scalar values). *Right*: evolution of $R(\varepsilon)$ as defined in (11) with ε in logarithmic scale, for the unbiased estimator and the dynamics biased by \hat{g}_1^ε and \hat{g}_2^ε . Linear and quadratic decays are plotted in dashed lines.

function gives larger importance to those trajectories ending around 2. We also set $x_0 = -1$, so tilted trajectories typically relax towards 0 before being pushed away to larger positive values. This system is non-trivial while satisfying Assumption 1. In this case the solution to the instanton equation (21) is actually explicit, while the Riccati equation (32) is stable and can be integrated with a simple Euler scheme. For more general dynamics, the instanton can be computed with the algorithm described in Section III A from [20] (which we actually use in practice). Finally, $f''(x) \leq 0$ and $b''(x) = 0$ for all $x \in \mathbb{R}$, so the conditions (28) and (29) are satisfied and Theorem 2 applies. As a result, \hat{g}_1^ε defines a 1-SVA of the system. Similarly, since the system is quite simple, we can show that \hat{g}_2^ε is a 2-SVA.

In order to perform numerical simulations, we set $T = 5$ and discretize the SDE, the instanton and Riccati equations with a time step $\Delta t = 5 \times 10^{-3}$. We numerically assess that this value is small enough to neglect the bias due to time discretization compared to the relative variance we are interested in (not shown). We draw 5×10^6 trajectories to estimate the relative variance ratio $R(\varepsilon)$ defined in (11) for a series of values of ε . The left panel in Figure 1 shows the evolution of the instanton and Riccati terms with time. While momentum adds an upward drift, the negative Riccati term forces the process to remain around the instanton. On the right panel we show the evolution of $R(\varepsilon)$ with ε for the different estimators. We observe that $R(\varepsilon)$ is roughly constant without biasing, meaning relative variance grows exponentially in this case. On the other hand, the linear and quadratic decay at log scale for the first and second order biases are respectively in accordance with the results of Theorem 2 and the discussion in Section 4.3.

5. DISCUSSION

This work is concerned with the issue of variance reduction for the Monte Carlo estimation of exponential-like expectations, as they often arise in statistical physics and other areas. Our main goal was to design a framework to assess whether a given approximation of the optimal Hamilton–Jacobi control indeed reduces relative variance as temperature becomes small, and at which scale.

To achieve this, we introduce the notion of stochastic viscosity approximation (SVA), inspired by recent developments on the theory of Feynman–Kac partial differential equations [26]. We believe our definition is meaningful: a SVA approximately solves the Hamilton–Jacobi–Bellman equation along relevant tilted paths. This is precisely what is needed in the Girsanov theorem to make stochastic terms small. We therefore prove an associated variance reduction property. This definition comes in sharp contrast with standard approximation techniques criteria, which typically control a non-local error to assess convergence as a mesh size decays to zero.

Here we show that even an approximation that does not depend on temperature or any parameter going to zero can have a vanishing log-efficiency in the zero temperature limit!

Since this study was in part motivated by earlier heuristics on importance sampling [12, 20], we specifically study these approaches. We show, under geometrical conditions, that the first and second order techniques proposed in [12, 20] define stochastic viscosity approximations of order one and two respectively, thus proving variance reduction properties for these heuristics. This also shows that our definition can indeed be applied to actual approximation schemes. A simple numerical example illustrates our results.

As for any new definition, connection between stochastic viscosity solutions and existing works should be explored further. In particular, concerning variance reduction, it seems interesting to understand better the link with the large deviation criteria designed in [21] as well as the more involved Isaacs equations [8]. Moreover, it seems natural to revisit existing approximation techniques and check whether they match our definition and, as already noted in [26], the relation with viscosity solutions of HJB problems should be made clearer, apart from the variance reduction problem.

Finally, although our work gives ground to the schemes based on instanton expansions as the ones presented in Section 4, such schemes do not apply as such when the solution to the instanton problem is not unique, which is the case in many situations. Two ways can be followed to overcome this issue. The first is to extend the approximation to a situation with several instantons, to build a more general abstract expansion of the HJB equation – this is the subject of ongoing investigation of the author. The second, more pragmatic, is to use tools such as the region of strong regularity [16] to update the computation of the instanton at certain times, in the flavour of [22]. Whatever approach is chosen, we believe the notion we introduced in this work is valuable as it quantitatively shows that, when the goal of control is variance reduction, one should focus on approximating the HJB equation along relevant low temperature paths.

Acknowledgements. The author particularly thanks the referee for providing many comments that helped to significantly improve the paper. The author is also grateful towards Charles Bertucci for reading a first version of the manuscript and providing a series of useful comments. He also thanks Tobias Grafke and Hugo Touchette for interesting discussions on large deviations, as well as Djilil Chafaï and Gabriel Stoltz for stimulating general discussions and encouragements.

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) 1185–1210.
- [2] G. Barles and E.R. Jakobsen, Error bounds for monotone approximation schemes for parabolic Hamilton–Jacobi–Bellman equations. *Math. Comput.* **76** (2007) 1861–1893.
- [3] G. Barles and E.R. Jakobsen, On the convergence rate of approximation schemes for Hamilton–Jacobi–Bellman equations. *ESAIM: Math. Model. Numer. Anal.* **36** (2002) 33–54.
- [4] A.N. Borodin and P. Salminen, Handbook of Brownian Motion–Facts and formulae. Springer Science and Business Media. Springer-Verlag, New York (2015).
- [5] J. Bucklew, Introduction to rare event simulation, in *Springer Series in Statistics*. Springer-Verlag, New York (2004).
- [6] M.G. Crandall and P.-L. Lions, Two approximations of solutions of Hamilton–Jacobi equations. *Math. Comput.* **43** (1984) 1–19.
- [7] P. Del Moral, Feynman–Kac Formulae. *Probability and its Applications*. Springer (2004).
- [8] P. Dupuis and H. Wang, Subsolutions of an Isaacs equation and efficient schemes for importance sampling. *Math. Oper. Res.* **32** (2007) 723–757.
- [9] P. Dupuis, A.D. Sezer and H. Wang, Dynamic importance sampling for queueing networks. *Ann. Appl. Prob.* **17** (2007) 1306–1346.
- [10] L.C. Evans, Partial differential equations, in *Graduate Studies in Mathematics*. Vol. 19. American Mathematical Society (2010).
- [11] 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) 349–380.
- [12] G. Ferré and T. Grafke, Approximate optimal controls via instanton expansion for low temperature free energy computation. *SIAM Multiscale Model. Simul.* **19** (2021) 1310–1332.
- [13] G. Ferré and G. Stoltz, Error estimates on ergodic properties of discretized Feynman–Kac semigroups. *Numer. Math.* **143** (2019) 261–313.
- [14] G. Ferré and G. Stoltz, Large deviations of empirical measures of diffusions in weighted topologies. *Electron. J. Prob.* **25** (2020) 1–52.

- [15] W.H. Fleming, Exit probabilities and optimal stochastic control. *Appl. Math. Optim.* **4** (1977) 329–346.
- [16] W.H. Fleming and M.R. James, Asymptotic series and exit time probabilities. *Ann. Probab.* **20** (1992) 1369–1384.
- [17] W.H. Fleming and H.M. Soner, Controlled Markov processes and viscosity solutions, in *Stochastic Modelling and Applied Probability*. Vol. 25. Springer Science & Business Media (2006).
- [18] M.I. Freidlin and A.D. Wentzell, Random perturbations of dynamical systems, in *Grundlehren der mathematischen Wissenschaften*. Vol. 260. Springer (1998).
- [19] P. Glasserman and Y. Wang, Counterexamples in importance sampling for large deviations probabilities. *Ann. Appl. Probab.* **7** (1997) 731–746.
- [20] T. Grafke and E. Vanden-Eijnden, Numerical computation of rare events via large deviation theory. *Chaos* **29** (2019) 063118.
- [21] A. Guyader and H. Touchette, Efficient large deviation estimation based on importance sampling. *J. Stat. Phys.* **181** (2020) 551–586.
- [22] M. Hairer and J. Weare, Improved diffusion Monte Carlo. *Comm. Pure Appl. Math.* **67** (2014) 1995–2021.
- [23] C. Hartmann and C. Schütte, Efficient rare event simulation by optimal nonequilibrium forcing. *J. Stat. Mech. Theory Exp.* **2012** (2012) 11004.
- [24] I. Karatzas and S. Shreve, Brownian motion and stochastic calculus, in *Graduate Texts in Mathematics*. vol. 113. Springer Science & Business Media (2012).
- [25] T. Lelièvre and G. Stoltz, Partial differential equations and stochastic methods in molecular dynamics. *Acta Numer.* **25** (2016) 681–880.
- [26] C. Léonard, Feynman-Kac formula under a finite entropy condition. Preprint [arXiv:2104.09171](https://arxiv.org/abs/2104.09171) (2021).
- [27] A.M. Oberman, Convergent difference schemes for degenerate elliptic and parabolic equations: Hamilton–Jacobi equations and free boundary problems. *SIAM J. Numer. Anal.* **44** (2006) 879–895.
- [28] H. Pham, Continuous-time stochastic control and optimization with financial applications, in *Stochastic Modelling and Applied Probability*. Vol. 61. Springer Science & Business Media (2009).
- [29] L. Rey-Bellet, Ergodic properties of Markov processes, in *Open Quantum Systems II*. Springer (2006) 1–39.
- [30] P.E. Souganidis, Approximation schemes for viscosity solutions of Hamilton–Jacobi equations. *J. Differ. Equ.* **59** (1985) 1–43.
- [31] E. Vanden-Eijnden and J. Weare, Rare event simulation of small noise diffusions. *Comm. Pure Appl. Math.* **65** (2012) 1770–1803.
- [32] 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) A2654–A2672.



Please help to maintain this journal in open access!

This journal is currently published in open access under the Subscribe to Open model (S2O). We are thankful to our subscribers and supporters for making it possible to publish this journal in open access in the current year, free of charge for authors and readers.

Check with your library that it subscribes to the journal, or consider making a personal donation to the S2O programme by contacting subscribers@edpsciences.org.

More information, including a list of supporters and financial transparency reports, is available at <https://edpsciences.org/en/subscribe-to-open-s2o>.