

IMPORTANCE SAMPLING FOR THE EMPIRICAL MEASURE OF WEAKLY INTERACTING DIFFUSIONS

Z.W. BEZEMEK AND M. HELDMAN

ABSTRACT. We construct an importance sampling method for computing statistics related to rare events for weakly interacting diffusions. Standard Monte Carlo methods behave exponentially poorly with the number of particles in the system for such problems. Our scheme is based on subsolutions of a Hamilton-Jacobi-Bellman (HJB) equation on Wasserstein space which arises in the theory of mean-field (McKean-Vlasov) control. We identify conditions under which such a scheme is asymptotically optimal. In the process, we make connections between the large deviations principle for the empirical measure of weakly interacting diffusions, mean-field control, and the HJB equation on Wasserstein space. We also provide evidence, both analytical and numerical, that with sufficient regularity of the HJB equation, our scheme can have vanishingly small relative error in the many particle limit.

1. INTRODUCTION

Consider the weakly interacting particle system:

$$(1) \quad dX_t^{i,N,s,y} = b(X_t^{i,N,s,y}, \mu_t^{N,s,y})dt + \sigma(X_t^{i,N,s,y}, \mu_t^{N,s,y})dW_t^i, \quad X_s^{i,N,s,y} = y_i$$

on some stochastic basis $(\Omega, \mathcal{F}, \mathbb{P}), \{\mathcal{F}_t\}_{t \in [s, T]}$ satisfying the usual conditions, where $b : \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}^d$, $\sigma : \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}^{d \times m}$, $\{W^i\}_{i \in \mathbb{N}}$ are independent m -dimensional standard $\mathcal{F}_t$ -Brownian motions initialized at $W_s^i = 0$, $i \in \{1, \dots, N\}$. Here $\mathcal{P}_2(\mathbb{R}^d)$ is the space of probability measures with finite second moment (see Definition A.1). In (1), N denotes the total number of particles, $i \in \{1, 2, \dots, N\}$ indexes a particular particle, and $\mu^{N,s,y} \in C([0, T]; \mathcal{P}(\mathbb{R}^d))$ is the empirical measure:

$$(2) \quad \mu_t^{N,s,y} := \frac{1}{N} \sum_{i=1}^N \delta_{X_t^{i,N,s,y}}, \quad t \in [s, T].$$

The i 'th entry of the vector $y = (y_1, y_2, \dots) \in \oplus_{i=1}^\infty \mathbb{R}^d$ contains the deterministic initial condition for the i 'th particle at time s satisfying $0 \leq s \leq t \leq T$. The entries y_i for $i > N$ are ignored for any particular value of N ; we include them in defining (1) for convenience, since our purpose will be to consider the sequence of solutions to (1) as $N \rightarrow \infty$. We will at times identify y with its projection onto $\mathbb{R}^{dN} = [\mathbb{R}^d]^N = \oplus_{i=1}^N \mathbb{R}^d$ without ambiguity.

In this paper, we design a control-based importance sampling scheme for estimating functionals of the form

$$(3) \quad \mathbb{E} \left[\exp \left(-NG(\mu_T^{N,s,y}) \right) \right], \quad G : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathbb{R}.$$

The basic method goes as follows: letting $v_i^N : [s, T] \times \mathbb{R}^{dN} \rightarrow \mathbb{R}^m$ be a chosen bounded function, or control, for each $i \in \{1, \dots, N\}$ and $N \in \mathbb{N}$, we consider the *controlled* version of the interacting particle system (1)

Date: November 17, 2023.

2010 Mathematics Subject Classification. 60F10, 60F05, 65C05.

Key words and phrases. interacting particle systems, empirical measure, large deviations, importance sampling.

Z.W.B. was partially supported by the National Science Foundation (DMS 2107856). The majority of this work was completed while Z.W.B. was a PhD candidate at Boston University. M.H. was partially supported by NSF-DMS 1902854, ARO W911NF-20-1-0244, and a subgrant of NSF-OAC 2139536. The authors of the paper would like to thank both reviewers for a their careful and constructive reviews of this article.

given by

$$(4) \quad \begin{aligned} d\hat{X}_t^{i,N,s,y} &= [b(\hat{X}_t^{i,N,s,y}, \hat{\mu}_t^{i,N,s,y}) + \sigma(\hat{X}_t^{i,N,s,y}, \hat{\mu}_t^{i,N,s,y})v_i^N(t, \hat{X}_t^{1,N,s,y}, \dots, \hat{X}_t^{N,N,s,y})]dt \\ &\quad + \sigma(\hat{X}_t^{i,N,s,y}, \hat{\mu}_t^{i,N,s,y})d\hat{W}_t^i, \quad \hat{X}_s^{i,N,s,y} = y_i. \end{aligned}$$

Girsanov's theorem allows us to convert statistics for the controlled system (4) into statistics for (1); therefore, in order to develop a more accurate Monte Carlo estimator for (1) our objective will be to identify controls v_i^N that reduce the variance of our target statistic compared with the uncontrolled version, or in the best case cause the relative variance to vanish as $N \rightarrow \infty$. By these measures, we find that subsolutions of a zero-viscosity Hamilton-Jacobi-Bellman (HJB) equation on Wasserstein space, which we derive starting with the Dawson-Gärtner large deviations principle (LDP) [16, 29], provide controls with good performance.

Our main results, stated rigorously in Section 2, show that under certain assumptions our scheme outperforms standard Monte Carlo methods by requiring a subexponential number of samples to achieve a given relative error in estimating (3) as $N \rightarrow \infty$ (Theorem 2.4). Under stronger assumptions, we show that in fact only a *vanishing* number of simulations suffices (Theorem 2.5), so that for large N our method requires only a single sample. Our theoretical results are confirmed by numerical experiments in Section 4.

Our importance sampling scheme is related to and inspired by previous work concerning the development of asymptotically optimal importance sampling methods for estimating expectations of the form

$$(5) \quad \mathbb{E} \left[\exp \left(-\frac{1}{\epsilon} G(X_T^{\epsilon,s,y}) \right) \right], \quad G \in C_b(\mathbb{R}^d),$$

where $X^{\epsilon,s,y}$ is a single particle satisfying, e.g., the small-noise SDE:

$$(6) \quad \begin{aligned} dX_t^{\epsilon,s,y} &= b(X_t^{\epsilon,s,y})dt + \sqrt{\epsilon}\sigma(X_t^{\epsilon,s,y})dW_t, \\ X_s^{\epsilon,s,y} &= y \in \mathbb{R}^d. \end{aligned}$$

We will draw parallels between the small-noise setting and ours throughout this article.

A general strategy for efficiently computing (5) involves exploiting the connection between the large deviations rate function of Freidlin-Wentzell (FW) [48] and a class of first-order Hamilton-Jacobi-Bellman (HJB) equations, whose subsolutions can then be leveraged to obtain optimal controls for the zero-noise process $X^{0,s,y}$. In turn, the controls can be used to design an optimal change of measure for estimating (5). The control formulation of the FW rate function, provided via the weak convergence approach to large deviations of Dupuis and Ellis [13, 34], makes the chain linking the rate function, zero-viscosity HJB equation, and optimal control of the zero-noise process more evident. The strategy described above has been applied for a diverse range of stochastic models [36–38, 74, 79].

Elevating the above framework for small-noise SDE importance sampling schemes to interacting particle systems is a major contribution of this paper. Moreover, the connections between the approaches used herein and in the small-noise setting allow us to conjecture (see Remark 2.6) about extending this research to problems related to metastable dynamics that have been previously considered for the small-noise case (see, e.g., [57] and the references therein) such as exit probabilities and mean first passage times. We therefore view this paper as the first step towards a new method of designing importance sampling schemes for studying properties of the empirical measure which are of interest to the greater scientific community, such as the free-energy differences for chemical and biological molecular systems [3, 25, 61, 78] or exit times of the empirical measure from domains of attraction in models related to social dynamics and consensus convergence [6, 28, 49, 50, 59].

Let us now make the comparison between the interacting particle system (1) and the small-noise system (6) more explicit. It is well known that $\mu^{N,s,y} \rightarrow \mathcal{L}(X^{s,\nu})$ in distribution when considered as $\mathcal{P}(C([s, T]; \mathbb{R}^d))$ -valued random variables. Here $\nu \in \mathcal{P}(\mathbb{R}^d)$ is the weak limit as $N \rightarrow \infty$ of the empirical measures $\frac{1}{N} \sum_{i=1}^N \delta_{y_i}$, and $X^{s,\nu}$ satisfies the McKean-Vlasov equation:

$$(7) \quad dX_t^{s,\nu} = b(X_t^{s,\nu}, \mathcal{L}(X_t^{s,\nu}))dt + \sigma(X_t^{s,\nu}, \mathcal{L}(X_t^{s,\nu}))dW_t, \quad X_s^{s,\nu} \sim \nu.$$

This phenomenon, known as the propagation of chaos, goes back to the works of Kac [63]. The empirical measure $\mu^{N,s,y}$ can also formally be seen to satisfy an *infinite-dimensional* small-noise equation, analogous

to (6), with parameter $\epsilon = \frac{1}{\sqrt{N}}$, as per pp. 249-250 of [29]. Thus, the problem of estimating (3) has the same basic structure as the more well-known problem of estimating (5), but with a different state space.

An LDP for the many-particle limit of $\{\mu^{N,s,y}\}_{N \in \mathbb{N}}$ was first established in the classical work of [29] in the case where the diffusion coefficient σ does not depend on the empirical measure, with the result extended via the weak convergence approach of Dupuis and Ellis [34] to the general setting in [16]. In the latter, the form of the rate function is given in terms of a control problem (see Theorem 1.2), which further motivates our search for an asymptotically optimal importance sampling scheme for (3) associated with an HJB equation as in the small-noise SDE setting. In contrast to the setting of small-noise SDEs, where the state space of the HJB equation is Euclidean space, we find that the appropriate HJB equation to consider in our setting of measure-valued random variables is posed on Wasserstein space.

Perhaps due to the computational intractability of both the Dawson-Gärtner [29] and Budhiraja-Dupuis-Fischer [16] rate functions, practical applications of the LDP for the many-particle limit of the empirical measure (2) associated with (1) have been few and far between (for some exceptions see, e.g., [30, 49, 50, 70]). On the other hand, in recent years the HJB equation on Wasserstein space of [73] (along with the related equations in [21, 67]) has received an immense amount of attention both in terms of numerical applications [20, 22, 47, 54, 55, 60, 66] and theoretical results [5, 18, 19, 26, 27, 56, 65, 73, 80]. The Dawson-Gärtner rate function has previously been related to the theory of mean field games and control through the observation that it can be viewed in terms of derivatives of the free energy associated to the limiting McKean-Vlasov equation viewed as a gradient flow on Wasserstein space in some settings [1, 2, 4, 41–44, 52]. However, to our knowledge, this paper is the first time that the connection between the LDP rate function (10) and the HJB equation on Wasserstein space (23) has been made explicit in the literature (see Remark 2.2).

As far as we are aware, this paper also represents the first time the problem of estimating (3) with an asymptotically optimal importance sampling scheme has been considered, though many other versions of the small-noise importance sampling problem have appeared in the literature. These include, for example, discrete-time Markov chains [37, 38], diffusions with multiscale structure [36, 69], and stochastic partial differential equations [40, 51]. See in addition [11, 35, 39, 53, 57, 79] for a diverse, but by no means comprehensive, collection of small-noise importance sampling research in a variety of settings. We in particular mention the work of [33] on the design of importance sampling schemes for small-noise McKean-Vlasov SDEs using a decoupling approach (see also [7, 8]) and a complete change of measure approach. The problem posed in [33] is significantly different from the one we consider here, principally because they seek to estimate statistics of the limiting McKean-Vlasov equation (7) rather than the interacting particle system (1). The change of measure they employ is therefore based on the Freidlin-Wentzell (small-noise) LDP for Brownian motion, rather than the Sanov-type (many-particle) LDP for the empirical measure (2) used in our setting. Note that Freidlin-Wentzell LDPs have been derived in different settings for McKean-Vlasov equations [32, 68], though they were not directly employed in the results of [7, 8, 33].

Finally, we note that methods of importance sampling for high-dimensional diffusions have also been established in previous research, e.g., [11, 71, 77]). While in theory these could be applied to estimate (3), they are designed for more general problems than ours and therefore do not exploit the exchangeability of the particles, and resulting propagation of chaos, to obtain an asymptotically optimal scheme. However, it may be that a combination of these non-asymptotic techniques with ours could allow us to treat a broader range of problems in the future, as we remark in Section 6.

The rest of the paper is organized as follows. We establish notation in Subsection 1.1 below, and go on to provide background on large deviations theory for the empirical measure in Subsection 1.2. Then, in Subsection 1.3 we develop measures of performance for the control-based importance sampling of diffusions for computing statistics of the form (3), some of which can be quantified using large deviations theory. In Section 2, we formally derive the HJB Equation on Wasserstein space and state our main results, Theorem 2.4 and Theorem 2.5. In Section 3, we introduce a class of linear-quadratic examples for which the HJB equation can be solved analytically and discuss simple cases in which our scheme can be thoroughly analyzed in comparison to standard Monte Carlo, along the way comparing with results for small-noise diffusion processes on $\mathbb{R}^d$. In Section 4, we demonstrate our method with numerical examples. Section 5 contains detailed proofs of the results stated in Section 2, before Section 6 provides concluding remarks. Lastly, in

Appendix A we provide additional background on the calculus of square-integrable probability measures, $\mathcal{P}_2(\mathbb{R}^d)$.

1.1. Notation. In the following paragraph, let $\mathbf{X}$ and $\mathbf{Y}$ be Polish spaces, and $(\Omega, \mathcal{F}, \mu)$ a measure space. In the course of this paper, we denote by $\mathcal{P}(\mathbf{X})$ the space of probability measures on $\mathbf{X}$ endowed with the topology of weak convergence. We use $\mathcal{P}_2(\mathbf{X}) \subset \mathcal{P}(\mathbf{X})$ to denote the subspace of $\mathcal{P}(\mathbf{X})$ consisting of square-integrable probability measures on $\mathbf{X}$ endowed with the 2-Wasserstein metric (see Definition A.1). We additionally let $\mathcal{B}(\mathbf{X})$ be the Borel σ -algebra associated with $\mathbf{X}$, $C(\mathbf{X}; \mathbf{Y})$ the space of continuous functions from $\mathbf{X}$ to $\mathbf{Y}$, $C_b(\mathbf{X})$ the space of bounded, continuous functions from $\mathbf{X}$ to $\mathbb{R}$ with norm $\|\phi\|_\infty := \sup_{x \in \mathbf{X}} |\phi(x)|$, and for $p \geq 1$ we define $L^p(\Omega, \mathcal{F}, \mu; \mathbb{R}^d)$ as the space of p -integrable functions on $(\Omega, \mathcal{F}, \mu)$ with values in $\mathbb{R}^d$ and norm

$$\|\phi\|_{L^p(\Omega, \mathcal{F}, \mu; \mathbb{R}^d)} = \left(\int_\Omega |\phi(z)|^p \mu(dz) \right)^{1/p}.$$

$C_{b,L}(\mathbb{R}^d) \subset C_b(\mathbb{R}^d)$ is the set of bounded, globally Lipschitz functions $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ with norm $\|\phi\|_{b,L} := \|\phi\|_\infty + \sup_{x,y \in \mathbb{R}^d, x \neq y} \frac{|\phi(x) - \phi(y)|}{|x - y|}$. For $\phi \in L^1(\mathbf{X}, \mu)$, $\mu \in \mathcal{P}(\mathbf{X})$, we define the pairing $\langle \phi, \mu \rangle$ by

$$\langle \phi, \mu \rangle := \int_{\mathbf{X}} \phi(x) \mu(dx).$$

We note that, as in the above definitions, the codomain for a given function space is assumed to be $\mathbb{R}$ unless otherwise specified. We will also occasionally refer to a measure on a topological space $\mathbf{X}$ without explicitly stating the corresponding σ -algebra; in this case, the measurable sets are assumed to come from $\mathcal{B}(\mathbf{X})$.

Partial derivatives of a function ϕ with respect to a variable x are normally denoted by $\partial_x \phi$; if the variable x is measure valued then the derivative ∂_x should be interpreted in the Lions sense (see Appendix A). For time derivatives $\partial_t \phi$ with respect to $t \in I \subseteq \mathbb{R}^+$, we also use the notation $\dot{\phi}$.

For $\phi : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$, we use $\phi \in C^{1,2}([0, T] \times \mathcal{P}_2(\mathbb{R}^d))$ to mean that $\mu \mapsto \phi(t, \mu)$ is fully C^2 in the sense of Definition A.2 for all t , with all derivatives in that definition jointly continuous in (t, z, z', μ) and the map $t \mapsto \phi(t, \mu)$ continuously differentiable in t . The subspace $C_b^{1,2}([0, T] \times \mathcal{P}_2(\mathbb{R}^d)) \subset C^{1,2}([0, T] \times \mathcal{P}_2(\mathbb{R}^d))$ denotes the class of functions ϕ such that further ϕ and all its derivatives from Definition A.2 are uniformly bounded.

For $x, y \in \mathbb{R}^d$, $x = (x_1, \dots, x_d)$, $y = (y_1, \dots, y_d)$, $x \cdot y$ denotes the standard inner product $x \cdot y := \sum_{i=1}^d x_i y_i$. For $A, B \in \mathbb{R}^{d \times d}$ matrices with entries $\{a_{ij}\}_{i,j=1}^d$ and $\{b_{ij}\}_{i,j=1}^d$ respectively, $A : B := \sum_{i,j=1}^d a_{ij} b_{ij}$.

For $x \in \oplus_{i=1}^N \mathbb{R}^d$ or $\oplus_{i=1}^\infty \mathbb{R}^d$, we will denote by μ_x^N associated empirical measure, i.e., the element of $\mathcal{P}(\mathbb{R}^d)$ given by

$$\mu_x^N := \frac{1}{N} \sum_{i=1}^N \delta_{x_i}.$$

Finally, for given sequences $\{f_n\}, \{g_n\}$ in $\mathbb{R}^+$, we use the following notation to compare their asymptotic behavior: $f_n = \mathcal{O}(g_n)$ if there is a constant $C > 0$ independent of n such that $f_n \leq C g_n$, and $f_n = o(g_n)$ if $\frac{f_n}{g_n} \rightarrow 0$ as $n \rightarrow \infty$. If instead of positive real numbers f_n and g_n are sequences on normed spaces, we use the same notation to refer to the behavior of the sequences $|f_n|$ and $|g_n|$.

1.2. Large deviations results and an HJB equation for the particle system. We now introduce some of the key tools for proving our main results in Section 2, namely, a stochastic control representation for the LDP associated with the empirical measure (2) as provided by [16] and a control representation for statistics of the form (3). The final tool, the HJB equation on Wasserstein space, will be introduced in Section 2.

The Laplace principle stated in Theorem 1.2 is based on the controlled version of (7):

$$(8) \quad \begin{aligned} d\hat{X}_t^{v,s,\nu} &= [b(\hat{X}_t^{v,s,\nu}, \mathcal{L}(\hat{X}_t^{v,s,\nu})) + \sigma(\hat{X}_t^{v,s,\nu}, \mathcal{L}(\hat{X}_t^{v,s,\nu}))v(t)]dt + \sigma(\hat{X}_t^{v,s,\nu}, \mathcal{L}(\hat{X}_t^{v,s,\nu}))dW_t, \\ \hat{X}_s^{v,s,\nu} &\sim \nu. \end{aligned}$$

For the LDP to hold, we need the following assumptions on the particle system (1), the controlled version of the limiting McKean-Vlasov equation (8), and the problem data:

Assumption 1.1. *Suppose:*

- (A1) *The initial conditions y for (1) are such that there exists $\nu \in \mathcal{P}_2(\mathbb{R}^d)$ with $\mu_y^N \rightarrow \nu$ as $N \rightarrow \infty$.*
- (A2) *The coefficients b and σ are continuous on $\mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$.*
- (A3) *For all $N \in \mathbb{N}$, existence and uniqueness of solutions holds in the strong sense for the system of SDEs (1).*
- (A4) *Weak uniqueness holds for solutions of (8) in the sense that if Θ^1, Θ^2 are probability measures such that for $i = 1, 2$:*
 - *Θ^i is the law of some process triple $(\hat{X}^{v_i, s, \nu}, v_i, W^i)$ satisfying (8) on some filtered probability space $(\Omega^i, \mathcal{F}^i, \mathbb{P}^i), \{\mathcal{F}_t^i\}_{t \in [s, T]}$ satisfying the usual conditions*
 - *v_i is an $\mathcal{F}_t^i$ -progressively measurable control with $\mathbb{E}^{\Theta^i} \left[\int_s^T |v_i(t)|^2 dt \right] < \infty$*
 - *W^i a standard $\mathcal{F}_t^i$ -Brownian motion with $W_s^i = 0$**and $\Theta^1 \circ \vartheta^{-1} = \Theta^2 \circ \vartheta^{-1}$ where ϑ is the identity mapping in the second and third coordinates and the evaluation map at time s in the first coordinate, then $\Theta^1 = \Theta^2$.*
- (A5) *For any sequence of controlled processes $\{\hat{X}^{i, N, s, y}\}$ from (4) with $\mathcal{F}_t$ -progressively measurable controls $v^N = (v_1^N, \dots, v_N^N)$ satisfying*

$$\sup_{N \in \mathbb{N}} \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \int_s^T |v_i^N(t)|^2 dt \right] < \infty,$$

$\hat{\mu}^N := \frac{1}{N} \sum_{i=1}^N \delta_{\hat{X}^{i, N, s, y}}$ *is tight as a sequence of $\mathcal{P}(C([s, T]; \mathbb{R}^d))$ -valued random variables.*

Assumption 1.1 is taken directly from [16] with the exception of (A4), which is a mildly less restrictive form of weak uniqueness under which the arguments from [16] still hold (see Lemma 3.4 in [15] or Definition 3.6 in [9]). For verifiable conditions on the data of the problem under which (A4) holds, from Appendix C in [45] we have that:

- (A1') b, σ are jointly globally Lipschitz on $\mathbb{R}^d \times \mathcal{P}(\mathbb{R}^d)$ where $\mathcal{P}(\mathbb{R}^d)$ is equipped with the bounded Lipschitz metric $d_{BL}(\mu, \nu) := \sup_{\{\phi \in C_{b,L}(\mathbb{R}^d) : \|\phi\|_{b,L} \leq 1\}} \int_{\mathbb{R}^d} \phi(x) [\mu(dx) - \nu(dx)]$.
- (A2') σ is bounded and $\sigma(x, \mu) = \sigma(\mu)$.

are enough to imply (A2)-(A5). Note that the proof given in [45] that (A1') and (A2') together imply (A4) (Proposition C.1 in [45]), is based on an erroneous localization argument. See [15] for a correct proof.

We are now ready to state the LDP for the empirical measure.

Theorem 1.2. *Under Assumption 1.1, for any $F \in C_b(C([s, T]; \mathcal{P}(\mathbb{R}^d)))$:*

$$(9) \quad \lim_{N \rightarrow \infty} -\frac{1}{N} \log \mathbb{E} \left[\exp(-NF(\mu^{N, s, y})) \right] = \mathcal{G}(s, \nu; F) := \inf_{\mu \in C([s, T]; \mathcal{P}(\mathbb{R}^d))} \left\{ S_{s, T}^\nu(\mu) + F(\mu) \right\},$$

where $S_{s, T}^\nu : C([s, T]; \mathcal{P}(\mathbb{R}^d)) \rightarrow [0, +\infty]$ is given by

$$(10) \quad S_{s, T}^\nu(\mu) := \inf_{v \in \mathcal{U} : \mu(t) = \mathcal{L}(\hat{X}_t^{v, s, \nu}), t \in [s, T]} \frac{1}{2} \mathbb{E} \left[\int_s^T |v(t)|^2 dt \right].$$

In the above we take $\inf \emptyset = +\infty$, and in the expression for $S_{s, T}^\nu(\mu)$ the admissible set of controls $\mathcal{U}$ is the set of all quadruples $((\Omega, \mathcal{F}, \mathbb{P}), \{\mathcal{F}_t\}_{t \in [s, T]}, v, W)$ such that:

1. $((\Omega, \mathcal{F}, \mathbb{P}), \{\mathcal{F}_t\})$ forms a stochastic basis satisfying the usual conditions,
2. W is a standard m -dimensional Brownian motion initialized at $W_s = 0$,
3. v is an $\mathbb{R}^m$ -valued $\{\mathcal{F}_t\}$ -progressively measurable process such that $\mathbb{E} \left[\int_s^T |v(t)|^2 dt \right] < \infty$.

We write $v \in \mathcal{U}$ to denote $((\Omega, \mathcal{F}, \mathbb{P}), \{\mathcal{F}_t\}_{t \in [s, T]}, v, W) \in \mathcal{U}$ in the above. For each $v \in \mathcal{U}$, $\hat{X}_t^{v, \nu}$ satisfies the controlled McKean-Vlasov equation (8) on the space $(\Omega, \mathcal{F}, \mathbb{P}), \{\mathcal{F}_t\}_{t \in [s, T]}$ with driving Brownian motion W .

Proof. The result follows from Theorem 3.1 in [16] (see also Theorem 5.1 in [45]) by applying the contraction principle to pose the rate function on $C([s, T]; \mathcal{P}(\mathbb{R}^d))$ rather than $\mathcal{P}(C([s, T]; \mathbb{R}^d))$ as in Proposition 5.6 of [9]. $\square$

The following corollary of Theorem 3.6 in [13], which provides a prelimit expression for the left-hand side of (9) which can be applied to prove Theorem 1.2, is also critical to our arguments in this paper. In Section 5, we apply it to several auxiliary systems of the form (1) with different choices of the drift term in the course of proving Theorem 2.4. The drift terms in the new system have a slightly different form from those in (1), since they are allowed here to directly depend on the time variable.

To distinguish these auxiliary systems and their solutions from those of the original system (1), we establish alternative notation here for the solutions of the controlled and uncontrolled auxiliary systems and their empirical measures that we will also use in Section 5.

Proposition 1.3. *Let $\tilde{\mu}_t^{N,s,y} = \frac{1}{N} \sum_{i=1}^N \delta_{\tilde{X}_t^{i,N,s,y}}, t \in [s, T]$ where $\tilde{X}_t^{i,N,s,y}$ solves the system (11) (i.e., (1) with b replaced by $\tilde{b}$):*

$$(11) \quad d\tilde{X}_t^{i,N,s,y} = \tilde{b}(t, \tilde{X}_t^{i,N,s,y}, \tilde{\mu}_t^{N,s,y})dt + \sigma(\tilde{X}_t^{i,N,s,y}, \tilde{\mu}_t^{N,s,y})d\tilde{W}_t^i, \quad \tilde{X}_s^{i,N,s,y} = y_i$$

and $\tilde{b}, \sigma$ are such that existence and uniqueness holds in the strong sense for the system of SDEs (11) and $\tilde{W}^i$ are IID standard m -dimensional Brownian motions initialized at $\tilde{W}_s^i = 0$. Then for $F \in C_b(C([s, T]; \mathcal{P}(\mathbb{R}^d)))$:

$$(12) \quad -\frac{1}{N} \log \mathbb{E} \left[\exp(-NF(\tilde{\mu}^{N,s,y})) \right] = \inf_{u^N \in \mathcal{U}^N} \left\{ \mathbb{E} \left[\frac{1}{2N} \sum_{i=1}^N \int_s^T |u_i^N(t)|^2 dt \right] + \mathbb{E}[F(\tilde{\mu}^{N,s,y})] \right\},$$

where $\mathcal{U}^N$ is the space of adapted controls $u^N = (u_1^N, \dots, u_N^N), u_i^N : [s, T] \rightarrow \mathbb{R}^m$ such that $\mathbb{E} \left[\int_s^T |u_i^N(t)|^2 dt \right] < \infty$, and $\tilde{\mu}^{N,s,y}$ is the empirical measure associated with the controlled version of (11), i.e.,

$$d\bar{X}_t^{i,N,s,y} = \tilde{b}(t, \bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y})dt + \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y})u(t)dt + \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y})d\tilde{W}_t^i, \quad \bar{X}_s^{i,N,s,y} = y_i.$$

1.3. Log-efficiency and asymptotic optimality in control-based importance sampling. Recall our definition of the controlled particle system (4), and let

$$(13) \quad Z^{i,N,s,y} := \exp \left(- \int_s^T v_i^N(t, \hat{X}_t^{i,N,s,y}, \dots, \hat{X}_t^{i,N,s,y}) \cdot d\hat{W}_t^i - \frac{1}{2} \int_s^T |v_i^N(t, \hat{X}_t^{i,N,s,y}, \dots, \hat{X}_t^{i,N,s,y})|^2 dt \right).$$

By Girsanov's theorem, an unbiased estimator for (3) is given by

$$(14) \quad \hat{\delta}_N := \frac{1}{M} \sum_{j=1}^M \exp \left(-NG(\hat{\mu}^{N,s,y,j}) \right) \prod_{i=1}^N Z^{i,N,s,y,j}$$

where $(\hat{\mu}^{N,s,y,j}, Z^{1,N,s,y,j}, \dots, Z^{N,N,s,y,j})$ are M independent samples of $(\hat{\mu}^{N,s,y}, Z^{1,N,s,y}, \dots, Z^{N,N,s,y})$, and

$$(15) \quad \hat{\mu}_t^{i,N,s,y} := \frac{1}{N} \sum_{i=1}^N \delta_{\hat{X}_t^{i,N,s,y}}.$$

We note that the standard Monte Carlo method estimator, which we denote by δ_N , corresponds to (14) with $v_i^N = 0$, so that $Z^{i,N,s,y,j} \equiv 1$.

We define the *relative error* of the estimator (14) by

$$(16) \quad \rho(\hat{\delta}_N) = \frac{1}{\sqrt{M}} \sqrt{\frac{\mathbb{E} \hat{\delta}_N^2 - [\mathbb{E} \hat{\delta}_N]^2}{[\mathbb{E} \hat{\delta}_N]^2}} = \frac{1}{\sqrt{M}} \sqrt{\frac{\mathbb{E} \left[\exp(-2NG(\hat{\mu}^{N,s,y})) \prod_{i=1}^N (Z^{i,N,s,y})^2 \right]}{\mathbb{E} \left[\exp(-NG(\mu^{N,s,y})) \right]^2}} - 1.$$

To control the size of (16) we either need to increase the number of samples M or reduce the variance by choosing v_i^N so that

$$(17) \quad R(\hat{\delta}_N) := \frac{\mathbb{E} \left[\exp(-2NG(\hat{\mu}^{N,s,y})) \prod_{i=1}^N (Z^{i,N,s,y})^2 \right]}{\mathbb{E} \left[\exp(-NG(\mu^{N,s,y})) \right]^2},$$

is close to 1. With this in mind, we define our first measure of the efficiency for an importance sampling scheme (14) in terms of the asymptotic behavior of $R(\hat{\delta}_N)$ as $N \rightarrow \infty$:

Definition 1.4. *An importance sampling scheme of the form (14) is called **log-efficient** if*

$$\lim_{N \rightarrow \infty} -\frac{1}{N} \log R(\hat{\delta}_N) = 0.$$

Note that, in our definition, we do not require that $R(\hat{\delta}_N) \rightarrow 1$ or even $R(\hat{\delta}_N) = \mathcal{O}(1)$; instead we only ask that as $N \rightarrow \infty$, $R(\hat{\delta}_N)$ does not grow exponentially in N .

The notion of log-efficiency can also be described in terms of large deviations theory. For the uncontrolled system, from (9) in Theorem 1.2 we know that if $G \in C_b(\mathcal{P}(\mathbb{R}^d))$ and Assumption 1.1 holds, then

$$(18) \quad \begin{aligned} \lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{E}[\exp(-NG(\mu_T^{N,s,y}))] &= -\mathcal{G}(s, \nu; \mu \mapsto G(\mu_T)) := -\gamma_1, \\ \lim_{N \rightarrow \infty} \frac{1}{N} \log \mathbb{E}[\exp(-2NG(\mu_T^{N,s,y}))] &= -\mathcal{G}(s, \nu; \mu \mapsto 2G(\mu_T)) := -\gamma_2. \end{aligned}$$

Thus, for the standard Monte Carlo scheme the relative error can be written in terms of γ_1 and γ_2 as follows:

$$(19) \quad \rho(\delta_N) = \frac{1}{\sqrt{M}} \sqrt{\exp(N[2\gamma_1 - \gamma_2 + o(1)]) - 1}$$

Starting from the inequality $\mathbb{E}[\exp(-2NG(\mu_T^{N,s,y}))] \geq \mathbb{E}[\exp(-NG(\mu_T^N))]^2$, taking logarithms, multiplying by $-\frac{1}{N}$, and taking the limit $N \rightarrow \infty$ shows that $\gamma_2 \leq 2\gamma_1$. Evidently, the relative error $\rho(\delta_N)$ grows exponentially in N if the inequality is strict, i.e., $\gamma := 2\gamma_1 - \gamma_2 > 0$. Exponentially many samples are therefore required to reduce the relative error below a given tolerance in that case.

Now, we assume that the relative error for the importance sampling scheme (16) admits the same representation as the standard Monte Carlo importance sampling scheme in (19), with analogous quantities $\hat{\gamma}_1$, $\hat{\gamma}_2$, and $\hat{\gamma} = 2\hat{\gamma}_1 - \hat{\gamma}_2$. That is:

$$(20) \quad \rho(\hat{\delta}_N) = \frac{1}{\sqrt{M}} \sqrt{\exp(N[\hat{\gamma} + o(1)]) - 1}.$$

Note that in the same way as with standard Monte Carlo, we can conclude via Jensen's inequality that $\hat{\gamma} \geq 0$. Then we can compare the relative error $\rho(\hat{\delta}_N)$ to the Monte Carlo relative error $\rho(\delta_N)$ by comparing the corresponding quantities $\hat{\gamma}$ and γ :

1. If $\hat{\gamma} < \gamma$, then the importance sampling scheme is more efficient than standard Monte Carlo in the sense that asymptotically (as $N \rightarrow \infty$), it yields a smaller relative error.
2. If $\hat{\gamma} = 0$, then the importance sampling scheme is log-efficient and, if $\gamma > 0$,

$$\rho(\hat{\delta}_N) = \frac{1}{\sqrt{M}} \sqrt{\exp(o(N)) - 1} = o(\rho(\delta_N)).$$

The object of our first result Theorem 2.4 presented in the next section is to show conditions under which, for our choice of controls, $\hat{\gamma} < \gamma$ or $\hat{\gamma} = 0$.

Though the assumptions underlying Theorem 2.4 preclude the relative error of our importance sampling scheme from growing exponentially, extremely fast subexponential growth rates may still be possible. To obtain a more refined result, one needs to study the behavior of the $o(1)$ terms in $R(\hat{\delta}_N)$. To that end, our second main result, Theorem 2.5, assumes an expansion of the prelimit quantities on the left-hand side of (18) leading to an expression for $R(\hat{\delta}_N)$ of the form

$$(21) \quad R(\hat{\delta}_N) = \exp \left(a_0 N + a_1 + a_2 N^{-1} + \dots \right).$$

Note that here a_0 plays the same role as $\hat{\gamma}$ in (20). Analysis of the coefficients in the power series appearing in the exponential shows that, under stronger regularity assumptions than required for our Theorem 2.4, $a_0 = a_1 = 0$, which implies that the importance sampling scheme has *vanishing* relative error:

Definition 1.5. *An estimator of the form (14) is said to have **vanishing relative error** if, for $\rho(\hat{\delta}_N)$ as in (16),*

$$\lim_{N \rightarrow \infty} \rho(\hat{\delta}_N) = 0.$$

We conclude this chapter by comparing to the small-noise diffusion regime. Our discussion in this section parallels that of the first section in [79], in which similar importance sampling scheme for small-noise diffusions is presented along with numerical schemes for estimating the optimal controls. However, we note that in our paper stating the LDP and its underlying assumptions is somewhat more involved due to the measure-valued state space.

In [79], the authors also provide sufficient conditions for a log-efficient importance sampling scheme for small-noise diffusions to have vanishing relative error. For our Theorem 2.5, we instead opt to perform an analysis in the spirit of the later [75]. The paper [75] furthered the analysis of the small-noise case by expanding the $o(1)$ terms in the relative error into powers of the small-noise parameter ϵ to recover log-efficiency from the first-order expansion, and the vanishing relative error result from of [79] from the second-order expansion (see also the related Section 3 of [62]).

2. THE HJB EQUATION AND STATEMENT OF MAIN RESULTS

We now turn to the full statement of our importance sampling scheme in terms of the solution to an HJB equation on Wasserstein space, explain its relationship to the rate function (10), and give a formal derivation of the importance sampling scheme and the HJB equation as the zero viscosity limit of a sequence of zero variance controls.

Starting with the Hamiltonian function which acts on $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, $p : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\Gamma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ by:

$$(22) \quad H_0(\mu, p, \Gamma) = - \int_{\mathbb{R}^d} b(x, \mu) \cdot p(x) - \frac{1}{2} p(x) \cdot [\sigma \sigma^\top(x, \mu) p(x)] + \frac{1}{2} \sigma \sigma^\top(x, \mu) : \Gamma(x) \mu(dx),$$

the HJB equation on Wasserstein space is given by:

$$(23) \quad \begin{aligned} -\partial_t \Psi(t, \nu) + H_0(\nu, \partial_\mu \Psi(t, \nu)[\cdot], \partial_z \partial_\mu \Psi(t, \nu)[\cdot]) &= 0, \quad t \in [0, T], \nu \in \mathcal{P}_2(\mathbb{R}^d), \\ \Psi(T, \nu) &= G(\nu), \quad \nu \in \mathcal{P}_2(\mathbb{R}^d). \end{aligned}$$

Recall here that $\partial_\mu \Psi(t, \nu)[\cdot] : \mathbb{R}^d \rightarrow \mathbb{R}^d$, the derivative of Ψ in the measure-valued variable μ , denotes the Lions derivative of Ψ at ν (Definition A.1 in Appendix A).

We will make use of the following notion of subsolutions to (23):

Definition 2.1. *We call $\Psi : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ a **classical subsolution** to (23) if:*

- (i) $\Psi \in C^{1,2}([0, T] \times \mathcal{P}_2(\mathbb{R}^d))$
- (ii) $\partial_t \Psi(t, \mu) - H_0(\mu, \partial_\mu \Psi(t, \mu)[\cdot], \partial_z \partial_\mu \Psi(t, \mu)[\cdot]) \geq 0$ for all $t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d)$
- (iii) $\Psi(T, \mu) \leq G(\mu), \mu \in \mathcal{P}_2(\mathbb{R}^d)$
- (iv) *There exists $C > 0$ such that*

$$\begin{aligned} \sup_{t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d), z \in \mathbb{R}^d} |\partial_\mu \Psi(t, \mu)[z]| &\leq C, \\ \sup_{t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d)} |\partial_z \partial_\mu \Psi(t, \mu)[z]| &\leq C(1 + |z|^2), \\ \sup_{t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d)} \left\| \partial_\mu^2 \Psi(t, \mu)[\cdot, \cdot] \right\|_{L^2(\mathbb{R}^d; \mu) \otimes L^2(\mathbb{R}^d; \mu)} &\leq C. \end{aligned}$$

Note that under the above definition, a classical subsolution Ψ to (23) is also a viscosity subsolution in the sense of Definition 3.5 in [26]. Moreover, a classical *solution* Ψ to (23) which also satisfies the bounds of item (iv) in Definition 2.1 is also a classical subsolution.

Remark 2.2. Under certain sufficient conditions stated in Proposition 3.1 of [73], the HJB equation (23) coincides with the large deviations rate function $S_{s,T}^\nu$ from (10) in the sense that Ψ solving (23) is given by:

$$(24) \quad \Psi(s, \nu) = \inf_{\mu \in C([s,T]; \mathcal{P}(\mathbb{R}^d))} \{S_{s,T}^\nu(\mu) + G(\mu_T)\}.$$

Stated more concretely, [73] shows that if σ is constant, b is Lipschitz on $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$, $G(\mu) \leq C \int_{\mathbb{R}^d} |z|^2 \mu(dz)$ for some $C > 0$, then any $\Psi \in C_b^{1,2}([0, T]; \mathcal{P}_2(\mathbb{R}^d))$ satisfying (23) in the classical sense must be given by (24). Note that while [73] restricts to the case of Markovian feedback controls for the infimization problem (9), we know thanks to the convexity of the running cost in the definition of the rate function that the restricted infimization problem is equivalent to the full one (see Proposition 5.8 in [9]).

Moreover, Definition 3.5 of [26] states a natural notion of viscosity solutions to (23), and Theorem 3.8 therein establishes sufficient conditions for the value function $(s, \nu) \mapsto \inf_{\mu \in C([s,T]; \mathcal{P}(\mathbb{R}^d))} \{S_{s,T}^\nu(\mu) + G(\mu_T)\}$ to be the unique viscosity solution to (23). Although that theorem does not allow for the linear growth in the control present in our drift term nor the quadratic running cost in the definition of $S_{s,T}^\nu$, it appears the results should hold under this same set of assumptions or similar; see Remarks 2.11 and 3.2 in [27].

Whether the value function is indeed the unique viscosity solution to (23) is of no consequence to the results of this paper. However, the connection between viscosity solutions of (23) and the rate function from Theorem 1.2 parallels the small-noise regime, and helped us to identify the correct PDE for designing our importance sampling scheme (compare (24) with, e.g., Equation (2.12) in [79] in the small-noise setting).

2.1. Formal derivation of the HJB equation (23). In this section, we will see formally how to arrive at the expression for the control used in our importance sampling scheme in terms of the HJB equation on Wasserstein space (23).

Following pp. 1778-1779 of [79], we let $\hat{\Phi}^N : [0, T] \times \oplus_{j=1}^N \mathbb{R}^d \rightarrow \mathbb{R}$ be the solution to the backward Kolmogorov equation with terminal condition $\exp(-NG^N(x_1, \dots, x_N)) = \exp(-NG(\mu_x^N))$ associated to the system of SDEs (1), so that $\hat{\Phi}^N$ satisfies:

$$(25) \quad \hat{\Phi}^N(t, x) = \mathbb{E}[\exp(-NG(\mu_T^{N,t,x}))], \quad x \in \mathbb{R}^{Nd}, t \in [0, T],$$

and

$$(26) \quad \begin{aligned} \partial_t \hat{\Phi}^N(t, x) + \sum_{i=1}^N b(x_i, \mu_x^N) \cdot \partial_{x_i} \hat{\Phi}^N(t, x) \\ + \frac{1}{2} \sigma \sigma^\top(x_i, \mu_x^N) : \partial_{x_i}^2 \hat{\Phi}^N(t, x) = 0, \quad t \in [0, T], x = (x_1, \dots, x_N) \in \mathbb{R}^{dN}, \\ \hat{\Phi}^N(T, x) = \exp(-NG(\mu_x^N)), x \in \mathbb{R}^{dN}. \end{aligned}$$

For fixed N , choosing controls

$$(27) \quad v_i^N(t, x_1, \dots, x_N) = \sigma^\top(x_i, \mu_x^N) \frac{\partial_{x_i} \hat{\Phi}^N(t, x_1, \dots, x_N)}{\hat{\Phi}^N(t, x_1, \dots, x_N)}$$

in the construction of $\hat{\delta}_N$ from (14) leads to a *deterministic* estimator for (3) with *zero* variance:

$$(28) \quad e^{-NG(\hat{\mu}_T^{N,s,y})} \prod_{i=1}^N Z^{i,N,s,y} = \hat{\Phi}^N(s, y_1, \dots, y_N).$$

The change of measure described above is commonly known as the Doob h -transform, and the equality (28) can be seen via an application of Itô's formula.

Next defining $\tilde{\Phi}^N(t, x_1, \dots, x_N) = -\frac{1}{N} \log \hat{\Phi}^N(t, x_1, \dots, x_N)$, we have that $\tilde{\Phi}^N$ is a solution to the second-order HJB equation:

$$(29) \quad \begin{aligned} -\partial_t \tilde{\Phi}^N(t, x_1, \dots, x_N) - \sum_{i=1}^N \left\{ b(x_i, \mu_x^N) \cdot \partial_{x_i} \tilde{\Phi}^N(t, x_1, \dots, x_N) \right. \\ \left. - \frac{N}{2} |\sigma^\top(x_i, \mu_x^N) \partial_{x_i} \tilde{\Phi}^N(t, x_1, \dots, x_N)|^2 \right. \\ \left. + \frac{1}{2} \sigma \sigma^\top(x_i, \mu_x^N) : \partial_{x_i}^2 \tilde{\Phi}^N(t, x_1, \dots, x_N) \right\} = 0, \quad x_1, \dots, x_N \in \mathbb{R}^d, t \in [0, T), \\ \tilde{\Phi}^N(T, x_1, \dots, x_N) = G(\mu_x^N), \quad x_1, \dots, x_N \in \mathbb{R}^d \end{aligned}$$

and the choice of controls from (27) is given in terms of $\tilde{\Phi}^N$ by:

$$(30) \quad v_i^N(t, x_1, \dots, x_N) = -N \sigma^\top(x_i, \mu_x^N) \partial_{x_i} \tilde{\Phi}^N(t, x_1, \dots, x_N).$$

As per (28), in (30) we have constructed a zero-variance estimator for (3) from solutions of (29):

$$(31) \quad \exp(-N \tilde{\Phi}^N(s, y_1, \dots, y_N)) = \mathbb{E}[e^{-NG(\mu_T^{N,s,y})}] = e^{-NG(\mu_T^{N,s,y})} \prod_{i=1}^N Z^{i,N,s,y},$$

so if one had access to $\tilde{\Phi}^N$ in order to construct v^N , one would already be able to compute the desired expectation explicitly.

Until this point, we have merely derived PDEs for the exact expectation that we wish to compute. To derive an importance sampling scheme that is effective for all large values of N , we assume that, given the symmetries in the coefficients and terminal condition in (29), there exists $\Phi^N : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ which is fully C^2 in the sense of Definition A.2 such that $\Phi^N(t, \mu_x^N) = \tilde{\Phi}^N(t, x_1, \dots, x_N)$ for all $t \in [0, T]$, $x_1, \dots, x_N \in \mathbb{R}^d$. Then by Proposition A.3, we get:

$$\begin{aligned} -\partial_t \Phi^N(t, \mu_x^N) - \int_{\mathbb{R}^d} b(z, \mu_x^N) \cdot \partial_\mu \Phi^N(t, \mu_x^N)[z] - \frac{1}{2} |\sigma^\top(z, \mu_x^N) \partial_\mu \Phi^N(t, \mu_x^N)[z]|^2 \\ + \frac{1}{2} \sigma \sigma^\top(z, \mu_x^N) : \partial_z \partial_\mu \Phi^N(t, \mu_x^N)[z] \\ + \frac{1}{2N} \sigma \sigma^\top(z, \mu_x^N) : \partial_\mu^2 \Phi^N(t, \mu_x^N)[z, z] \mu_x^N(dz) = 0, \quad x_1, \dots, x_N \in \mathbb{R}^d, t \in [0, T), \\ \Phi^N(T, \mu_x^N) = G(\mu_x^N), \quad x_1, \dots, x_N \in \mathbb{R}^d. \end{aligned}$$

The zero-variance optimal control from (30) can then be expressed as

$$v_i^N(t, x_1, \dots, x_N) = -\sigma^\top(x_i, \mu_x^N) \partial_\mu \Phi^N(t, \mu_x^N)[x_i].$$

Finally, supposing that the equation for Φ^N holds not only for measures of the form μ_x^N , but for all $\nu \in \mathcal{P}_2(\mathbb{R}^d)$, the above becomes:

$$(32) \quad \begin{aligned} -\partial_t \Phi^N(t, \nu) - \int_{\mathbb{R}^d} b(z, \nu) \cdot \partial_\mu \Phi^N(t, \nu)[z] - \frac{1}{2} |\sigma^\top(z, \nu) \partial_\mu \Phi^N(t, \nu)[z]|^2 \\ + \frac{1}{2} \sigma \sigma^\top(z, \nu) : \partial_z \partial_\mu \Phi^N(t, \nu)[z] \\ + \frac{1}{2N} \sigma \sigma^\top(z, \nu) : \partial_\mu^2 \Phi^N(t, \nu)[z, z] \nu(dz) = 0, \quad \nu \in \mathcal{P}_2(\mathbb{R}^d), t \in [0, T), \\ \Phi^N(T, \nu) = G(\nu), \quad \nu \in \mathcal{P}_2(\mathbb{R}^d). \end{aligned}$$

As $N \rightarrow \infty$, we expect that, up to leading order, solutions to (32) should be well approximated by solutions to the *zero-viscosity* HJB equation (23), which we arrive at by simply setting the $O(1/N)$ term in (32) to 0, leading us finally to the HJB equation on Wasserstein space (23). Rigorous justifications for this limit in the case where there is a common driving noise between the particles can be found in [56].

Our discussion here parallels the construction in Section 3.1 of [79] of a log-efficient control in the small-noise setting, in which the authors start from the deterministic scheme given by the Doob h -transform and take a formal zero-viscosity limit in the small-noise parameter. Given the expression (31), we expect then that

at least the $O(N)$ terms in an expression for $\log R(\hat{\delta}_N)$ with the choice of controls from Theorem 2.4 should vanish, since they are 0 when considering the controls from (30). The log-efficiency proved in Theorem 2.4 shows that this is not only the case for classical solutions to (23), but even certain classical subsolutions.

2.2. Statement of main results. We are now ready to state our main results, Theorems 2.4 and 2.5, the proofs of which are contained in Section 5. The first is on designing a log-efficient importance sampling scheme, and parallels Theorem 8.1 in [38] and Theorem 4.1 in [36]. It requires the additional two assumptions:

Assumption 2.3. *Suppose:*

- (B1) G in equations (3) and (23) is bounded and continuous on $\mathcal{P}(\mathbb{R}^d)$
- (B2) σ is bounded.

Theorem 2.4. *Suppose Ψ is a classical subsolution to (23) in the sense of Definition 2.1 and that Assumptions 1.1 and 2.3 hold. Then the importance sampling scheme from (14) with the choice of controls*

$$(33) \quad v_i^N(t, x_1, \dots, x_N) = -\sigma^\top \left(x_i, \frac{1}{N} \sum_{j=1}^N \delta_{x_j} \right) \partial_\mu \Psi \left(t, \frac{1}{N} \sum_{j=1}^N \delta_{x_j} \right) [x_i]$$

has the property that

$$\lim_{N \rightarrow \infty} -\frac{1}{N} \log R(\hat{\delta}_N) \geq \Psi(s, \nu) - \gamma_1,$$

where $R(\hat{\delta}_N)$ is as in (17), γ_1 is as in (18), s is the initial time in (1), and ν is as in (A1). In particular, if $\Psi(s, \nu) \geq \gamma_2 - \gamma_1$, then this importance sampling scheme performs better than the standard Monte Carlo estimator in the sense that it admits an expansion of the form (20) such that $\hat{\gamma} \leq 2\gamma_1 - \gamma_2$ (recalling here (19)).

Moreover, if $\Psi(s, \nu) = \gamma_1$, then this importance sampling scheme is log-efficient in the sense of Definition 1.4.

As we discussed in Section 1, [79] provides sufficient conditions in terms of a concept of “uniform log-efficiency” for the log-efficient importance sampling estimator in their small-noise setting to have vanishing relative error (Definition 1.5). It is not immediately clear how to verify such conditions in our setting; their proof relies heavily on local regularity results from the theory of (standard) first-order HJB equations which we cannot directly apply to our infinite-dimensional HJB equations. In [75], the author takes a different approach, where an asymptotic expansion in the small-noise parameter of the HJB equation associated to the Doob h -transform (see (32) herein) is used to establish that their log-efficient importance sampling estimator in fact has vanishing relative error. This is done by establishing an expression of the form (21), where the a_k ’s are expressed in terms of solutions to different PDEs. Vanishing relative error can then be established by showing $a_0 = a_1 = 0$.

In order to carry out a similar analysis here, we will likewise derive PDEs whose solutions can be used to carry out an asymptotic expansion analysis in Theorem 2.5. Firstly, let ϕ_0 denote a classical $C_b^{1,2}([0, T] \times \mathcal{P}_2(\mathbb{R}^d))$ solution to (23). Using the control from (33) with ϕ_0 in the place of Ψ , we obtain an expression for the numerator of $R(\hat{\delta}_N)$ from (17), which results from applying Girsanov’s theorem, the Feynman-Kac’s formula, and a log transformation. Concretely, letting Ξ^N satisfy:

$$(34) \quad \begin{aligned} & \partial_t \Xi^N(t, \nu) + \int_{\mathbb{R}^d} [b(z, \nu) - \sigma(z, \nu)v(t, \nu, z)] \cdot \partial_\mu \Xi^N(t, \nu)[z] + \frac{1}{2} \sigma \sigma^\top(z, \nu) : \partial_z \partial_\mu \Xi^N(t, \nu)[z] \nu(dz) \\ & - \int_{\mathbb{R}^d} \frac{1}{2} |\sigma^\top(z, \nu) \partial_\mu \Xi^N(t, \nu)[z]|^2 + |v(t, \nu, z)|^2 \nu(dz) \\ & + \frac{1}{2N} \int_{\mathbb{R}^d} \sigma \sigma^\top(z, \nu) : \partial_\mu^2 \Xi^N(t, \nu)[z, z] \nu(dz) = 0, \quad t \in [0, T], \nu \in \mathcal{P}_2(\mathbb{R}^d), \\ & \Xi^N(T, \nu) = 2G(\nu), \quad \nu \in \mathcal{P}_2(\mathbb{R}^d), \\ & v(t, \nu, x) = -\sigma^\top(x, \nu) \partial_\mu \phi_0(t, \nu)[x], \quad t \in [0, T], x \in \mathbb{R}^d, \nu \in \mathcal{P}_2(\mathbb{R}^d), \end{aligned}$$

we have

$$\Xi^N(t, \mu_x^N) = -\frac{1}{N} \log \mathbb{E} \left[\exp(-2NG(\hat{\mu}_T^{N,t,x})) \prod_{i=1}^N (Z^{i,N,t,x})^2 \right], \quad t \in [0, T], x \in \mathbb{R}^{dN}.$$

Note that this corresponds to Equation (3.3) in the small-noise setting of [75].

Expanding $\Xi^N(t, \mu) = \xi_0(t, \mu) + o(1)$, we expect that ξ_0 solves the “zero viscosity HJB equation” obtained by setting the final $\mathcal{O}(1/N)$ term in (34) equal to 0, that is:

$$\begin{aligned} (35) \quad & \partial_t \xi_0(t, \mu) + \int_{\mathbb{R}^d} [b(z, \mu) - \sigma(z, \mu)v(t, \mu, z)] \cdot \partial_\mu \xi_0(t, \mu)[z] \\ & + \frac{1}{2} \sigma \sigma^\top(z, \mu) : \partial_z \partial_\mu \xi_0(t, \mu)[z] \mu(dz) - \int_{\mathbb{R}^d} \frac{1}{2} |\sigma^\top(z, \mu) \partial_\mu \xi_0(t, \mu)[z]|^2 \\ & + |\sigma^\top(z, \mu) \partial_\mu \phi_0(t, \mu)[z]|^2 \mu(dz) = 0, \quad t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d), \\ & \xi_0(T, \mu) = 2G(\mu), \quad \mu \in \mathcal{P}_2(\mathbb{R}^d), \end{aligned}$$

where v is as in (34). Note that, by the same logic, expanding Φ^N from (32) as $\Phi^N(t, \mu) = \phi_0(t, \mu) + o(1)$, we expect that $\phi_0(t, \mu)$ solves (23).

We will also consider higher order terms in the following series expansions of Φ^N from (32) and Ξ^N from (34):

$$(36) \quad \Phi^N(t, \mu) = \phi_0(t, \mu) + \frac{1}{N} \phi_1(t, \mu) + \frac{1}{N^2} \phi_2(t, \mu) + \dots + \frac{1}{N^k} \phi_k(t, \mu) + o(1/N^k),$$

and

$$(37) \quad \Xi^N(t, \mu) = \xi_0(t, \mu) + \frac{1}{N} \xi_1(t, \mu) + \frac{1}{N^2} \xi_2(t, \mu) + \dots + \frac{1}{N^k} \xi_k(t, \mu) + o(1/N^k)$$

for some $k \in \mathbb{N}$ and all $t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d)$.

Matching terms of the same order upon inserting this ansatz into the equations for Φ^N and Ξ^N , we expect that

$$\begin{aligned} (38) \quad & \partial_t \phi_k(t, \mu) + \int_{\mathbb{R}^d} b(z, \mu) \cdot \partial_\mu \phi_k(t, \mu)[z] + \frac{1}{2} \sigma \sigma^\top(z, \mu) : \partial_z \partial_\mu \phi_k(t, \mu)[z] \mu(dz) \\ & + \int_{\mathbb{R}^d} \frac{1}{2} \sigma \sigma^\top(z, \mu) : \partial_\mu^2 \phi_{k-1}(t, \mu)[z, z] \mu(dz) \\ & - \int_{\mathbb{R}^d} \sum_{i+j=k, 0 \leq i < j} \langle \sigma^\top(z, \mu) \partial_\mu \phi_i(t, \mu)[z], \sigma^\top(z, \mu) \partial_\mu \phi_j(t, \mu)[z] \rangle \\ & + \frac{1}{2} |\sigma^\top(z, \mu) \partial_\mu \phi_{k/2}(t, \mu)[z]|^2 \mathbb{1}_{k/2 \in \mathbb{N}} \mu(dz) = 0, \quad \mu \in \mathcal{P}_2(\mathbb{R}^d), t \in [0, T], \\ & \phi_k(T, \mu) = 0, \quad \mu \in \mathcal{P}_2(\mathbb{R}^d) \end{aligned}$$

and

$$\begin{aligned} (39) \quad & \partial_t \xi_k(t, \mu) + \int_{\mathbb{R}^d} [b(z, \mu) + \sigma(z, \mu) \sigma^\top(z, \mu) \partial_\mu \phi_0(t, \mu)[z]] \cdot \partial_\mu \xi_k(t, \mu)[z] \\ & + \frac{1}{2} \sigma \sigma^\top(z, \mu) : \partial_z \partial_\mu \xi_k(t, \mu)[z] \mu(dz) \\ & + \int_{\mathbb{R}^d} \frac{1}{2} \sigma \sigma^\top(z, \mu) : \partial_\mu^2 \xi_{k-1}(t, \mu)[z, z] \mu(dz) \\ & - \int_{\mathbb{R}^d} \sum_{i+j=k, 0 \leq i < j} \langle \sigma^\top(z, \mu) \partial_\mu \xi_i(t, \mu)[z], \sigma^\top(z, \mu) \partial_\mu \xi_j(t, \mu)[z] \rangle \\ & + \frac{1}{2} |\sigma^\top(z, \mu) \partial_\mu \xi_{k/2}(t, \mu)[z]|^2 \mathbb{1}_{k/2 \in \mathbb{N}} \mu(dz) = 0, \quad \mu \in \mathcal{P}_2(\mathbb{R}^d), t \in [0, T], \\ & \xi_k(T, \mu) = 0, \quad \mu \in \mathcal{P}_2(\mathbb{R}^d) \end{aligned}$$

for all $k \geq 1$. Note that these correspond to the PDEs found in Theorem 3.3 of [75].

Both of these equations fit the form of Equation (1.2) in [23], so under sufficient regularity assumptions we would have by Theorem 3.8 therein that for $k \geq 1$:

(40)

$$\begin{aligned} \phi_k(t, \mu) &= \mathbb{E} \left[\int_t^T \frac{1}{2} \sigma \sigma^\top(Z_\tau^{t, \mu}, \mathcal{L}(Z_\tau^{t, \mu})) : \partial_\mu^2 \phi_{k-1}(\tau, \mathcal{L}(Z_\tau^{t, \mu})) [Z_\tau^{t, \mu}, Z_\tau^{t, \mu}] \right. \\ &\quad - \sum_{i+j=k, 0 \leq i < j < k} \langle \sigma^\top(Z_\tau^{t, \mu}, \mathcal{L}(Z_\tau^{t, \mu})) \partial_\mu \phi_i(\tau, \mathcal{L}(Z_\tau^{t, \mu})) [Z_\tau^{t, \mu}], \sigma^\top(Z_\tau^{t, \mu}, \mathcal{L}(Z_\tau^{t, \mu})) \partial_\mu \phi_j(\tau, \mathcal{L}(Z_\tau^{t, \mu})) [Z_\tau^{t, \mu}] \rangle \\ &\quad \left. - \frac{1}{2} |\sigma^\top(Z_\tau^{t, \mu}, \mathcal{L}(Z_\tau^{t, \mu})) \partial_\mu \phi_{k/2}(\tau, \mathcal{L}(Z_\tau^{t, \mu})) [Z_\tau^{t, \mu}]|^2 \mathbb{1}_{k/2 \in \mathbb{N}} d\tau \right], \\ dZ_\tau^{t, \mu} &= \left[b(Z_\tau^{t, \mu}, \mathcal{L}(Z_\tau^{t, \mu})) - \sigma \sigma^\top(Z_\tau^{t, \mu}, \mathcal{L}(Z_\tau^{t, \mu})) \partial_\mu \phi_0(\tau, \mathcal{L}(Z_\tau^{t, \mu})) [Z_\tau^{t, \mu}] \right] d\tau + \sigma(Z_\tau^{t, \mu}, \mathcal{L}(Z_\tau^{t, \mu})) dW_\tau, Z_t^{t, \mu} \sim \mu \end{aligned}$$

and

$$\begin{aligned} \xi_k(t, \mu) &= \mathbb{E} \left[\int_t^T \frac{1}{2} \sigma \sigma^\top(Y_\tau^{t, \mu}, \mathcal{L}(Y_\tau^{t, \mu})) : \partial_\mu^2 \xi_{k-1}(\tau, \mathcal{L}(Y_\tau^{t, \mu})) [Y_\tau^{t, \mu}, Y_\tau^{t, \mu}] \right. \\ &\quad - \sum_{i+j=k, 0 \leq i < j < k} \langle \sigma^\top(Y_\tau^{t, \mu}, \mathcal{L}(Y_\tau^{t, \mu})) \partial_\mu \xi_i(\tau, \mathcal{L}(Y_\tau^{t, \mu})) [Y_\tau^{t, \mu}], \\ &\quad \left. \sigma^\top(Y_\tau^{t, \mu}, \mathcal{L}(Y_\tau^{t, \mu})) \partial_\mu \xi_j(\tau, \mathcal{L}(Y_\tau^{t, \mu})) [Y_\tau^{t, \mu}] \rangle \right. \\ &\quad \left. - \frac{1}{2} |\sigma^\top(Y_\tau^{t, \mu}, \mathcal{L}(Y_\tau^{t, \mu})) \partial_\mu \xi_{k/2}(\tau, \mathcal{L}(Y_\tau^{t, \mu})) [Y_\tau^{t, \mu}]|^2 \mathbb{1}_{k/2 \in \mathbb{N}} d\tau \right] \\ dY_\tau^{t, \mu} &= \left[b(Y_\tau^{t, \mu}, \mathcal{L}(Y_\tau^{t, \mu})) + \sigma \sigma^\top(Y_\tau^{t, \mu}, \mathcal{L}(Y_\tau^{t, \mu})) \partial_\mu \phi_0(\tau, \mathcal{L}(Y_\tau^{t, \mu})) [Y_\tau^{t, \mu}], \right. \\ &\quad \left. - \sigma \sigma^\top(Y_\tau^{t, \mu}, \mathcal{L}(Y_\tau^{t, \mu})) \partial_\mu \xi_0(\tau, \mathcal{L}(Y_\tau^{t, \mu})) [Y_\tau^{t, \mu}] \right] d\tau + \sigma(Y_\tau^{t, \mu}, \mathcal{L}(Y_\tau^{t, \mu})) dW_\tau, Y_t^{t, \mu} \sim \mu. \end{aligned} \tag{41}$$

As we will see, the a_k 's from (21) will be realized in terms of the solutions to the PDEs (38) and (39) via the representation:

$$a_k = 2\phi_k(s, \mu_y^N) - \xi_k(s, \mu_y^N), \quad k \in \mathbb{N},$$

where $s \in [0, T]$ is the initial time for the particles (1), $y \in \oplus_{i=1}^\infty \mathbb{R}^d$ encodes their initial conditions, and we use the empirical measure notation from the end of Section 1.1. See the proof of Theorem 2.4 in Subsection 5.2 for more details.

In order to generate a single realization of $\mu^{N, s, y}$, we need to simulate a system of N SDEs. Thus, with $R(\hat{\delta}_N)$ as in (17), the quantity

$$T(N) := N[R(\hat{\delta}_N) - 1], \tag{42}$$

is proportional to the maximum number of particles which need to be simulated (and therefore the computational work required) in order for the relative error from (16) to be below a certain threshold. Put explicitly, to achieve a certain relative error c , we require at least $\frac{R(\hat{\delta}_N)-1}{c^2}$ samples of $(\hat{\mu}^{N, s, y}, Z^{1, N, s, y}, \dots, Z^{N, N, s, y})$. Therefore, the total number of particles simulated will need to be at least $T(N)/c^2$. Note, of course, that even if $R(\hat{\delta}_N) = 1$, that we will need to simulate at least N particles to get a realization of $\hat{\mu}^{N, s, y}$.

This motivates us to go beyond studying ϕ_0, ξ_0, ϕ_1 , and ξ_1 in order to prove $a_0 = a_1 = 0$ in (21), and to further study a_2 . We find that if we assume third-order expansions of the form (36) and (37) that have sufficient regularity properties, then not only does the relative error vanish, but it vanishes fast enough that $T(N)$ is bounded as $N \rightarrow \infty$. This further suggests that the relative error of our importance sampling scheme is expected to vanish like $\mathcal{O}(1/\sqrt{N})$ as $N \rightarrow \infty$ if the coefficients of the interacting particle system (1) and the target function G from (3) are sufficiently regular.

Theorem 2.5. *Let Assumption 2.3 and Assumption 1.1 (A1)-(A4) hold. Assume also that $|b(x, \mu)|^2 \leq C(1 + |x|^2 + \int |z|^2 \mu(dz))$ for all $x \in \mathbb{R}^d$ and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. Consider $\Phi^N, \Xi^N : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ from (32) and (34) respectively. Suppose these are unique $C_b^{1,2}([0, T] \times \mathcal{P}_2(\mathbb{R}^d))$ solutions admitting expansions of*

the form (36) and (37) respectively up to $k = 1$, and that $\phi_0, \xi_0, \phi_1, \xi_1$, are the unique $C_b^{1,2}([0, T] \times \mathcal{P}_2(\mathbb{R}^d))$ solutions to (23), (35), (38), (39) with $k = 1$ respectively, with ϕ_1, ξ_1 admitting the stochastic representations (40), (41) with $k = 1$ respectively. Then the importance sampling scheme from Theorem 2.4 with ϕ_0 in the place of Ψ has vanishing relative error in the sense of Definition 1.5.

If further we assume expansions of the form (36) and (37) up to $k = 2$ and that ϕ_2, ξ_2 are the unique $C_b^{1,2}([0, T] \times \mathcal{P}_2(\mathbb{R}^d))$ solutions to (38), (39) with $k = 2$ respectively admitting the representations (40), (41) with $k = 2$ respectively, then

$$(43) \quad \lim_{N \rightarrow \infty} T(N) = \mathbb{E} \left[\int_s^T |\sigma^\top(\hat{X}_t^{u,s,\nu}, \mathcal{L}(\hat{X}_t^{u,s,\nu})) \partial_\mu \phi_1(t, \mathcal{L}(\hat{X}_t^{u,s,\nu})) [\hat{X}_t^{u,s,\nu}]|^2 dt \right]$$

where $\hat{X}^{u,s,\nu}$ is as in (8) with $u(t) = -\sigma^\top(\hat{X}_t^{u,s,\nu}, \mathcal{L}(\hat{X}_t^{u,s,\nu})) \partial_\mu \phi_0(t, \mathcal{L}(\hat{X}_t^{u,s,\nu})) [\hat{X}_t^{u,s,\nu}]$, $t \in [s, T]$.

Remark 2.6. The results of Theorems 2.4 and 2.5 can easily be extended to the situation where we modify the desired expectation (3) to

$$\mathbb{E} \left[\exp \left(-N \left[G(\mu_T^{N,s,y}) + \int_s^T f(\mu_t^{N,s,y}) dt \right] \right) \right]$$

for sufficiently regular $f : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$. In this case one takes $F(\mu) = G(\mu_T) + \int_s^T f(\mu_t^{N,s,y}) dt$ in (9) and modifies H_0 from (22) to

$$H_0(\mu, p, \Gamma) = - \int_{\mathbb{R}^d} b(x, \mu) \cdot p(x) - \frac{1}{2} p(x) \cdot [\sigma \sigma^\top(x, \mu) p(x)] + \frac{1}{2} \sigma \sigma^\top(x, \mu) : \Gamma(x) \mu(dx) - f(\mu).$$

One should also be able to extend these methods in order to design importance sampling schemes for the probabilities of rare events. That is, rather than estimating (3), one may want to estimate

$$\mathbb{P}(\mu_T^{N,s,y} \in A)$$

for some $A \subset \mathcal{P}(\mathbb{R}^d)$. This formally corresponds to taking

$$G(\mu) = \begin{cases} +\infty, & \mu \in A^c \\ 0, & \mu \in A \end{cases}$$

in (3) and (23). It is well known that this extension can be made in the small-noise setting — see, e.g., [36] Proposition 4.2. We refrain from performing this analysis here, but we plan to extend the methods presented to not only probabilities at finite time, but also exit probabilities and mean first passage times in future work. Thanks to the many parallels to the small-noise setting, we expect that designing importance sampling schemes related to exit events will require the study of HJB equations of the form (23) restricted to some subset of the space $\mathcal{P}_2(\mathbb{R}^d)$ with boundary conditions (compare with, e.g. Equations (2.4)-(2.5) in [35]). Such equations are already beginning to be studied in the context of optimal stopping problems for McKean-Vlasov equations — see, e.g., [76].

Remark 2.7. The additional assumption that b has at most linear growth in the statement of Theorem 2.5 ensures that the unique solutions of the Feynman-Kac equations (26) and (62) are given by their appropriate stochastic representations, i.e. the denominator and numerator of $R(\hat{\delta}_N)$ from (17), respectively. This can of course hold under weaker conditions.

Moreover, although we do not necessarily assume (A5) for Theorem 2.5, it is unlikely that there are situations in which the existence and uniqueness of classical solutions to the HJB equations (32) and (34) hold for which (A5) doesn't hold. In particular, the required assumptions for Theorem 2.5 should be much stricter than those imposed for Theorems 1.2 and 2.4.

The unique stochastic representations assumed for ϕ_k, ξ_k , $k = 1, 2$ solving the PDEs (38) and (39) are known to hold under fairly weak assumptions — see Theorem 3.8 in [23]. Moreover, these representations for the prefactor terms in the expansion yield parallels to those found in the small-noise setting — see the discussion towards the end of Subsection 5.2. The reason for the assumptions in Theorem 2.5 being stated as such is that little is available in the current literature in terms of sufficient conditions for existence and uniqueness of classical solutions to (23), let alone on obtaining the desired formal expansions in N . Indeed, even the

convergence of Φ^N to ϕ_0 (corresponding to the first-order expansion) has only been studied in the case where the particles have common noise — see [56].

The assumptions made in terms of this expansion essentially mimic the conclusion of Theorem 3.3 in [75], which is a consequence of Theorem 5.1 in [46]. We expect that analogous conditions to those of these theorems should be able to be found to be sufficient for the expansion analysis of Theorem 2.5 to go through. This is an interesting avenue for future research.

More generally, on p. 1781-1782 of [79], it is discussed how if a discontinuity of their optimal control is anything more exotic than a single curve, they are unable to prove log-efficiency. In general the interplay between regularity of solutions to the zero-viscosity HJB equation and the properties of the relative error of the importance sampling scheme is an interesting open problem even in the small-noise setting, though there is a wealth of numerical evidence suggesting that in many situations importance sampling schemes derived from non-differentiable subsolutions of the zero-viscosity HJB equation may have bounded or even vanishing relative error.

We contribute to these numerical findings in the context of weakly interacting diffusions in the numerical examples of Subsection 4.2, where we construct subsolutions for modifications of the linear-quadratic regime discussed in Section 3 such that the corresponding optimal controls (33) are discontinuous. We find that, depending on the nature of the discontinuity, the relative error of our importance sampling scheme can be expected to grow sublinearly, or even vanish — see Tables 3 and 2, respectively.

3. A CLASS OF EXAMPLES: THE LINEAR-QUADRATIC REGIME

We now consider a class of HJB equations of the form (23) for which explicit solutions are known. Consider the setting where

$$(44) \quad \begin{aligned} b(x, \mu) &= b_0 + Bx + \bar{B} \int_{\mathbb{R}^d} z \mu(dz), \\ \sigma(x, \mu) &= \sigma, \\ G(\mu) &= \int_{\mathbb{R}^d} z^\top P_2 z + p_1 \cdot z \mu(dz) + \left[\int_{\mathbb{R}^d} z \mu(dz) \right]^\top \bar{P}_2 \int_{\mathbb{R}^d} z \mu(dz) + p_2 \end{aligned}$$

for $b_0, p_1 \in \mathbb{R}^d$, $\sigma \in \mathbb{R}^{d \times 1}$ (so $m = 1$), $p_2 \in \mathbb{R}$, and $P_2, B, \bar{B}, \bar{P}_2 \in \mathbb{R}^{d \times d}$ such that $P_2, \bar{P}_2$ are symmetric and positive semidefinite. By Section 4 of [73], the unique classical solution to (23) is given by

$$\Psi(t, \mu) = \int_{\mathbb{R}^d} z^\top \Lambda(t) z \mu(dz) + \left[\int_{\mathbb{R}^d} z \mu(dz) \right]^\top [\Gamma(t) - \Lambda(t)] \int_{\mathbb{R}^d} z \mu(dz) + \gamma(t) \cdot \int_{\mathbb{R}^d} z \mu(dz) + \chi(t),$$

where:

$$(45) \quad \begin{aligned} \dot{\Lambda}(t) + \Lambda(t)B + B^\top \Lambda(t) - 2\Lambda(t)\sigma\sigma^\top \Lambda^\top(t) &= 0, \Lambda(T) = P_2, \\ \dot{\Gamma}(t) + \Gamma(t)[B + \bar{B}] + [B + \bar{B}]^\top \Gamma(t) - 2\Gamma(t)\sigma\sigma^\top \Gamma^\top(t) &= 0, \Gamma(T) = P_2 + \bar{P}_2, \\ \dot{\gamma}(t) + [B + \bar{B}]^\top \gamma(t) - 2\Gamma(t)\sigma\sigma^\top \gamma(t) + 2\Gamma(t)b_0 &= 0, \gamma(T) = p_1, \\ \dot{\chi}(t) - \frac{1}{2}\gamma^\top(t)\sigma\sigma^\top \gamma(t) + \gamma(t) \cdot b_0 + \sigma^\top \Lambda(t)\sigma &= 0, \chi(T) = p_2. \end{aligned}$$

The above Riccati equations admit unique solutions $\gamma \in \mathbb{R}^d, \chi \in \mathbb{R}, \Lambda, \Gamma \in \mathbb{R}^{d \times d}$ such that Λ, Γ are symmetric and positive semidefinite for all t .

Then, using:

$$\partial_\mu \Psi(t, \mu)[z] = 2\Lambda(t)z + 2[\Gamma(t) - \Lambda(t)] \int_{\mathbb{R}^d} z \mu(dz) + \gamma(t)$$

the controls from (33) in Theorem 2.4 are given by

$$(46) \quad \begin{aligned} v_i^N(t, x_1, \dots, x_N) &= -\sigma^\top \partial_\mu \Psi(t, \mu_x^N)[x_i] \\ &= -\sigma^\top \left[2\Lambda(t)x_i + 2[\Gamma(t) - \Lambda(t)] \int_{\mathbb{R}^d} z \mu_x^N(dz) + \gamma(t) \right]. \end{aligned}$$

Note that using Remark 5.2 in [5], we can extend to the case where $m > 1$, though we refrain from doing so for simplicity. We could also allow for time dependence in the coefficients, but for the sake of simplicity we do not make this extension here.

Despite the fact that G from (44) does not satisfy the Assumption 2.3 (B1), as we will see in Remark 3.6, the controls (46) yield not only the log-efficiency proved in Theorem 1.4, but in fact zero relative error for all N .

Our first example from the linear-quadratic regime is chosen to have symmetries such that the relative error of the standard Monte Carlo estimator δ_N and of the importance sampling estimator $\hat{\delta}_N$ (14) with the choice of control (46) are both easily computable.

Example 3.1. Let $G(\mu) = \int_{\mathbb{R}} x\mu(dx)$, $d = m = 1$, and

$$X_t^{i,N,s,y} = y_i + \int_s^t X_\tau^{i,N,s,y} - \frac{1}{N} \sum_{j=1}^N X_\tau^{j,N,s,y} d\tau + \sigma W_t^i.$$

Symmetries in this problem allow us to calculate the target expectation (3), as well as the Monte Carlo relative error $\rho(\delta_N)$ and importance sampling relative error $\rho(\hat{\delta}_N)$ (16), explicitly.

We have

$$\begin{aligned} \mathbb{E} \left[\exp(-NG(\mu_T^N)) \right] &= \mathbb{E} \left[\exp \left(- \sum_{i=1}^N \left\{ y_i + \int_s^T X_t^{i,N,s,y} - \frac{1}{N} \sum_{j=1}^N X_t^{j,N,s,y} dt + \sigma W_T^i \right\} \right) \right] \\ &= \exp \left(- \sum_{i=1}^N y_i \right) \mathbb{E} \left[\exp \left(- \sigma W_T^1 \right) \right]^N \\ &= \exp \left(- \sum_{i=1}^N y_i \right) \exp(N\sigma^2(T-s)/2), \end{aligned}$$

and similarly:

$$\begin{aligned} \mathbb{E} \left[\exp(-2NG(\mu_T^N)) \right] &= \mathbb{E} \left[\exp \left(-2 \sum_{i=1}^N \left\{ y_i + \int_s^T X_t^{i,N,s,y} - \frac{1}{N} \sum_{j=1}^N X_t^{j,N,s,y} dt + \sigma W_T^i \right\} \right) \right] \\ &= \exp \left(-2 \sum_{i=1}^N y_i \right) \exp(2N\sigma^2(T-s)), \end{aligned}$$

where in the last step we have used the formula for the moment generating function of a normal random variable. Thus:

$$\rho(\delta_N) = \frac{1}{\sqrt{M}} \sqrt{\frac{\exp(2N\sigma^2(T-s))}{\exp(N\sigma^2(T-s))}} - 1 = \frac{1}{\sqrt{M}} \sqrt{\exp(N\sigma^2(T-s))} - 1.$$

Now, to construct the control (46) for our importance sampling scheme, we have that in the setup of (44), $b_0 = P_2 = \bar{P}_2 = \bar{p}_1 = p_2 = 0$, $B = p_1 = 1$, and $\bar{B} = -1$. Thus the solution to the system of ODEs (45) is $\Lambda(t) = \Gamma(t) = 0$, $\gamma(t) = 1$, and $\chi(t) = \frac{1}{2}\sigma^2(t-T)$, and so the solution to the HJB equation (23) is

$$\Psi(t, \mu) = \int_{\mathbb{R}} x\mu(dx) + \frac{1}{2}\sigma^2(t-T).$$

Our control from Theorem 2.4 is given by

$$\bar{v}(s, x, \mu) \equiv -\sigma.$$

Then, to compute the relative error of the importance sampling scheme, we have

$$\begin{aligned}\rho(\hat{\delta}_N) &= \frac{1}{\sqrt{M}} \sqrt{\frac{\mathbb{E} \left[\exp \left(-2 \sum_{i=1}^N \hat{X}_T^{i,N,s,y} \right) \prod_{i=1}^N \exp \left(2\sigma \hat{W}_T^i - \int_s^T \sigma^2 dt \right) \right]}{\mathbb{E} \left[\exp(-NG(\mu^N)) \right]^2} - 1} \\ &= \frac{1}{\sqrt{M}} \sqrt{\frac{\mathbb{E} \left[\exp \left(-2 \sum_{i=1}^N (\hat{X}_T^{i,N,s,y} - \sigma \hat{W}_T^i) \right) \right] \exp \left(-N\sigma^2(T-s) \right)}{\mathbb{E} \left[\exp(-NG(\mu_T^N)) \right]^2} - 1}\end{aligned}$$

where

$$\hat{X}_t^{i,N,s,y} = y_i + \int_s^t \hat{X}_\tau^{i,N,s,y} - \frac{1}{N} \sum_{j=1}^N \hat{X}_\tau^{j,N,s,y} - \sigma^2 d\tau + \sigma \hat{W}_t^i.$$

Continuing, we have

$$\begin{aligned}\rho(\hat{\delta}_N) &= \frac{1}{\sqrt{M}} \sqrt{\frac{\mathbb{E} \left[\exp \left(-2 \sum_{i=1}^N \left\{ y_i + \int_s^T \hat{X}_t^{i,N,s,y} - \frac{1}{N} \sum_{j=1}^N \hat{X}_t^{j,N,s,y} - \sigma^2 dt \right\} \right) \right] \exp \left(-N\sigma^2(T-s) \right)}{\mathbb{E} \left[\exp(-NG(\mu_T^N)) \right]^2} - 1} \\ &= \frac{1}{\sqrt{M}} \sqrt{\frac{\mathbb{E} \left[\exp \left(-2 \sum_{i=1}^N (y_i - (T-s)\sigma^2) \right) \right] \exp \left(-N\sigma^2(T-s) \right)}{\mathbb{E} \left[\exp(-NG(\mu_T^N)) \right]^2} - 1} \\ &= \frac{1}{\sqrt{M}} \sqrt{\frac{\exp \left(-2 \sum_{i=1}^N y_i \right) \exp \left(2N\sigma^2(T-s) \right) \exp \left(-N\sigma^2(T-s) \right)}{\exp \left(-2 \sum_{i=1}^N y_i \right) \exp(N\sigma^2(T-s))} - 1} \\ &= 0.\end{aligned}$$

This shows that where the Monte Carlo estimator has relative error which increases exponentially as $N \rightarrow \infty$ or $T \rightarrow \infty$, our importance sampling estimator actually has zero variance for all N and T .

The reader familiar with small-noise importance sampling schemes for SDEs may realize at this point that the controls computed in Example 3.1 can also be derived using the classical small-noise theory. In that vein, we make the following remark relating a special case of our importance sampling method to the classical small-noise importance sampling scheme:

Remark 3.2. In the linear-quadratic regime, $Y_t^{N,s,y} = \frac{1}{N} \sum_{i=1}^N X_t^{i,N,s,y}$ satisfies the small-noise SDE:

$$\begin{aligned}dY_t^{N,s,y} &= [b_0 + (B + \bar{B})Y_t^{N,s,y}]dt + \frac{\sigma}{\sqrt{N}}d\mathcal{W}_t^N, \\ Y_s^{N,s,y} &= \frac{1}{N} \sum_{i=1}^N y_i,\end{aligned}$$

where for each N , $\mathcal{W}^N$ is the standard one-dimensional Brownian motion $\mathcal{W}_t^N = \frac{1}{\sqrt{N}} \sum_{i=1}^N W_t^i$. Thus, when $G(\mu) = g(\int_{\mathbb{R}^d} x\mu(dx))$ for some $g : \mathbb{R}^d \rightarrow \mathbb{R}^d$, estimating functionals of the form

$$\mathbb{E}[\exp(-NG(\mu_T^{N,s,y}))] = \mathbb{E}[\exp(-Ng(Y_T^{N,s,y}))]$$

can also be done using standard small-noise importance sampling for SDEs, i.e., the appropriate control is $\hat{v}(t, y) = -\sigma^\top \partial_y \psi(t, y)$ where ψ satisfies the standard first-order HJB equation:

$$(47) \quad \begin{aligned} -\partial_t \psi(t, y) - [b_0 + (B + \bar{B})y] \cdot \partial_y \psi(t, y) + \frac{1}{2} |\sigma^\top \partial_y \psi(t, y)|^2 &= 0, \quad y \in \mathbb{R}^d, t \in [0, T], \\ \psi(T, y) &= g(y), \quad y \in \mathbb{R}^d \end{aligned}$$

— see, e.g., equations (1.12) and (2.11) in [79]. The resulting importance sampling estimator is given by

$$(48) \quad \hat{\delta}_N = \frac{1}{M} \sum_{j=1}^M \exp(-Ng(\hat{Y}_T^{N,s,y,j})) \exp\left(-\sqrt{N} \int_s^T \hat{v}(t, \hat{Y}_t^{N,s,y,j}) \cdot d\hat{\mathcal{W}}_t^N - \frac{N}{2} \int_s^T |\hat{v}(t, \hat{Y}_t^{N,s,y,j})|^2 dt\right)$$

where $\hat{Y}_t^{N,s,y,j}$ denotes the j 'th realization of $\hat{Y}_t^{N,s,y}$ solving

$$d\hat{Y}_t^{N,s,y} = [b_0 + (B + \bar{B})\hat{Y}_t^{N,s,y} + \sigma \hat{v}(t, \hat{Y}_t^{N,s,y})]dt + \frac{\sigma}{\sqrt{N}} d\hat{\mathcal{W}}_t^N, \quad \hat{Y}_s^{N,s,y} = \frac{1}{N} \sum_{i=1}^N y_i.$$

In our linear-quadratic regime (44), the choice of parameters $P_2 = 0, p_1 = [p, \dots, p]^\top$ (so that $g(y) = py + y^\top \bar{P}_2 y + p_2$) is therefore covered by the small-noise theory.

To show that the two importance sampling estimators (48) and (14) are always identical in this case, we make the ansatz $\Psi(t, \nu) = \psi\left(t, \int_{\mathbb{R}^d} z\nu(dz)\right)$. Then $\partial_\mu \Psi(t, \nu)[z] = \partial_y \psi\left(t, \int_{\mathbb{R}^d} x\nu(dx)\right)$, $\partial_z \partial_\mu \Psi(t, \nu)[z] \equiv 0$, and (23) is given by:

$$\begin{aligned} -\partial_t \psi\left(t, \int_{\mathbb{R}^d} z\nu(dz)\right) - \left(b_0 + [B + \bar{B}] \int_{\mathbb{R}^d} z\nu(dz)\right) \cdot \partial_y \psi\left(t, \int_{\mathbb{R}^d} z\nu(dz)\right) + \frac{1}{2} \left| \sigma^\top \partial_y \psi\left(t, \int_{\mathbb{R}^d} z\nu(dz)\right) \right|^2 &= 0, \\ t \in [0, T], \nu \in \mathcal{P}_2(\mathbb{R}^d), \end{aligned}$$

$$\psi\left(T, \int_{\mathbb{R}^d} z\nu(dz)\right) = g\left(\int_{\mathbb{R}^d} z\nu(dz)\right), \quad \nu \in \mathcal{P}_2(\mathbb{R}^d),$$

which we see holds for any ν by substituting $y = \int_{\mathbb{R}^d} z\nu(dz)$ into (47). Then the controls from (33) in Theorem 2.4 are given by

$$v_i^N(t, x_1, \dots, x_N) = -\sigma^\top \partial_y \psi\left(t, \frac{1}{N} \sum_{i=1}^N x_i\right), \forall i \in \{1, \dots, N\}$$

and the importance sampling estimator (14) is given by

$$\begin{aligned} \hat{\delta}_N &= \frac{1}{M} \sum_{j=1}^M \exp\left(-Ng\left(\frac{1}{N} \sum_{j=1}^N \hat{X}_t^{j,N,s,y}\right)\right) \exp\left(\sum_{i=1}^N \left\{ \int_s^T \sigma^\top \partial_y \psi\left(t, \frac{1}{N} \sum_{j=1}^N \hat{X}_t^{j,N,s,y}\right) \cdot d\hat{W}_t^i \right. \right. \\ &\quad \left. \left. - \frac{1}{2} \int_s^T \left| \sigma^\top \partial_y \psi\left(t, \frac{1}{N} \sum_{j=1}^N \hat{X}_t^{j,N,s,y}\right) \right|^2 dt \right\} \right) \\ &= \frac{1}{M} \sum_{j=1}^N \exp(-Ng(\hat{Y}_T^{N,s,y,j})) \exp\left(-\sqrt{N} \int_s^T \hat{v}(t, \hat{Y}_t^{N,s,y,j}) \cdot d\hat{\mathcal{W}}_t^N - \frac{N}{2} \int_s^T |\hat{v}(t, \hat{Y}_t^{N,s,y,j})|^2 dt\right), \end{aligned}$$

where in the first line above

$$\begin{aligned} d\hat{X}_t^{i,N,s,y} &= \left[b_0 + B\hat{X}_t^{i,N,s,y} + \frac{1}{N} \bar{B} \sum_{j=1}^N \hat{X}_t^{j,N,s,y} - \sigma \sigma^\top \partial_y \psi\left(t, \frac{1}{N} \sum_{j=1}^N \hat{X}_t^{j,N,s,y}\right) \right] dt + \sigma d\hat{W}_t^i, \\ \hat{X}_s^{i,N,s,y} &= y_i. \end{aligned}$$

Note, however, that not all examples in the linear-quadratic case (44) can be framed as a small-noise SDE problem. For instance, even when $d = 1$, $P_2 \neq 0$, one needs to consider both $Y^{N,s,y}$ and $\tilde{Y}_t^{N,s,y} = \frac{1}{N} \sum_{i=1}^N (X_t^{i,N,s,y})^2$. For the latter, the martingale term will be $\frac{2\sigma}{N} \sum_{i=1}^N \int_s^T X_t^{i,N,s,y} dW_t^i$, which cannot be written in terms of $\mathcal{W}^N$.

Our next example 3.3 shows that our importance sampling scheme can yield zero relative error even in the situation where the particles from (1) are non-interacting, and hence IID. Again, the fact that the relative error of the importance sampling scheme is zero can be viewed as the consequence of a more general principle discussed in Remark 3.4.

Example 3.3. Consider now the system:

$$X_t^{i,N,y} = y - \int_0^t X_s^{i,N,y} + W_t^i$$

where $d = m = 1$, and we fix $s = 0$ and suppress it in the notation for simplicity. We are now dealing with IID diffusions, that is, there is no interaction between the particles.

Suppose we take $G(\mu) = \int_{\mathbb{R}} x^2 \mu(dx)$. Then in the notation of (44), $\bar{P}_2 = \bar{p}_1 = p_1 = \bar{B} = b_0 = p_2 = 0$, $\sigma = P_2 = 1$, $B = -1$ and $y_i = y$ for all $i \in \mathbb{N}$. Using the known density of the Ornstein–Uhlenbeck process $X_T^{1,N,y}$, we get:

$$\begin{aligned} \mathbb{E} \left[\exp(-NG(\mu_T^{N,y})) \right] &= \mathbb{E} \left[\exp \left(- \sum_{i=1}^N (X_T^{i,N,y})^2 \right) \right] \\ &= \mathbb{E} \left[\exp \left(-(X_T^{1,N,y})^2 \right) \right]^N \\ &= \exp \left(Ny^2 / (1 - 2e^{2T}) \right) (2 - e^{-2T})^{-N/2} \end{aligned}$$

and similarly

$$\mathbb{E} \left[\exp \left(-2NG(\mu_T^{N,y}) \right) \right] = \exp \left(2Ny^2 / (2 - 3e^{2T}) \right) (3 - 2e^{-2T})^{-N/2}.$$

Then:

$$\rho(\delta_N) = \frac{1}{\sqrt{M}} \sqrt{\exp \left(2Ny^2 [(2 - 3e^{2T})^{-1} - (1 - 2e^{2T})^{-1}] \right) \left(\frac{(2 - e^{-2T})^2}{3 - 2e^{-2T}} \right)^{N/2} - 1}.$$

Meanwhile, the optimal control can be found from solving the ODEs (45) to get $\Lambda(t) = \Gamma(t) = \frac{e^{2t}}{2e^{2T} - e^{2t}}$, $\gamma = 0$, $\chi(t) = \log(\sqrt{2 - e^{2t-2T}})$, so the control (46) is given by:

$$v_i^N(t, x_1, \dots, x_N) = \frac{2e^{2t}}{e^{2t} - 2e^{2T}} x_i.$$

Then the relative error $\rho(\hat{\delta}_N)$ is given by

$$\begin{aligned} & \frac{1}{\sqrt{M}} \sqrt{\frac{\mathbb{E} \left[\exp \left(-2 \sum_{i=1}^N (\hat{X}_T^{i,N,y})^2 \right) \prod_{i=1}^N \exp \left(-4 \int_0^T \frac{e^{2t}}{e^{2t} - 2e^{2T}} \hat{X}_t^{i,N,y} d\hat{W}_t^i - 4 \int_0^T \frac{e^{4t}}{(e^{2t} - 2e^{2T})^2} (\hat{X}_t^{i,N,y})^2 dt \right) \right]}{\mathbb{E} \left[\exp(-NG(\mu_T^{N,y})) \right]^2} - 1} \\ &= \frac{1}{\sqrt{M}} \sqrt{\frac{\mathbb{E} \left[\exp \left(-2 \sum_{i=1}^N (\hat{X}_T^{i,N,y})^2 + 2 \int_0^T \frac{e^{2t}}{e^{2t} - 2e^{2T}} \hat{X}_t^{i,N,y} d\hat{W}_t^i + 2 \int_0^T \frac{e^{4t}}{(e^{2t} - 2e^{2T})^2} (\hat{X}_t^{i,N,y})^2 dt \right) \right]}{\mathbb{E} \left[\exp(-NG(\mu^N)) \right]^2} - 1} \end{aligned}$$

where

$$\hat{X}_t^{i,N,y} = y + \int_0^t \left[\frac{2e^{2s}}{e^{2s} - 2e^{2T}} - 1 \right] \hat{X}_s^{i,N,y} ds + \hat{W}_t^i.$$

By Itô's formula, we have

$$\begin{aligned} & (\hat{X}_t^{i,N,y})^2 \frac{e^{2t}}{e^{2t} - 2e^{2T}} \\ &= y^2 \frac{1}{1 - 2e^{2T}} + 2 \int_0^t \frac{e^{4s}}{(e^{2s} - 2e^{2T})^2} (\hat{X}_s^{i,N,y})^2 ds + 2 \int_0^t \hat{X}_s^{i,N,y} \frac{e^{2s}}{e^{2s} - 2e^{2T}} d\hat{W}_s^i + \frac{1}{2} \log \left(\frac{2e^{2T} - e^{2t}}{2e^{2T} - 1} \right). \end{aligned}$$

Setting $t = T$, we get

$$(\hat{X}_T^{i,N,y})^2 + 2 \int_0^T \frac{e^{4s}}{(e^{2s} - 2e^{2T})^2} (\hat{X}_s^{i,N,y})^2 ds + 2 \int_0^T \hat{X}_s^{i,N,y} \frac{e^{2s}}{e^{2s} - 2e^{2T}} d\hat{W}_s^i = -y^2 \frac{1}{1 - 2e^{2T}} - \frac{1}{2} \log \left(\frac{e^{2T}}{2e^{2T} - 1} \right),$$

so, continuing:

$$\rho(\hat{\delta}_N) = \frac{1}{\sqrt{M}} \sqrt{\frac{\exp \left(\frac{2N}{1-2e^{2T}} y^2 \right) \frac{e^{2NT}}{(2e^{2T}-1)^N}}{\exp(2Ny^2/(1-2e^{2T}))} (2 - e^{-2T})^N - 1} = 0.$$

This shows that in this situation as well, while the standard Monte Carlo method yields a relative error which grows exponentially in N for fixed $y \in \mathbb{R}$, $T > 0$, and grows exponentially in $|y|$ for fixed $N \in \mathbb{N}$, $T > 0$, our importance sampling scheme has zero error.

In this trivial case, the controlled particles are IID, so for both the importance sampling scheme and standard Monte Carlo we only need to simulate one particle (making the analysis in N a bit useless in practice, but the analysis in $|y|$ still holds true when $N = 1$).

Remark 3.4. This last example can be viewed through the following lens. Suppose that $b(x, \mu) = b(x)$, $\sigma(x, \mu) = \sigma(x)$ (so that the particles are IID when given the same initial condition) and that $G(\mu) = \langle g, \mu \rangle$ for some $g : \mathbb{R}^d \rightarrow \mathbb{R}$. Making the ansatz $\Psi(t, \mu) = \langle \psi(t, \cdot), \mu \rangle$ in (23), we get:

$$\begin{aligned} \langle -\dot{\psi}(t, \cdot) - b(\cdot) \cdot \partial_x \psi(t, \cdot) + \frac{1}{2} |\sigma^\top(\cdot) \partial_x \psi(t, \cdot)|^2 - \frac{1}{2} \sigma \sigma^\top(\cdot) : \partial_x^2 \psi(t, \cdot), \mu \rangle &= 0, \quad t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d), \\ \langle \psi(T, \cdot), \mu \rangle &= \langle g(\cdot), \mu \rangle, \quad \mu \in \mathcal{P}_2(\mathbb{R}^d), \end{aligned}$$

which of course is satisfied if we have a unique solution ψ to

$$\begin{aligned} (49) \quad -\dot{\psi}(t, x) - b(x) \cdot \partial_x \psi(t, x) + \frac{1}{2} |\sigma^\top(x) \partial_x \psi(t, x)|^2 - \frac{1}{2} \sigma \sigma^\top(x) : \partial_x^2 \psi(t, x) &= 0, \quad t \in [0, T], x \in \mathbb{R}^d, \\ \psi(T, x) &= g(x), \quad x \in \mathbb{R}^d. \end{aligned}$$

For a related discussion, see Remark 3.3 in [73]. Equation (49) is the HJB equation corresponding to the zero-variance estimator resulting from the Doob h -transform in the single particle setting — see Equation (2.7) in [79]. It can thus be seen, as in, e.g., [79] pp. 1778-1779, that letting $v(t, x) = -\sigma^\top(x) \partial_x \psi(t, x)$,

$$\mathbb{E} \left[\exp \left(-2g(\hat{X}_T^{s,y}) \right) \exp \left(-2 \int_s^T v(t, \hat{X}_t^{s,y}) \cdot d\hat{W}_t - \int_0^T |v(t, \hat{X}_t^{s,y})|^2 dt \right) \right] = \mathbb{E}[\exp(-g(X_T^{s,y}))]^2$$

where

$$dX_t^{s,y} = b(X_t^{s,y}) dt + \sigma(X_t^{s,y}) dW_t, X_s^{s,y} = y$$

and

$$d\hat{X}_t^{s,y} = [b(\hat{X}_t^{s,y}) + \sigma v(t, \hat{X}_t^{s,y})] dt + \sigma(\hat{X}_t^{s,y}) d\hat{W}_t, \hat{X}_s^{s,y} = y.$$

Here $y \in \mathbb{R}^d$ and $W, \hat{W}$ are standard m -dimensional Brownian motions initialized at $W_s = \hat{W}_s = 0$. Then we have, for $\mu_T^{N,s,y}, \hat{\mu}_T^{N,s,y}$ the empirical measure on IID copies of solutions $X_T^{s,y}, \hat{X}_T^{s,y}$ respectively (which we denote by $X_T^{i,s,y}, \hat{X}_T^{i,s,y}, i \in \mathbb{N}$ and by $W^i, \hat{W}^i$ their driving Brownian motions) and using the controls from (33)

in Theorem 2.4 are given by $v_i^N(t, x_1, x_2, \dots, x_N) = -\sigma^\top(x_i)\partial_\mu\Psi(t, \mu_x^N)[x_i] = -\sigma^\top(x_i)\partial_x\psi(t, x_i) = v(t, x_i)$, the importance sampling relative error from (17) satisfies:

$$\begin{aligned}
& \frac{\mathbb{E}\left[\exp\left(-2NG(\hat{\mu}_T^{N,s,y})\right)\prod_{i=1}^N(Z^{i,N,s,y})^2\right]}{\mathbb{E}\left[\exp\left(-NG(\mu_T^{N,s,y})\right)\right]^2} - 1 \\
&= \frac{\mathbb{E}\left[\exp\left(\sum_{i=1}^N -2g(\hat{X}_T^{i,s,y}) - 2\int_s^T v(t, \hat{X}_t^{i,s,y})d\hat{W}_t^i - \int_s^T |v(t, \hat{X}_t^{i,s,y})|^2 dt\right)\right]}{\mathbb{E}\left[\exp\left(-\sum_{i=1}^N g(X_T^{i,s,y})\right)\right]^2} - 1 \\
&= \left(\mathbb{E}\left[\exp\left(-2g(\hat{X}_T^{s,y})\right)\exp\left(-2\int_s^T v(t, \hat{X}_t^{s,y})\cdot d\hat{W}_t - \int_s^T |v(t, \hat{X}_t^{s,y})|^2 dt\right)\right] \cdot \mathbb{E}[\exp(-g(X_T^{s,y}))]^{-2}\right)^N - 1 \\
&= 1 - 1 \\
&= 0.
\end{aligned}$$

Thus our importance sampling scheme yields an estimator with zero relative error in the setting of non-interacting diffusions and linear G .

This also shows that in the non-interacting regime, even when G is linear, solving (23) is as difficult as solving for the exact (non-zero viscosity) solution to the HJB equation associated with the importance sampling scheme for one particle (49). We take this remark as motivation to construct importance sampling schemes for the empirical measure of weakly interacting diffusions based on the joint small-noise and large N limit large deviation principles derived in [72] and [14]. We expect that using such a scheme, in the non-interacting and linear G regime, the small noise importance sampling scheme of, e.g., [79] will be recovered, and solving the resulting first-order HJB equation on Wasserstein space will be equivalent to solving the standard zero-viscosity HJB equation found as Equation (2.10) therein.

Remark 3.5. Note, despite the framework of our linear-quadratic example and the discussion in Remarks 3.2 and 3.4, that in Theorem 2.4, we have the freedom to choose $G(\mu)$ to be arbitrarily non-linear so long as it is sufficiently smooth to have a classical subsolution in the sense of Definition (2.1). That is, it need not take the form $G(\mu) = \langle g, \mu \rangle$ or even $G(\mu) = g_1(\langle g_2, \mu \rangle)$. Also, as evidenced in the examples of Subsection 4.2, we expect that even if G is not smooth, so that only a (weak) subsolution to (23) can be recovered, our importance sampling scheme can be expected to still yield sub-exponentially growing (log-efficient – see Definition 1.4), or even vanishing relative error (Definition 1.5) — see Tables 3 and 2, respectively.

Moreover, our importance sampling scheme treats the empirical measure of IID diffusions and of weakly interacting particles uniformly. The former case can in some sense be seen as a “Sanov’s Theorem” type generalization of the “Cramér’s Theorem” type results found in [37], since using the ideas outlined in that paper would essentially correspond to taking independent particles and $G(\mu) = g(\int_{\mathbb{R}^d} z\mu(dz))$.

We end this section by observing that the above examples are not anomalous within the linear-quadratic setting. Namely, the exact desired expectation (3), and hence the relative error of the standard Monte Carlo estimator, can always easily be found by solving a system of ODEs. Moreover, from this we can make the observation that the relative error of the importance sampling estimator with control (46) will in fact always have zero relative error.

Note that in the linear-quadratic regime, the right hand side of prelimit representation from (12) in Proposition 1.3 with $\tilde{b} = b$ and $F(\mu) = G(\mu_T)$ takes the form of a standard linear-quadratic stochastic control problem. We see then that

$$-\frac{1}{N}\log\mathbb{E}[\exp(-NG(\mu_T^{N,s,y}))] = \bar{y}^\top\Lambda_N(s)\bar{y} + \gamma_N(s) \cdot \bar{y} + \chi_N(s),$$

so the desired expectation from (3) is given by

$$\mathbb{E} \left[\exp \left(-NG(\mu_T^{N,s,y}) \right) \right] = \exp \left(-N[\bar{y}^\top \Lambda_N(s) \bar{y} + \gamma_N(s) \cdot \bar{y} + \chi_N(s)] \right).$$

In the above we denote by $(y_1, \dots, y_N) = \bar{y} \in \mathbb{R}^{dN}$, and $\Lambda_N \in \mathbb{R}^{dN \times dN}$, $\gamma_N \in \mathbb{R}^{dN}$, $\chi_N \in \mathbb{R}$ satisfy the Riccati equations:

$$\begin{aligned} \dot{\Lambda}_N(t) + \Lambda_N(t)B_N + B_N^\top \Lambda_N(t) - 2N\Lambda_N(t)\sigma_N\sigma_N^\top \Lambda_N(t) &= 0, \Lambda_N(T) = P_{N2}, \\ \dot{\gamma}_N(t) + B_N^\top \gamma_N(t) - 2N\Lambda_N(t)\sigma_N\sigma_N^\top \gamma_N(t) + 2\Lambda_N(t)b_{0N} &= 0, \gamma_N(T) = p_{N1}, \\ \dot{\chi}_N(t) - \frac{N}{2}\gamma_N^\top(t)\sigma_N\sigma_N^\top \gamma_N(t) + \gamma_N(t) \cdot b_{0N} + \sum_{i=1}^N \bar{\sigma}_{i,N}^\top \Lambda_N(t) \bar{\sigma}_{i,N} &= 0, \chi_N(T) = p_2. \end{aligned}$$

Here

$$\begin{aligned} b_{0N} &= [b_0^\top, \dots, b_0^\top]^\top \in \mathbb{R}^{dN}, \\ \sigma_N &= \text{diag}(\sigma) \in \mathbb{R}^{dN \times N}, \\ p_{1N} &= \frac{1}{N}[p_1^\top, \dots, p_1^\top]^\top \in \mathbb{R}^{dN}, \\ B_N &= \text{diag}(B) + \frac{1}{N}\bar{B}_N \in \mathbb{R}^{dN \times dN}, \\ P_{2N} &= \frac{1}{N}\text{diag}(P_2) + \frac{1}{N^2}\bar{P}_{2N} \in \mathbb{R}^{dN \times dN} \end{aligned}$$

and $\bar{B}_N$ is the block matrix with every $d \times d$ entry given by $\bar{B}$, $\bar{P}_{2N}$ is the block matrix with every $d \times d$ entry given by $\bar{P}_2$, $\bar{\sigma}_{i,N} \in \mathbb{R}^{dN}$ has i 'th entry σ and the rest 0.

It may seem initially that we need to solve $dN \times dN$ Riccati equations in order to get this value, but in fact we can find that, letting $\Lambda_N^{i,j} \in \mathbb{R}^{d \times d}$ denote the $d \times d$ block matrix in the i, j position, that $\Lambda_N^{i,j} = \bar{\Lambda}_N$ for all $i \neq j$ and $\Lambda_N^{i,i} = \hat{\Lambda}_N$ for all $i \in \{1, \dots, N\}$, and hence $\gamma_N^i = \bar{\gamma}_N \in \mathbb{R}^d$ for all i . Thus we need only solve:

$$\begin{aligned} \dot{\hat{\Lambda}}_N + \bar{\Lambda}_N[\bar{B} + B] + \frac{1}{N}[\hat{\Lambda}_N - \bar{\Lambda}_N]\bar{B} + [B + \bar{B}]^\top \bar{\Lambda}_N + \frac{1}{N}\bar{B}^\top [\hat{\Lambda}_N - \bar{\Lambda}_N] - 2N[(N-2)\bar{\Lambda}_N\sigma\sigma^\top \bar{\Lambda}_N^\top, \\ + \hat{\Lambda}_N\sigma\sigma^\top \bar{\Lambda}_N^\top + \bar{\Lambda}_N\sigma\sigma^\top \hat{\Lambda}_N^\top] &= 0, \\ \bar{\Lambda}_N(T) &= \frac{1}{N^2}\bar{P}_2, \\ \dot{\bar{\Lambda}}_N + (1 - \frac{1}{N})\bar{\Lambda}_N\bar{B} + (1 - \frac{1}{N})\bar{B}^\top \bar{\Lambda}_N + \hat{\Lambda}_N[\bar{B}/N + B] + [\bar{B}/N + B]^\top \hat{\Lambda}_N - 2N[(N-1)\bar{\Lambda}_N\sigma\sigma^\top \bar{\Lambda}_N^\top, \\ + \hat{\Lambda}_N\sigma\sigma^\top \hat{\Lambda}_N^\top] &= 0, \\ \hat{\Lambda}_N(T) &= \frac{1}{N^2}\bar{P}_2 + \frac{1}{N}P_2, \\ \dot{\bar{\gamma}}_N + [\bar{B} + B]^\top \bar{\gamma}_N - 2N[(N-1)\bar{\Lambda}_N + \hat{\Lambda}_N]\sigma\sigma^\top \bar{\gamma}_N + 2[(N-1)\bar{\Lambda}_N + \hat{\Lambda}_N]b_0 &= 0, \\ \bar{\gamma}_N(T) &= p_1/N, \\ \dot{\chi}_N(t) - \frac{N^2}{2}\bar{\gamma}_N^\top\sigma\sigma^\top \bar{\gamma}_N + N\bar{\gamma}_N \cdot b_0 + N\sigma^\top \hat{\Lambda}_N\sigma &= 0, \\ \chi_N(T) &= p_2. \end{aligned}$$

Then we have

$$\begin{aligned}
\mathbb{E} \left[\exp \left(-NG(\mu_T^{N,s,y}) \right) \right] &= \exp \left(-N \left[\bar{y}^\top \Lambda_N(s) \bar{y} + \gamma_N(s) \cdot \bar{y} + \chi_N(s) \right] \right) \\
&= \exp \left(-N \left[\sum_{i \neq j}^N y_i^\top \bar{\Lambda}_N(s) y_j + \sum_{i=1}^N y_i^\top \hat{\Lambda}_N(s) y_i + \bar{\gamma}^\top(s) \sum_{i=1}^N y_i + \chi_N(s) \right] \right) \\
&= \exp \left(-N \left[\left(\sum_{i=1}^N y_i \right)^\top \bar{\Lambda}_N(s) \left(\sum_{i=1}^N y_i \right) + \sum_{i=1}^N y_i^\top \left[\hat{\Lambda}_N(s) - \bar{\Lambda}_N(s) \right] y_i \right. \right. \\
&\quad \left. \left. + \bar{\gamma}^\top(s) \sum_{i=1}^N y_i + \chi_N(s) \right] \right)
\end{aligned}$$

Note that this is useful not only for establishing the true desired expectation (3), but also allows for us to calculate the relative error for the standard Monte Carlo estimator $\rho(\delta_N)$ by substituting $2G$ for G in (44) with P_2, p_1 , and $\bar{P}_2$ modified to $2P_2, 2p_1$, and $2\bar{P}_2$, respectively. We will use this fact in the numerical example in Subsection 4.1.

Also note that we can see directly how to write the right-hand side of the previous display as a function of μ_y^N . Letting

$$\begin{aligned}
(50) \quad \Phi^N(t, \nu) &= N^2 \left[\int_{\mathbb{R}^d} z \mu(dz) \right]^\top \bar{\Lambda}_N(t) \left[\int_{\mathbb{R}^d} z \mu(dz) \right] + N \int_{\mathbb{R}^d} z^\top [\hat{\Lambda}_N(t) - \bar{\Lambda}_N(t)] z \mu(dz) \\
&\quad + N \bar{\gamma}_N^\top \int_{\mathbb{R}^d} z \mu(dz) + \tilde{\chi}_N(t) \\
&= \int_{\mathbb{R}^d} z^\top \tilde{\Lambda}_N(t) z \mu(dz) + \left(\int_{\mathbb{R}^d} z \mu(dz) \right)^\top [\tilde{\Gamma}_N(t) - \tilde{\Lambda}_N(t)] \left(\int_{\mathbb{R}^d} z \mu(dz) \right) + \tilde{\gamma}_N^\top \int_{\mathbb{R}^d} z \mu(dz) + \tilde{\chi}_N(t),
\end{aligned}$$

where $\tilde{\Lambda}_N(t) = N [\hat{\Lambda}_N(t) - \bar{\Lambda}_N(t)]$, $\tilde{\Gamma}_N = N(N-1)\bar{\Lambda}_N(t) + N\hat{\Lambda}(t)$, $\tilde{\gamma}_N(t) = N\gamma(t)$, we have

$$\exp \left(-N\Phi^N(s, \mu_x^N) \right) = \mathbb{E} \left[\exp \left(-NG(\mu_T^{N,s,x}) \right) \right],$$

and get the Riccati equations:

$$\begin{aligned}
(51) \quad \dot{\tilde{\Lambda}}_N + \tilde{\Lambda}_N B + B^\top \tilde{\Lambda}_N - 2\tilde{\Lambda}_N \sigma \sigma^\top \tilde{\Lambda}_N^\top &= 0, \quad \tilde{\Lambda}_N(T) = P_2, \\
\dot{\tilde{\Gamma}}_N + \tilde{\Gamma}_N [B + \bar{B}] + [B + \bar{B}]^\top \tilde{\Gamma}_N - 2\tilde{\Gamma}_N \sigma \sigma^\top \tilde{\Gamma}_N^\top &= 0, \quad \tilde{\Gamma}_N(T) = \bar{P}_2 + P_2, \\
\dot{\tilde{\gamma}}_N(t) + [B + \bar{B}]^\top \tilde{\gamma}_N(t) - 2\tilde{\Gamma}_N(t) \sigma \sigma^\top \tilde{\gamma}_N(t) + 2\tilde{\Gamma}_N(t) b_0 &= 0, \quad \tilde{\gamma}_N(T) = p_1, \\
\dot{\tilde{\chi}}_N(t) - \frac{1}{2} \tilde{\gamma}_N^\top \sigma \sigma^\top \tilde{\gamma}_N + \tilde{\gamma}_N \cdot b_0 + \sigma^\top \tilde{\Lambda}_N \sigma + \frac{1}{N} \sigma^\top [\tilde{\Gamma}_N - \tilde{\Lambda}_N] \sigma &= 0, \quad \tilde{\chi}_N(T) = p_2.
\end{aligned}$$

Note firstly that it is easily verified that $\Phi^N(t, \nu)$ from (50) satisfies the prelimit HJB equation (32) for all ν . Also observe that, other than the perturbation of the ODE for $\tilde{\chi}_N$ by the addition of the term $\frac{1}{N} \sigma^\top [\tilde{\Gamma}_N - \tilde{\Lambda}_N] \sigma$, these are the same Riccati equations as (45) for the “zero viscosity” HJB equation (23). The zero variance control from (30) is thus:

$$\begin{aligned}
v_i^N(t, x_1, \dots, x_N) &= -\sigma^\top \left[2\tilde{\Lambda}(t) x_i + 2[\tilde{\Gamma}(t) - \tilde{\Lambda}(t)] \int_{\mathbb{R}^d} z \mu_x^N(dz) + \tilde{\gamma}(t) \right] \\
&= -\sigma^\top \left[2\Lambda(t) x_i + 2[\Gamma(t) - \Lambda(t)] \int_{\mathbb{R}^d} z \mu_x^N(dz) + \gamma(t) \right],
\end{aligned}$$

that is, it is the same as the control from Theorem 2.4 in the linear-quadratic regime (see (46)).

Remark 3.6. The above discussion implies that in the linear-quadratic regime, the relative error of our importance sampling scheme is always zero.

This can also be derived via the following observation: In (32), when we use the ansatz that the solution takes the form of (50), since $\partial_\mu^2 \Phi^N(\mu)[z_1, z_2] = 2[\tilde{\Gamma}_N(t) - \tilde{\Lambda}_N(t)]$, the only term appearing from the second-order term $\int_{\mathbb{R}^d} \frac{1}{2N} \sigma \sigma^\top : \partial_\mu^2 \Phi^N(t, \mu)[z, z] \mu(dz)$ is $\frac{1}{N} \sigma \sigma^\top : [\tilde{\Gamma}_N(t) - \tilde{\Lambda}_N(t)] = \frac{1}{N} \sigma^\top [\tilde{\Gamma}_N(t) - \tilde{\Lambda}_N(t)] \sigma$. Thus, when deriving the form of the Riccati equations (51), it will only show up for $\tilde{\chi}_N$, which collects the constant in μ terms.

More generally, if we have a solution to (23) such that $\partial_\mu \Psi = \partial_\mu \Phi^N$ for all N , where Φ^N solves (32), then the importance sampling scheme has zero relative error for all N , as per the discussion in Section 2.1. Supposing, as in the linear-quadratic case, that we have a solution Ψ to (23) such that $\partial_\mu^2 \Psi(t, \mu)[z_1, z_2] = C(t) \in \mathbb{R}^{d \times d}$ and that σ does not depend on x, μ , then $\Psi(T, \nu) = G(\nu)$ and:

$$\begin{aligned} -\partial_t \Psi(t, \nu) - \int_{\mathbb{R}^d} b(z, \nu) \cdot \partial_\mu \Psi(t, \nu)[z] - \frac{1}{2} |\sigma^\top \partial_\mu \Psi(t, \nu)[z]|^2 \\ + \frac{1}{2} \sigma \sigma^\top(z, \nu) : \partial_z \partial_\mu \Psi(t, \nu)[z] + \frac{1}{2N} \sigma \sigma^\top : \partial_\mu^2 \Psi(t, \nu)[z, z] \nu(dz) = -\frac{1}{2N} \sigma \sigma^\top : C(t), \end{aligned}$$

so letting $\lambda \in C([0, T]; \mathbb{R})$ solve $\dot{\lambda}(t) = -\frac{1}{2N} \sigma \sigma^\top : C(t)$, $\lambda(T) = 0$, $\Phi^N(t, \mu) = \Psi(t, \mu) + \lambda(t)$ is a solution to (32), and $\partial_\mu \Phi^N = \partial_\mu \Psi$ for all N .

It thus clear that letting $m > 1$ or having the coefficients $b_0, B, \bar{B}, \sigma$ in (44) depend on time does not change the fact that our importance sampling scheme yields a zero-variance estimator in the linear-quadratic regime. Also note that this discussion also applies in the linear-quadratic setting for small noise importance sampling schemes — compare Equations (2.7) and (2.10) in [79] and recall that if b is linear in x , σ is constant, and G is quadratic in x therein that solutions to equation (2.10) will have second derivatives which are constant in x .

This reflects a difficulty present in both our setting and the small noise setting for importance sampling schemes based on large deviation principles and zero-viscosity HJB equations — the hurdle of actually computing the derivative of a solution to the HJB equation and hence the control to be used, numerically or otherwise, persists, and in the main class of problems for which a solution can be computed analytically (the linear-quadratic regime), the solution to the zero-viscosity HJB equation is no easier to solve for than the desired expectation. This means that efficient methods for numerically computing $\partial_\mu \Psi$ appearing in the definition of the control (33) in Theorem 2.4, or otherwise numerically obtaining the control, are highly desirable. Designing and implementing such schemes is an interesting avenue for future research.

4. NUMERICAL RESULTS

Here we present numerical results based on examples in the linear-quadratic regime of (3) and perturbations thereof. In all three examples, the importance sampling scheme greatly outperforms standard Monte Carlo methods. In the standard linear-quadratic regime, we observe that the importance sampling relative error is near 0 for all N , as expected, in Table 1. When a “global” lack of differentiability is introduced to G , we observe in Table 2 that the importance sampling relative error appears to be vanishing. When a “local” lack of differentiability is introduced to G , we observe in Table 3 that the importance sampling relative error does not seem to vanish, but grows sublinearly, still a massive improvement over the standard Monte Carlo relative error which increases exponentially with N in all three examples.

We simulate all the SDE systems using the Euler-Maruyama method with time step $\Delta t = .01/N$, which we chose empirically by iteratively refining until the estimates δ_N and $\hat{\delta}_N$ stabilized (i.e., consecutive estimates were nearly equal). We use $M \approx 10^7$ samples for all our simulations. In each case we also use the same final time $T = 1$ and diffusion constant $\sigma = .5$.

We note that our tabled results stop at relatively modest values of N ; although the importance sampling relative error remained small for much larger values of N than those presented in the tables, the exponential growth in the standard Monte Carlo relative errors required too many simulations to resolve for large N for us to make accurate comparisons.

A Python code implementing the examples in this section is provided on the public Gitlab repository <https://gitlab.com/mheldman/IS-interacting-particles>.

4.1. The Linear-Quadratic Regime. For our first example, we consider

$$(52) \quad G(\mu) = \int_{\mathbb{R}} z^2 \mu(dz).$$

This corresponds to the linear-quadratic regime in Section 3 with $d = m = 1, s = 0, B = -1, \bar{B} = 2, P_2 = 1, b_0 = \bar{P}_2 = p_1 = p_2 = 0$ and $y_i = y$ for all $i \in \mathbb{N}$.

We seek to estimate

$$\mathbb{E} \left[\exp \left(-NG(\mu_T^{N,y}) \right) \right] = \mathbb{E} \left[\exp \left(-\sum_{i=1}^N |X_T^{i,N,y}|^2 \right) \right]$$

where

$$(53) \quad dX_t^{i,N,y} = \left[-X_t^{i,N,y} + \frac{2}{N} \sum_{j=1}^N X_t^{j,N,y} \right] dt + \sigma dW_t^i, X_0^{i,N,y} = y.$$

Note that, as discussed in Remark 3.2, the importance sampling scheme for this case cannot be derived from previous work on small-noise SDEs since $P_2 \neq 0$.

The typical dynamics from (7) will have $\mathbb{E}[X_t^{\delta_y}] = ye^t$ for all $t \in [0, T]$ (we suppress the initial condition s in the notation here) and thus will be given by the time-inhomogeneous Ornstein–Uhlenbeck-like process with drift coefficient $b(t, x) = -x + 2ye^t$ and diffusion coefficient σ . The result of evaluating G from (52) at $\mu^{N,y}$:

$$(54) \quad G(\mu_T^{N,y}) = \frac{1}{N} \sum_{i=1}^N |X_T^{i,N,y}|^2$$

is thus expected to be close to the second moment of such a process at time T as N becomes large, which is

$$(55) \quad \mathbb{E} \left[|X_T^{\delta_y}|^2 \right] = y^2 e^{2T} + \frac{\sigma^2}{2} (1 - \exp(-2T)).$$

The atypical event we are seeking to sample, represented by G being close to zero, is that the particles interact in such a way that they are closer to the origin than expected via (55).

The role of the importance sampling scheme from Theorem 2.4 will thus be to modify the interaction and dynamics of each particle to predictably remain closer to the origin. Indeed, as discussed in Remark 3.6, the theoretical relative error of the importance sampling scheme is zero, so any contributing relative error will be due to numerical discretization. Indeed, we observed in numerical experiments that taking $\Delta t \downarrow 0$ also decreases the relative error of the importance scheme to 0.

Solving the Riccati equations (45), we get

$$\begin{aligned} \Lambda(t) &= \frac{e^{2t}}{e^{2T}(1 + \sigma^2) - e^{2t}\sigma^2}, \\ \Gamma(t) &= \frac{e^{2T}}{e^{2T}\sigma^2 - e^{2t}(\sigma^2 - 1)}, \\ \gamma(t) &= 0, \\ \chi(t) &= \frac{1}{2} \log \left(\frac{e^{2T}}{e^{2T}(1 + \sigma^2) - e^{2t}\sigma^2} \right) \end{aligned}$$

Thus the optimal control v from (46) and Theorem 2.4 is given by:

$$v_i^N(t, x_1, \dots, x_N) = -2\sigma \left[\frac{e^{2t}}{e^{2T}(1 + \sigma^2) - e^{2t}\sigma^2} x_i + \frac{1}{N} \left(\frac{e^{2T}}{e^{2T}\sigma^2 - e^{2t}(\sigma^2 - 1)} - \frac{e^{2t}}{e^{2T}(1 + \sigma^2) - e^{2t}\sigma^2} \right) \sum_{i=1}^N x_i \right].$$

In Table 1 below, we provide for various N the estimated expectation from standard Monte Carlo δ_N and importance sampling $\hat{\delta}_N$ (14), the empirical relative error from standard Monte Carlo $\tilde{\rho}(\delta_N)$ and for

importance sampling $\tilde{\rho}(\hat{\delta}_N)$, and the exact value of the expectation (3) from (50) and (51). Here $\tilde{\rho}(\hat{\delta}_N)$ is as in (16) but where $M = 1$ and the expectation and second moment are computed empirically. That is:

$$(56) \quad \tilde{\rho}(\hat{\delta}_N) := \sqrt{\frac{\frac{1}{M} \sum_{j=1}^M \exp(-2NG(\hat{\mu}_T^{N,s,y,j})) \left(\prod_{i=1}^N Z^{i,N,s,y,j} \right)^2}{\hat{\delta}_N^2}} - 1.$$

and similarly:

$$(57) \quad \tilde{\rho}(\delta_N) := \sqrt{\frac{\frac{1}{M} \sum_{j=1}^M \exp(-2NG(\mu_T^{N,s,y,j}))}{\delta_N^2}} - 1,$$

	IS Scheme		Standard Monte Carlo		
N	$\hat{\delta}_N$	$\tilde{\rho}(\hat{\delta}_N)$	δ_N	$\tilde{\rho}(\delta_N)$	Exact Value
5	$2.3816 \cdot 10^{-1}$	$3.4601 \cdot 10^{-2}$	$2.3807 \cdot 10^{-1}$	1.0380	$2.3747 \cdot 10^{-1}$
10	$8.2550 \cdot 10^{-2}$	$2.7101 \cdot 10^{-2}$	$8.2551 \cdot 10^{-2}$	1.5721	$8.2412 \cdot 10^{-2}$
15	$2.8641 \cdot 10^{-2}$	$2.4001 \cdot 10^{-2}$	$2.8661 \cdot 10^{-2}$	2.1957	$2.8600 \cdot 10^{-2}$
20	$9.9373 \cdot 10^{-3}$	$2.2369 \cdot 10^{-2}$	$9.9165 \cdot 10^{-3}$	2.9589	$9.9254 \cdot 10^{-3}$
25	$3.4486 \cdot 10^{-3}$	$2.1310 \cdot 10^{-2}$	$3.4824 \cdot 10^{-3}$	3.9086	$3.4445 \cdot 10^{-3}$
30	$1.1968 \cdot 10^{-3}$	$2.0550 \cdot 10^{-2}$	$1.1951 \cdot 10^{-1}$	5.1415	$1.1954 \cdot 10^{-3}$
50	$1.7361 \cdot 10^{-5}$	$1.8934 \cdot 10^{-2}$	$1.7569 \cdot 10^{-5}$	14.5202	$1.7339 \cdot 10^{-5}$
80	$3.0339 \cdot 10^{-8}$	$1.8003 \cdot 10^{-2}$	$3.2426 \cdot 10^{-8}$	74.1871	$3.0289 \cdot 10^{-8}$

TABLE 1. Comparison of the standard Monte Carlo estimator δ_N with the importance sampling estimator $\hat{\delta}_N$ (14) and their corresponding relative errors $\tilde{\rho}(\delta_N)$ (57) and $\tilde{\rho}(\hat{\delta}_N)$ (56) with G as in (52) and particles obeying (53) with initial condition $y = .2$.

In Figure 1(a) we plot an average of trajectories of the sum (54) for the uncontrolled and controlled particles, respectively, computed using the same noise. In Figure 1(b), we plot the analytical relative error for the Monte Carlo scheme ($\rho(\delta_N)$ from (16) with $v_i^N \equiv 0$) on a log scale for various values of N , again as computed via (50) and (51). One can see that the first few values of N considered in Table 1 are in a region where the growth for the standard Monte Carlo relative error has not yet shifted from linear to exponential, where importance sampling becomes even more valuable.

4.2. Examples with Nonsmooth G . Here we provide two examples where G is not differentiable in the Lions sense (see Definition A.1). The first example is constructed via a perturbation of the linear-quadratic regime which reduces the smoothness of G in a “global” way, whereas the second reduces the smoothness of G in a “local” way. In both examples, despite the lack of regularity we observe that the relative error of the importance sampling scheme is much smaller than that of the standard Monte Carlo method for all N . In Example 4.2, the relative error grows sublinearly, and in Example 4.1, the importance sampling relative error even appears to be vanishing, exhibiting that the assumptions in Theorems 2.4 and 2.5 are sufficient but not necessary. Thanks to the simple nature of the target function G , we are able to provide an interpretation of the role of the controls in each example — see Remark 4.3.

Example 4.1. In our second example, we consider

$$(58) \quad G(\mu) = \left| \int_{\mathbb{R}} z \mu(dz) \right|,$$

with drift and diffusion conforming to the linear-quadratic regime (44) with $d = m = 1$.

Note that while G is continuous, it is not Lions differentiable at measures μ such that $\int_{\mathbb{R}} z \mu(dz) = 0$, so we are not able to obtain an explicit solution to the HJB equation (23) as in Section 3. However, for $\mu \in \mathcal{P}_2(\mathbb{R})$ with nonzero mean, it is plain to see that $\partial_{\mu} G(\mu)[z] = \text{sign}\left(\int_{\mathbb{R}} x \mu(dx)\right)$ for all $z \in \mathbb{R}$, so $\partial_z \partial_{\mu} G(\mu) \equiv 0$ and




FIGURE 1. Left: An average of 5 realizations of the sum (54) for the uncontrolled (red) and controlled (green) particles respectively against the variance of the limiting McKean-Vlasov equation from (55) (black) for $T \in [0, 1.5]$, $y = .2$, and $N = 50$. Right: the log-scaled exact relative error $\rho(\delta_N)$ as N varies in the setting of Table 1, with $M = 1$ fixed in (16) (black) against the exact asymptotic rate $\approx e^{.051N}$ (red).

$\partial_\mu^2 G(\mu) \equiv 0$. Using this ansatz, taking $b_0 = 0$, and performing a similar computation to [73] Section 4, we obtain the formal (weak) viscosity solution to (23) in this setting:

$$\Psi(t, \mu) = e^{(B+\bar{B})(T-t)} \left| \int_{\mathbb{R}} z \mu(dz) \right| + \frac{\sigma^2}{4(B+\bar{B})} [1 - e^{2(B+\bar{B})(T-t)}]$$

so that we expect the control from (33) in Theorem 2.4 to be of the form

$$\begin{aligned} v(t, x_1, \dots, x_N) &= -\sigma \partial_\mu \Psi(t, \mu_x^N)[x_i] \\ &= -\sigma e^{(B+\bar{B})(T-t)} \text{sign} \left(\frac{1}{N} \sum_{j=1}^N x_j \right). \end{aligned}$$

Indeed, via the equivalence with the small-noise regime as discussed in Remark 3.2, such a choice of control can be seen to solve the minimization problem (9) with $F(\mu) = G(\mu_T)$.

In Table 2, we provide for various N the estimated expectation and empirical relative error for our importance sampling scheme compared with standard Monte Carlo, selecting $B = -1$, $\bar{B} = 2$, $s = 0$, and $y_i = .4$ for all $i \in \mathbb{N}$. As we see, the relative error for the importance sampling scheme seems to be vanishing as N increases, whereas the standard Monte Carlo relative error increases exponentially.

	IS Scheme		Standard Monte Carlo	
N	$\hat{\delta}_N$	$\tilde{\rho}(\hat{\delta}_N)$	δ_N	$\tilde{\rho}(\delta_N)$
5	$2.6644 \cdot 10^{-2}$	$3.7070 \cdot 10^{-1}$	$2.6653 \cdot 10^{-2}$	2.9930
10	$9.0804 \cdot 10^{-4}$	$3.1427 \cdot 10^{-1}$	$9.2009 \cdot 10^{-4}$	12.8032
15	$3.0269 \cdot 10^{-5}$	$2.6350 \cdot 10^{-1}$	$3.0645 \cdot 10^{-5}$	53.9410
20	$9.9533 \cdot 10^{-7}$	$2.2227 \cdot 10^{-1}$	$1.0625 \cdot 10^{-6}$	337.0401
25	$3.2471 \cdot 10^{-8}$	$1.8900 \cdot 10^{-1}$	$3.5791 \cdot 10^{-8}$	561.9164
30	$1.0537 \cdot 10^{-9}$	$1.6154 \cdot 10^{-1}$	$1.2253 \cdot 10^{-9}$	1016.1656

TABLE 2. Comparison of the standard Monte Carlo estimator δ_N with the importance sampling estimator $\hat{\delta}_N$ (14) and their corresponding relative errors $\tilde{\rho}(\delta_N)$ (57) and $\tilde{\rho}(\hat{\delta}_N)$ (56) with G as in (58) and particles obeying (53) with initial condition $y = .4$ (Example 4.1).

Example 4.2.

	IS Scheme		Standard Monte Carlo	
N	$\hat{\delta}_N$	$\tilde{\rho}(\hat{\delta}_N)$	δ_N	$\tilde{\rho}(\delta_N)$
5	$2.1355 \cdot 10^{-2}$	$4.7870 \cdot 10^{-1}$	$2.1382 \cdot 10^{-2}$	2.3263
10	$5.6703 \cdot 10^{-4}$	$6.0605 \cdot 10^{-1}$	$5.6806 \cdot 10^{-4}$	6.6374
15	$1.5087 \cdot 10^{-5}$	$7.2452 \cdot 10^{-1}$	$1.5040 \cdot 10^{-5}$	17.3337
20	$4.0157 \cdot 10^{-7}$	$8.3416 \cdot 10^{-1}$	$4.1260 \cdot 10^{-7}$	51.1058
25	$1.0687 \cdot 10^{-8}$	$9.4327 \cdot 10^{-1}$	$1.0296 \cdot 10^{-8}$	97.8820
30	$2.8470 \cdot 10^{-10}$	1.0524	$2.8800 \cdot 10^{-10}$	280.5945

TABLE 3. Comparison of the standard Monte Carlo estimator δ_N with the importance sampling estimator $\hat{\delta}_N$ (14) and their corresponding relative errors $\tilde{\rho}(\delta_N)$ (57) and $\tilde{\rho}(\hat{\delta}_N)$ (56) with G as in (59) and particles obeying (53) with initial condition $y = .4$ (Example 4.2).

Finally, we consider the same regime as Example 4.1, but now taking G in the target expectation (3) to be:

$$(59) \quad G(\mu) = \int_{\mathbb{R}} |z| \mu(dz).$$

Once again G is not everywhere Lions differentiable, in this case at measures such that the singleton $\{0\} \in \mathbb{R}$ has positive measure. Indeed, for $\mu \in \mathcal{P}_2(\mathbb{R})$ such that $\mu(\{0\}) = 0$, it is clear to see that $\partial_\mu G(\mu)[z] = \text{sign}(z)$, so $\partial_\mu^2 G(\mu) = 0$ and for $z \neq 0$, $\partial_z \partial_\mu G(\mu)[z] = 0$ (this follows via a similar computation to [19] Example 1 in Section 5.2.2).

Again we follow the methods of [73], taking $b_0 = 0$, and obtain the formal (weak) viscosity solution to (23):

$$\begin{aligned} \Psi(t, \mu) &= e^{B(T-t)} \int_{\mathbb{R}} |z| \mu(dz) + [e^{(B+\bar{B})(T-t)} - e^{B(T-t)}] \int_{\mathbb{R}} z \mu(dz) \int_{\mathbb{R}} \text{sign}(z) \mu(dz) \\ &\quad + \frac{\sigma^2}{4B(B+\bar{B})} [e^{2B(T-t)}(B+\bar{B}) - B e^{2(B+\bar{B})(T-t)} - \bar{B}] \left[\int_{\mathbb{R}} \text{sign}(z) \mu(dz) \right]^2 + \frac{\sigma^2}{4B} [1 - e^{2B(T-t)}]. \end{aligned}$$

Using that for $\mu \in \mathcal{P}_2(\mathbb{R})$ with $\mu(\{0\}) = 0$, $\mu \mapsto \int_{\mathbb{R}} \text{sign}(z) \mu(dz)$ is Lions differentiable with derivative 0, we expect the control from (33) in Theorem 2.4 to be of the form

$$\begin{aligned} v(t, x_1, \dots, x_N) &= -\sigma \partial_\mu \Psi(t, \mu_x^N)[x_i] \\ &= -\sigma \left[e^{B(T-t)} \text{sign}(x_i) + \left(e^{(B+\bar{B})(T-t)} - e^{B(T-t)} \right) \frac{1}{N} \sum_{j=1}^N \text{sign}(x_j) \right]. \end{aligned}$$

In Table 3, we provide for various N the estimated expectation and empirical relative error for our importance sampling scheme compared with standard Monte Carlo, selecting $B = -1$, $\bar{B} = 2$, $s = 0$, and $y_i = .4$ for all $i \in \mathbb{N}$. As we see, the importance sampling scheme has relative error which increases sublinearly as N becomes large, in contrast to the standard Monte Carlo relative error which increases exponentially.

Remark 4.3. We can think about the controls from Examples 4.1 and 4.2 in terms of the role they play in sampling the “rare event” of G being close to 0 for the particle system in each case. When the absolute value appears outside the empirical mean, as in Example 4.1, we only need to be concerned with controlling the magnitude of the empirical mean, a characteristic of the particle ensemble. This is reflected in the controls, which are identical for every particle and force them in the opposite direction of the sign of the empirical mean. Meanwhile, in the situation of Example (4.2), to make G close to zero we need the position of each individual particle to be near 0. Thus, the control uses information about the position of each particle relative to the origin. However, information about the ensemble of particles also still must appear because of the global interaction in the particles’ dynamics, which we see in the form of the difference in the proportions of the particles which are positive and negative at a given time.

We note that a similar combination of global information on the ensemble of particles and local information about the i 'th particle is also used for controls in the standard linear-quadratic regime of Section 3. See in particular the control used in our simple numerical example in Subsection 4.1. Numerical experiments confirm that the role of both the global and local information captured in the control is necessary for variance reduction in such cases. For example, using the simplified control $\tilde{v}(t, x_1, \dots, x_N) = -\sigma e^{B(T-t)} \text{sign}(x_i)$ in Example 4.2, we observed performance that was comparable or even worse than standard Monte Carlo.

It is also worth commenting on the choice of initial conditions in the examples of Subsection 4.2. We choose $y_i = .4$ for all $i \in \mathbb{N}$, so that $\nu = \delta_{.4}$ is away from the region of discontinuity of $\partial_\mu \Psi(t, \cdot)[z]$. We observed that if we took initial conditions such that ν is closer to the region of discontinuity many more rogue trajectories appeared in the controlled dynamics and an extremely high sample size was needed before the importance sampling scheme could be seen to perform better than standard Monte Carlo. This is once again similar to the case of small-noise diffusions—see Table 3.1 compared to Table 3.2 in [79] and the discussion on pp. 1786-1787 therein.

We end this section by remarking on the interesting problem of verifying that the formal viscosity solutions to the HJB equation (23) constructed in Examples 4.1 and 4.2 are indeed viscosity solutions in the sense of [26]. To our knowledge, there are no existing examples in the current literature of explicit viscosity solutions which are not also classical solutions to an HJB equation on Wasserstein space.

5. PROOF OF THE MAIN RESULTS

We begin with a lemma which allows us to express the numerator of $R(\hat{\delta}_N)$ (see (17)) in an alternative way:

Lemma 5.1. *Let Assumption 1.1 (A3) and Assumption 2.3 (B1) hold. Then, for any $s \in [0, T]$, $y_1, \dots, y_N \in \mathbb{R}^d$, and for $\hat{\mu}$ as in (15), $Z^{i,N,s,y}$ as in (13), and $v_i^N : [s, T] \times \mathbb{R}^{dN} \rightarrow \mathbb{R}^m$, $i = 1, \dots, N$ bounded:*

$$\mathbb{E} \left[\exp \left(-2NG(\hat{\mu}_T^{N,s,y}) \right) \prod_{i=1}^N (Z^{i,N,s,y})^2 \right] = \mathbb{E} \left[\exp \left(-2NG(\tilde{\mu}_T^{N,s,y}) \right) \exp \left(\int_s^T \sum_{i=1}^N |v_i^N(t, \tilde{X}_t^{1,N,s,y}, \dots, \tilde{X}_t^{N,N,s,y})|^2 dt \right) \right],$$

where $\tilde{\mu}_t^{N,s,y} = \frac{1}{N} \sum_{i=1}^N \delta_{\tilde{X}_t^{i,N,s,y}}$ and $\tilde{X}^{i,N,s,y}$ satisfy:

$$d\tilde{X}_t^{i,N,s,y} = [b(\tilde{X}_t^{i,N,s,y}, \tilde{\mu}_t^{N,s,y}) - \sigma(\tilde{X}_t^{i,N,s,y}, \tilde{\mu}_t^{N,s,y}) v_i^N(t, \tilde{X}_t^{1,N,s,y}, \dots, \tilde{X}_t^{N,N,s,y})] dt + \sigma(\tilde{X}_t^{i,N,s,y}, \tilde{\mu}_t^{N,s,y}) d\tilde{W}_t^i, \\ \tilde{X}_s^{i,N,s,y} = y_i,$$

and $\tilde{W}^i$ are IID standard m -dimensional Brownian motions initialized at $\tilde{W}_s^i = 0$.

Proof. This follows by a simple application of the Girsanov theorem, which is valid due to the boundedness of the v_i^N 's and G . We have:

$$\begin{aligned} & \mathbb{E} \left[\exp \left(-2NG(\hat{\mu}_T^{N,s,y}) \right) \prod_{i=1}^N (Z^{i,N,s,y})^2 \right] \\ &= \mathbb{E} \left[\exp \left(-2NG(\hat{\mu}_T^{N,s,y}) \right) \exp \left(-2 \sum_{i=1}^N \int_s^T v_i^N(t, \hat{X}_t^{1,N,s,y}, \dots, \hat{X}_t^{N,N,s,y}) \cdot d\hat{W}_t^i \right. \right. \\ & \quad \left. \left. - \int_s^T \sum_{i=1}^N |v_i^N(t, \hat{X}_t^{1,N,s,y}, \dots, \hat{X}_t^{N,N,s,y})|^2 dt \right) \right] \\ &= \mathbb{E} \left[\exp \left(-2NG(\hat{\mu}_T^{N,s,y}) \right) \exp \left(-2 \sum_{i=1}^N \int_s^T v_i^N(t, \hat{X}_t^{1,N,s,y}, \dots, \hat{X}_t^{N,N,s,y}) \cdot d\hat{W}_t^i \right. \right. \\ & \quad \left. \left. - \frac{1}{2} \int_s^T \sum_{i=1}^N |2v_i^N(t, \hat{X}_t^{1,N,s,y}, \dots, \hat{X}_t^{N,N,s,y})|^2 dt \right) \exp \left(\int_s^T \sum_{i=1}^N |v_i^N(t, \hat{X}_t^{1,N,s,y}, \dots, \hat{X}_t^{N,N,s,y})|^2 dt \right) \right] \\ &= \mathbb{E} \left[\exp \left(-2NG(\tilde{\mu}_T^{N,s,y}) \right) \exp \left(\int_s^T \sum_{i=1}^N |v_i^N(t, \tilde{X}_t^{1,N,s,y}, \dots, \tilde{X}_t^{N,N,s,y})|^2 dt \right) \right]. \end{aligned}$$

□

5.1. Proof of Theorem 2.4. Let Ψ be a classical subsolution to (23). We make the choice of control from (33) for (4) and (13) throughout this proof. Due to the assumed boundedness of these controls, we have via Lemma 5.1:

$$\begin{aligned} & \mathbb{E} \left[\exp \left(-2NG(\hat{\mu}_T^{N,s,y}) \right) \prod_{i=1}^N (Z^{i,N,s,y})^2 \right] \\ &= \mathbb{E} \left[\exp \left(-2NG(\tilde{\mu}_T^{N,s,y}) \right) \exp \left(N \int_s^T \int_{\mathbb{R}^d} |\sigma^\top(x, \tilde{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \tilde{\mu}_t^N)[x]|^2 \tilde{\mu}_t^{N,s,y}(dx) dt \right) \right], \end{aligned}$$

where $\tilde{\mu}_t^{N,s,y} = \frac{1}{N} \sum_{i=1}^N \delta_{\tilde{X}_t^{i,N,s,y}}$ and $\tilde{X}^{i,N,s,y}$ satisfy:

$$\begin{aligned} d\tilde{X}_t^{i,N,s,y} &= [b(\tilde{X}_t^{i,N,s,y}, \tilde{\mu}_t^{N,s,y}) + \sigma \sigma^\top(\tilde{X}_t^{i,N,s,y}, \tilde{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \tilde{\mu}_t^{N,s,y})[\tilde{X}_t^{i,N,s,y}]] dt \\ &\quad + \sigma(\tilde{X}_t^{i,N,s,y}, \tilde{\mu}_t^{N,s,y}) d\tilde{W}_t^i, \quad \tilde{X}_s^{i,N,s,y} = y_i. \end{aligned}$$

Strong existence and uniqueness of the above system of SDEs of all N follows from the assumed boundedness of σ and $\partial_\mu \Psi$ along with Assumption (A3) and an application of Girsanov theorem—see the discussion on p.81 of [16].

Applying Proposition (1.3) with $F(\mu) = 2G(\mu_T) - \int_s^T \int_{\mathbb{R}^d} |\sigma^\top(z, \mu_t) \partial_\mu \Psi(t, \mu_t)[z]|^2 \mu_t(dz) dt$ (F is bounded by assumption and its continuity follows from, e.g., [34] Theorem A.3.18) and $\tilde{b}(t, x, \nu) = b(x, \nu) + \sigma \sigma^\top(x, \nu) \partial_\mu \Psi(t, \nu)[x]$, we get

(60)

$$\begin{aligned} -\frac{1}{N} \log \mathbb{E} \left[\exp \left(-2NG(\hat{\mu}_T^{N,s,y}) \right) \prod_{i=1}^N (Z^{i,N,s,y})^2 \right] &= \inf_{u^N \in \mathcal{U}^N} \left\{ \mathbb{E} \left[\frac{1}{2N} \sum_{i=1}^N \int_s^T |u_i^N(t)|^2 dt \right] + 2\mathbb{E} \left[G(\bar{\mu}_T^{N,s,y}) \right] \right. \\ &\quad \left. - \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \int_s^T |\sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y})[\bar{X}_t^{i,N,s,y}]|^2 dt \right] \right\}, \end{aligned}$$

where $\bar{\mu}_t^{N,s,y} = \frac{1}{N} \sum_{i=1}^N \delta_{\bar{X}_t^{i,N,s,y}}$ and

$$\begin{aligned} d\bar{X}_t^{i,N,s,y} &= [b(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) + \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y})[\bar{X}_t^{i,N,s,y}]] dt \\ &\quad + \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) d\tilde{W}_t^i, \\ \bar{X}_s^{i,N,s,y} &= y_i. \end{aligned}$$

Fix any $u^N \in \mathcal{U}^N$ and let $\{\bar{X}^{i,N,s,y}\}_{i=1}^N$ satisfy the above controlled equations with this choice of control.

Letting $\tilde{\Psi}(t, x_1, \dots, x_N) = \Psi(t, \frac{1}{N} \sum_{i=1}^N \delta_{x_i})$, where Ψ is our classical subsolution, we apply Itô's formula to get

$$\begin{aligned} d\Psi(t, \bar{\mu}_t^{N,s,y}) &= \left[\partial_t \Psi(t, \bar{\mu}_t^{N,s,y}) + \sum_{i=1}^N \left\{ [b(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) + \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y})[\bar{X}_t^{i,N,s,y}]] \right. \right. \\ &\quad \left. \left. + \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) u_i^N(t) \right] \cdot \partial_{x_i} \tilde{\Psi}(t, \bar{X}_t^{i,N,s,y}, \dots, \bar{X}_t^{N,N,s,y}) \right. \\ &\quad \left. + \frac{1}{2} \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_{x_i}^2 \tilde{\Psi}(t, \bar{X}_t^{i,N,s,y}, \dots, \bar{X}_t^{N,N,s,y}) \right\} \Big] dt \\ &\quad + \sum_{i=1}^N [\partial_{x_i} \tilde{\Psi}]^\top(t, \bar{X}_t^{i,N,s,y}, \dots, \bar{X}_t^{N,N,s,y}) \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) d\tilde{W}_t^i \end{aligned}$$

$$\begin{aligned}
&= \left[\partial_t \Psi(t, \bar{\mu}_t^{N,s,y}) + \frac{1}{N} \sum_{i=1}^N \left\{ [b(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) + \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]] \right. \right. \\
&\quad \left. \left. + \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) u_i^N(t) \right] \cdot \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}] + \frac{1}{2} \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_z \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}] \right\} \Big] dt \\
&\quad + \frac{1}{N} \sum_{i=1}^N [\partial_\mu \Psi]^\top(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}] \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) d\tilde{W}_t^i \\
&\quad + \frac{1}{N^2} \sum_{i=1}^N \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{i,N,s,y}] dt,
\end{aligned}$$

where in the second step we applied Proposition A.3. Note here we have used the assumed regularity of Ψ from (i) in Definition 2.1.

Now using that Ψ satisfies (ii) in Definition 2.1, we have:

$$\begin{aligned}
&\partial_t \Psi(t, \bar{\mu}_t^{N,s,y}) + \frac{1}{N} \sum_{i=1}^N \left\{ [b(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \right. \\
&\quad \left. + \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]] \cdot \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}] \right. \\
&\quad \left. + \frac{1}{2} \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_z \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}] \right\} \\
&\geq \frac{3}{2} \frac{1}{N} \sum_{i=1}^N \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}] \cdot [\sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]] \\
&= \frac{3}{2} \frac{1}{N} \sum_{i=1}^N |\sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]|^2.
\end{aligned}$$

Combining the above two displays:

$$\begin{aligned}
\Psi(T, \bar{\mu}_T^{N,s,y}) &\geq \Psi \left(s, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) + \frac{1}{N} \sum_{i=1}^N \int_s^T u_i^N(t) \cdot [\sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]] dt \\
&\quad + \frac{1}{N} \sum_{i=1}^N \int_s^T [\partial_\mu \Psi]^\top(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}] \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) d\tilde{W}_t^i \\
&\quad + \frac{1}{N^2} \sum_{i=1}^N \int_s^T \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{i,N,s,y}] dt \\
&\quad + \int_s^T \frac{3}{2} \frac{1}{N} \sum_{i=1}^N |\sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]|^2 dt \\
&= \Psi \left(s, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) + \frac{1}{2N} \sum_{i=1}^N \int_s^T |u_i^N(t) + \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]|^2 dt \\
&\quad + \int_s^T \frac{1}{N} \sum_{i=1}^N |\sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]|^2 dt \\
&\quad - \int_s^T \frac{1}{2N} \sum_{i=1}^N |u_i^N(t)|^2 dt + \frac{1}{N} \sum_{i=1}^N \int_s^T [\partial_\mu \Psi]^\top(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}] \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) d\tilde{W}_t^i \\
&\quad + \frac{1}{N^2} \sum_{i=1}^N \int_s^T \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{i,N,s,y}] dt.
\end{aligned}$$

Then, for any choice of $u^N \in \mathcal{U}^N$:

$$\begin{aligned}
& \mathbb{E} \left[\frac{1}{2N} \sum_{i=1}^N \int_s^T |u_i^N(t)|^2 dt \right] + 2\mathbb{E} \left[G(\bar{\mu}_T^{N,s,y}) \right] - \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \int_s^T |\sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]|^2 dt \right] \\
& \geq \mathbb{E}[G(\bar{\mu}_T^{N,s,y})] + \mathbb{E}[\Psi(T, \bar{\mu}_T^{N,s,y})] + \mathbb{E} \left[\frac{1}{2N} \sum_{i=1}^N \int_s^T |u_i^N(t)|^2 dt \right] \\
& \quad - \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \int_s^T |\sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]|^2 dt \right] \\
& \geq \mathbb{E}[G(\bar{\mu}_T^{N,s,y})] + \Psi \left(s, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) + \mathbb{E} \left[\frac{1}{2N} \sum_{i=1}^N \int_s^T |u_i^N(t) + \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]|^2 dt \right] \\
& \quad + \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \int_s^T [\partial_\mu \Psi]^\top(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}] \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) d\tilde{W}_t^i \right] \\
& \quad + \mathbb{E} \left[\frac{1}{N^2} \sum_{i=1}^N \int_s^T \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{i,N,s,y}] dt \right],
\end{aligned}$$

where in the first inequality we used (iii) from Definition 2.1.

$u_i^N(t) + \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]$, $i = 1, \dots, N$ is a valid choice of control in the prelimit representation (12) from Proposition 1.3 since σ , $\partial_\mu \Psi$ are bounded, and taking $F(\mu) = G(\mu_T)$, $\tilde{b} = b$ therein (using $G \in C_b(\mathcal{P}(\mathbb{R}^d))$ and (A3)) we can continue:

$$\begin{aligned}
& \mathbb{E} \left[\frac{1}{2N} \sum_{i=1}^N \int_s^T |u_i^N(t)|^2 dt \right] + 2\mathbb{E} \left[G(\bar{\mu}_T^{N,s,y}) \right] - \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \int_s^T |\sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) \partial_\mu \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}]|^2 dt \right] \\
& \geq \Psi \left(s, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) + \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \int_s^T [\partial_\mu \Psi]^\top(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}] \sigma(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) d\tilde{W}_t^i \right] \\
& \quad + \mathbb{E} \left[\frac{1}{N^2} \sum_{i=1}^N \int_s^T \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{i,N,s,y}] dt \right] \\
& \quad - \frac{1}{N} \log \mathbb{E} \left[\exp \left(-NG(\bar{\mu}_T^{N,s,y}) \right) \right] \\
& = \Psi \left(s, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) + \mathbb{E} \left[\frac{1}{N^2} \sum_{i=1}^N \int_s^T \sigma \sigma^\top(\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{i,N,s,y}] dt \right] \\
& \quad - \frac{1}{N} \log \mathbb{E} \left[\exp \left(-NG(\bar{\mu}_T^{N,s,y}) \right) \right]
\end{aligned}$$

since the martingale term is bounded in square expectation for all N .

Infimizing over $u^N \in \mathcal{U}^N$ and using (60) we get:

$$\begin{aligned}
& \liminf_{N \rightarrow \infty} -\frac{1}{N} \log \left(\mathbb{E} \left[\exp \left(-2NG(\bar{\mu}_T^{N,s,y}) \right) \prod_{i=1}^N (Z^{i,N,s,y})^2 \right] \right) \\
& \geq \liminf_{N \rightarrow \infty} \left\{ \Psi \left(s, \frac{1}{N} \sum_{i=1}^N \delta_{y_i} \right) \right\}
\end{aligned}$$

$$\begin{aligned}
& + \inf_{u^N \in \mathcal{U}^N} \mathbb{E} \left[\frac{1}{N^2} \sum_{i=1}^N \int_s^T \sigma \sigma^\top (\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{i,N,s,y}] dt \right] \\
& - \frac{1}{N} \log \mathbb{E} \left[\exp(-NG(\mu_T^{N,s,y})) \right] \Big\} \\
& = \Psi(s, \nu) + \gamma_1,
\end{aligned}$$

where we used the assumed continuity of Ψ and the convergence of $\frac{1}{N} \sum_{i=1}^N \delta_{y_i}$ to ν , (18), and the fact that by boundedness of σ and (iv) in Definition 2.1, there is a constant C independent on the choice of $u^N \in \mathcal{U}^N$ (which may change from line to line) such that:

$$\begin{aligned}
& \mathbb{E} \left[\frac{1}{N^2} \sum_{i=1}^N \int_s^T \sigma \sigma^\top (\bar{X}_t^{i,N,s,y}, \bar{\mu}_t^{N,s,y}) : \partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{i,N,s,y}] dt \right] \\
& \leq C \mathbb{E} \left[\frac{1}{N^2} \sum_{i=1}^N \int_s^T |\partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{i,N,s,y}]| dt \right] \\
& = C \mathbb{E} \left[\frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \int_s^T |\partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{j,N,s,y}]| \mathbb{1}_{i=j} dt \right] \\
& \leq C \mathbb{E} \left[\int_s^T \left(\frac{1}{N^2} \sum_{i,j=1}^N |\partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{j,N,s,y}]|^2 \right)^{1/2} \left(\frac{1}{N^2} \sum_{i,j=1}^N \mathbb{1}_{i=j} \right)^{1/2} dt \right] \\
& = C \mathbb{E} \left[\frac{1}{N^{1/2}} \int_s^T \left(\frac{1}{N^2} \sum_{i,j=1}^N |\partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) [\bar{X}_t^{i,N,s,y}, \bar{X}_t^{j,N,s,y}]|^2 \right)^{1/2} dt \right] \\
& = C \mathbb{E} \left[\frac{1}{N^{1/2}} \int_s^T \left\| \partial_\mu^2 \Psi(t, \bar{\mu}_t^{N,s,y}) \right\|_{L^2(\mathbb{R}^d; \bar{\mu}_t^{N,s,y}) \otimes L^2(\mathbb{R}^d; \bar{\mu}_t^{N,s,y})} dt \right] \\
& \leq \frac{C}{N^{1/2}} (T - s).
\end{aligned}$$

Note that if we, instead of just the bound $\sup_{t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d)} \left\| \partial_\mu^2 \Psi(t, \mu) [\cdot, \cdot] \right\|_{L^2(\mathbb{R}^d; \mu) \otimes L^2(\mathbb{R}^d; \mu)} \leq C$, have further that e.g. $\sup_{t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d)} \int_{\mathbb{R}^d} |\partial_\mu^2 \Psi(t, \mu) [z, z]|^2 \mu(dz) \leq C$, then this remainder is $O(1/N)$. Also note that the Laplace Principle of Theorem 1.2 is used here to establish the equality (18), which yields the appearance of γ_1 in the limit.

This gives that

$$\begin{aligned}
& \lim_{N \rightarrow \infty} -\frac{1}{N} \log R(\hat{\delta}_N) \\
& = \lim_{N \rightarrow \infty} \left\{ -\frac{1}{N} \log \mathbb{E} \left[\exp(-2NG(\hat{\mu}^{N,s,y})) \prod_{i=1}^N (Z^{i,N,s,y})^2 \right] + \frac{2}{N} \log \mathbb{E} \left[\exp(-NG(\mu^{N,s,y})) \right] \right\} \\
& = \lim_{N \rightarrow \infty} -\frac{1}{N} \log \mathbb{E} \left[\exp(-2NG(\hat{\mu}^{N,s,y})) \prod_{i=1}^N (Z^{i,N,s,y})^2 \right] - 2\gamma_1 \quad (\text{by (18)}) \\
& \geq \Psi(s, \nu) - \gamma_1.
\end{aligned}$$

We thus have that $\hat{\delta}_N$ admits an expansion of the form (20) with $\hat{\gamma} \leq \gamma_1 - \Psi(s, \nu)$. Note that the case $\Psi \equiv 0$ corresponds to the standard Monte Carlo estimator δ_N in the above, from which we obtain $\gamma = 2\gamma_1 - \gamma_2 \leq \gamma_1$ as expected – see (19) and the discussion thereafter. The rest of the claims stated in the theorem are now immediate. In particular, by Jensen's inequality we know $-\frac{1}{N} \log R(\hat{\delta}_N) \leq 0$, so if $\Psi(s, \nu) = \gamma_1$ we have log-efficiency in the sense of Definition 1.4.

5.2. Proof of Theorem 2.5. First, we will establish a prelimit PDE expression for the numerator in the expression for $R(\hat{\delta}_N)$ in (17) in a manner along the lines of the discussion in Section 2.1. Define

$$\hat{\Xi}^N(t, x_1, \dots, x_N) = \mathbb{E} \left[\exp \left(-2NG(\hat{\mu}_T^{N,t,x}) \right) \prod_{i=1}^N (Z^{i,N,t,x})^2 \right], \quad x_1, \dots, x_N \in \mathbb{R}^d, t \in [0, T]$$

where $Z^{i,N,t,x}$, $i = 1, \dots, N$ and $\hat{\mu}_T^{N,t,x}$ are as in (15) and (13) with the choice of controls from (33) in Theorem 2.4. Define $v(t, \mu, x) := -\sigma^\top(x, \mu) \partial_\mu \phi_0(t, \mu)[x]$. Then $v(t, \mu_x^N, x_i) = v_i^N(t, x_1, \dots, x_N)$ for all N , where v_i^N are from (33) with ϕ_0 in the place of Ψ . Recall that in (33) we had only assumed Ψ be a classical subsolution of (23), whereas now we assume ϕ_0 is a classical solution.

From Lemma 5.1 (using the assumed boundedness of σ and $\partial_\mu \phi_0$), we have:

$$(61) \quad \hat{\Xi}^N(t, x_1, \dots, x_N) = \mathbb{E} \left[\exp(-2NG(\hat{\mu}_T^{N,t,x})) \exp \left(N \int_t^T \int_{\mathbb{R}^d} |v(\tau, \tilde{\mu}_\tau^{N,t,x}, z)|^2 \tilde{\mu}_\tau^{N,t,x}(dz) d\tau \right) \right]$$

where $\tilde{\mu}_\tau^{N,t,x} = \frac{1}{N} \sum_{i=1}^N \delta_{\tilde{X}_\tau^{i,N,t,x}}$,

$$d\tilde{X}_\tau^{i,N,t,x} = [b(\tilde{X}_\tau^{i,N,t,x}, \tilde{\mu}_\tau^{N,t,x}) - \sigma(\tilde{X}_\tau^{i,N,t,x}, \tilde{\mu}_\tau^{N,t,x})v(\tau, \tilde{\mu}_\tau^{N,t,x}, \tilde{X}_\tau^{i,N,t,x})]d\tau + \sigma(\tilde{X}_\tau^{i,N,t,x}, \tilde{\mu}_\tau^{N,t,x})d\tilde{W}_\tau^i, \\ \tilde{X}_t^{i,N,t,x} = x_i.$$

Applying Feynman-Kac, we have $\hat{\Xi}^N$ is the unique solution to:

$$(62) \quad \begin{aligned} \partial_t \hat{\Xi}^N(t, x_1, \dots, x_N) + \sum_{i=1}^N \left\{ [b(x_i, \mu_x^N) - \sigma(x_i, \mu_x^N)v(t, \mu_x^N, x_i)] \cdot \partial_{x_i} \hat{\Xi}^N(t, x_1, \dots, x_N) \right. \\ \left. + \frac{1}{2} \sigma \sigma^\top(x_i, \mu_x^N) : \partial_{x_i}^2 \hat{\Xi}^N(t, x_1, \dots, x_N) + |v(t, \mu_x^N, x_i)|^2 \hat{\Xi}^N(t, x_1, \dots, x_N) \right\} = 0, \\ t \in [0, T), x_1, \dots, x_N \in \mathbb{R}^d, \end{aligned}$$

$$\hat{\Xi}^N(T, x_1, \dots, x_N) = \exp(-2NG(\mu_x^N)), \quad x_1, \dots, x_N \in \mathbb{R}^d.$$

This is where the linear growth condition on b is used—see, e.g., Theorem 7.6 in [64].

Then $\tilde{\Xi}^N(t, x_1, \dots, x_N) := -\frac{1}{N} \log \hat{\Xi}^N(t, x_1, \dots, x_N)$ satisfies

$$(63) \quad \begin{aligned} \partial_t \tilde{\Xi}^N(t, x_1, \dots, x_N) + \sum_{i=1}^N \left\{ [b(x_i, \mu_x^N) - \sigma(x_i, \mu_x^N)v(t, \mu_x^N, x_i)] \cdot \partial_{x_i} \tilde{\Xi}^N(t, x_1, \dots, x_N) \right. \\ \left. + \frac{1}{2} \sigma \sigma^\top(x_i, \mu_x^N) : \partial_{x_i}^2 \tilde{\Xi}^N(t, x_1, \dots, x_N) - \frac{1}{N} |v(t, \mu_x^N, x_i)|^2 \right. \\ \left. - \frac{N}{2} |\sigma^\top(x_i, \mu_x^N) \partial_{x_i} \tilde{\Xi}^N(t, x_1, \dots, x_N)|^2 \right\} = 0, t \in [0, T), x_1, \dots, x_N \in \mathbb{R}^d, \\ \tilde{\Xi}^N(T, x_1, \dots, x_N) = 2G(\mu_x^N), \quad x_1, \dots, x_N \in \mathbb{R}^d. \end{aligned}$$

We consider now the PDE (34) from Subsection 2.2. Note that this agrees with (32) if we set $v \equiv 0$, other than the fact that the terminal condition is $2G$ rather than G . Under the current regularity assumptions, we see via Proposition A.3 that for all N , $\tilde{\Xi}^N$ is the empirical projection of Ξ^N from (34). That is:

$$(63) \quad \Xi^N(t, \mu_x^N) = \tilde{\Xi}^N(t, x_1, \dots, x_N) = -\frac{1}{N} \log \mathbb{E} \left[\exp(-2NG(\hat{\mu}_T^{N,t,x})) \prod_{i=1}^N (Z^{i,N,t,x})^2 \right].$$

We now observe how to uncover the log-efficiency established in Theorem 2.4 (though under stronger regularity assumptions) using the method of [75], and set up how higher order terms in such an expansion can be used to prove the results in Theorem 2.5.

Suppose, as in [75], that we have sufficient regularity to have the asymptotic expansions for (32) and (34) of the forms (36) and (37) for some k .

Then by (61) and (25) (noting that indeed the representation (25) holds for the unique solution $\hat{\Phi}^N$ to (26), again by Theorem 7.6 in [64]), we have

$$\begin{aligned} R(\hat{\delta}_N) &= \frac{\hat{\Xi}^N(s, y_1, \dots, y_N)}{[\hat{\Phi}^N(s, y_1, \dots, y_N)]^2} \\ &= \exp\left(N[2\phi_0(s, \mu_y^N) - \xi_0(s, \mu_y^N)] + [2\phi_1(s, \mu_y^N) - \xi_1(s, \mu_y^N)] + [2\phi_2(s, \mu_y^N) - \xi_2(s, \mu_y^N)]/N + \dots \right. \\ &\quad \left. + [2\phi_k(s, \mu_y^N) - \xi_k(s, \mu_y^N)]/N^{k-1} + o(1/N^{k-1})\right), \end{aligned}$$

where $R(\hat{\delta}_N)$ is as in (17). Here we have used, by the same logic as to obtain (63),

$$\Phi^N(t, \mu_x^N) = \tilde{\Phi}^N(t, x_1, \dots, x_N) = -\frac{1}{N} \log \mathbb{E}[\exp(-NG(\mu_T^{N,t,x}))],$$

where Φ^N is as in (32) and $\tilde{\Phi}^N$ is as in (29).

We now see that $2\phi_0$ (recalling here that ϕ_0 is the unique classical solution to (23)) and ξ_0 from (35) satisfy the same equation. Indeed, inserting this ansatz into (35) and using $v(t, \mu, x) = -\sigma^\top(x, \mu)\partial_\mu\phi_0(t, \mu)[x]$, we get

$$\begin{aligned} &2\partial_t\phi_0(t, \mu) + 2 \int_{\mathbb{R}^d} [b(z, \mu) + \sigma\sigma^\top(z, \mu)\partial_\mu\phi_0(t, \mu)[z]] \cdot \partial_\mu\phi_0(t, \mu)[z] + \frac{1}{2}\sigma\sigma^\top(z, \mu) : \partial_z\partial_\mu\phi_0(t, \mu)[z]\mu(dz) \\ &\quad - \int_{\mathbb{R}^d} 2|\sigma^\top(z, \mu)\partial_\mu\phi_0(t, \mu)[z]|^2 + |\sigma^\top(z, \mu)\partial_\mu\phi_0(t, \mu)[z]|^2\mu(dz) \\ &= 2\partial_t\phi_0(t, \mu) + 2 \int_{\mathbb{R}^d} b(z, \mu) \cdot \partial_\mu\phi_0(t, \mu)[z] + \frac{1}{2}\sigma\sigma^\top(z, \mu) : \partial_z\partial_\mu\phi_0(t, \mu)[z]\mu(dz) \\ &\quad - \int_{\mathbb{R}^d} |\sigma^\top(z, \mu)\partial_\mu\phi_0(t, \mu)[z]|^2\mu(dz) = 0, \quad t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d), \\ &2\phi_0(T, \mu) = 2G(\mu), \quad \mu \in \mathcal{P}_2(\mathbb{R}^d). \end{aligned}$$

which dividing by 2 gives the same equation as (23).

This shows that indeed $2\phi_0(t, \mu) = \xi_0(t, \mu)$ for all $t \in [0, T], \mu \in \mathcal{P}_2(\mathbb{R}^d)$ via our uniqueness assumption.

Thus, since $2\phi_0(s, \mu_y^N) = \xi_0(s, \mu_y^N)$ for all N , we have $-\frac{1}{N} \log R(\hat{\delta}_N) \rightarrow 0$ as $N \rightarrow \infty$.

Moreover, when $k = 1$, knowing $2\phi_0(s, \mu_y^N) = \xi_0(s, \mu_y^N)$, we have for any M

$$\lim_{N \rightarrow \infty} \rho(\hat{\delta}_N) = \frac{1}{\sqrt{M}} \sqrt{\lim_{N \rightarrow \infty} \exp(2\phi_1(s, \mu_y^N) - \xi_1(s, \mu_y^N) + o(1)) - 1}.$$

As we will show, under our current assumptions $2\phi_1(s, \mu_y^N) = \xi_1(s, \mu_y^N)$ for all N , so we get $\lim_{N \rightarrow \infty} \rho(\hat{\delta}_N) = 0$, yielding vanishing relative error (Definition 1.5).

Lastly, knowing $2\phi_0(s, \mu_y^N) = \xi_0(s, \mu_y^N)$ and $2\phi_1(s, \mu_y^N) = \xi_1(s, \mu_y^N)$, letting $k = 2$ we have:

$$\begin{aligned} \lim_{N \rightarrow \infty} T(N) &= \lim_{N \rightarrow \infty} N[\exp([2\phi_2(s, \mu_y^N) - \xi_2(s, \mu_y^N)]/N + o(1/N)) - 1] \\ &= \lim_{N \rightarrow \infty} [2\phi_2(s, \mu_y^N) - \xi_2(s, \mu_y^N)] + o(1) \\ &= \lim_{N \rightarrow \infty} [2\phi_2(s, \mu_y^N) - \xi_2(s, \mu_y^N)] \\ &= 2\phi_2(s, \nu) - \xi_2(s, \nu) \end{aligned}$$

where $T(N)$ is as in (42) and here we have used the assumed continuity of ϕ_2, ξ_2 . This will yield the limit (43).

We now insert the ansatz of the expansions (36) for $\Phi^N(t, \mu)$ and (37) for $\Xi^N(t, \mu)$ into their respective equations (32) and (34) and match $\mathcal{O}(1/N^k)$ to obtain formal recursive formulas for ξ_k and ϕ_k as in Theorem 3.3 in [75]. We arrive at the expressions (38) and (39) from Subsection 2.2.

Now recalling $2\phi_0(t, \mu) = \xi_0(t, \mu)$, we see that in (40) and (41) from Subsection 2.2, $Y^{t, \mu} \stackrel{d}{=} Z^{t, \mu}$ (using (A4)), and moreover we note that $Z^{t, \mu}$ is the solution to the (optimally) controlled limiting McKean-Vlasov

equation initialized at time t with distribution μ from Theorem 1.2, that is $Z^{t,\mu} \stackrel{d}{=} \hat{X}^{u,t,\mu}$ from (8) with $u(\tau) = -\sigma(\hat{X}_\tau^{u,t,\nu}, \mathcal{L}(\hat{X}_\tau^{u,t,\nu}))\partial_\mu\phi_0(\tau, \mathcal{L}(\hat{X}_\tau^{u,t,\nu}))[\hat{X}_\tau^{u,t,\nu}]$, $\tau \in [t, T]$.

Under the assumption that the expressions (40),(41) are valid with $k = 1$, we have:

$$\phi_1(t, \mu) = \frac{1}{2}\mathbb{E}\left[\int_t^T \sigma\sigma^\top(Z_\tau^{t,\mu}, \mathcal{L}(Z_\tau^{t,\mu})) : \partial_\mu^2\phi_0(\tau, \mathcal{L}(Z_\tau^{t,\mu}))[Z_\tau^{t,\mu}, Z_\tau^{t,\mu}]d\tau\right]$$

and

$$\begin{aligned}\xi_1(t, \mu) &= \mathbb{E}\left[\int_t^T \frac{1}{2}\sigma\sigma^\top(Z_\tau^{t,\mu}, \mathcal{L}(Z_\tau^{t,\mu})) : \partial_\mu^2\xi_0(\tau, \mathcal{L}(Z_\tau^{t,\mu}))[Z_\tau^{t,\mu}, Z_\tau^{t,\mu}]d\tau\right] \\ &= \mathbb{E}\left[\int_t^T \sigma\sigma^\top(Z_\tau^{t,\mu}, \mathcal{L}(Z_\tau^{t,\mu})) : \partial_\mu^2\phi_0(\tau, \mathcal{L}(Z_\tau^{t,\mu}))[Z_\tau^{t,\mu}, Z_\tau^{t,\mu}]d\tau\right],\end{aligned}$$

so indeed $\xi_1(t, \mu) = 2\phi_1(t, \mu)$ for all $t \in [0, T]$ and $\mu \in \mathcal{P}_2(\mathbb{R}^d)$. As discussed previously, this establishes that we indeed have vanishing relative error (Definition 1.5). Also, this is the expected expression to arrive at for these first-order correction terms—compare with equations (A.7) and (A.8) in [79], where the analogous expression is given in terms of the second derivative of their zero-viscosity (standard) HJB equation evaluated at the optimally controlled path for the limiting ODE in the small-noise regime.

Under the additional assumption that our representations (40),(41) for ϕ_k, ξ_k hold for the next order correction $k = 2$, we have:

$$\begin{aligned}\phi_2(t, \mu) &= \frac{1}{2}\mathbb{E}\left[\int_t^T \sigma\sigma^\top(Z_\tau^{t,\mu}, \mathcal{L}(Z_\tau^{t,\mu})) : \partial_\mu^2\phi_1(\tau, \mathcal{L}(Z_\tau^{t,\mu}))[Z_\tau^{t,\mu}, Z_\tau^{t,\mu}] \right. \\ &\quad \left. - |\sigma^\top(Z_\tau^{t,\mu}, \mathcal{L}(Z_\tau^{t,\mu}))\partial_\mu\phi_1(\tau, \mathcal{L}(Z_\tau^{t,\mu}))[Z_\tau^{t,\mu}]|^2d\tau\right]\end{aligned}$$

and

$$\begin{aligned}\xi_2(t, \mu) &= \mathbb{E}\left[\int_t^T \frac{1}{2}\sigma\sigma^\top(Z_\tau^{t,\mu}, \mathcal{L}(Z_\tau^{t,\mu})) : \partial_\mu^2\xi_1(\tau, \mathcal{L}(Z_\tau^{t,\mu}))[Z_\tau^{t,\mu}, Z_\tau^{t,\mu}] \right. \\ &\quad \left. - \frac{1}{2}|\sigma^\top(Z_\tau^{t,\mu}, \mathcal{L}(Z_\tau^{t,\mu}))\partial_\mu\xi_1(\tau, \mathcal{L}(Z_\tau^{t,\mu}))[Z_\tau^{t,\mu}]|^2d\tau\right] \\ &= \mathbb{E}\left[\int_t^T \sigma\sigma^\top(Z_\tau^{t,\mu}, \mathcal{L}(Z_\tau^{t,\mu})) : \partial_\mu^2\phi_1(\tau, \mathcal{L}(Z_\tau^{t,\mu}))[Z_\tau^{t,\mu}, Z_\tau^{t,\mu}]d\tau \right. \\ &\quad \left. - 2|\sigma^\top(Z_\tau^{t,\mu}, \mathcal{L}(Z_\tau^{t,\mu}))\partial_\mu\phi_1(\tau, \mathcal{L}(Z_\tau^{t,\mu}))[Z_\tau^{t,\mu}]|^2d\tau\right],\end{aligned}$$

so

$$2\phi_2(t, \mu) - \xi_2(t, \mu) = \mathbb{E}\left[\int_t^T |\sigma^\top(Z_\tau^{t,\mu}, \mathcal{L}(Z_\tau^{t,\mu}))\partial_\mu\phi_1(\tau, \mathcal{L}(Z_\tau^{t,\mu}))[Z_\tau^{t,\mu}]|^2d\tau\right].$$

By our previous discussion and our identification of $Z^{t,\mu}$ as related to $\hat{X}^{u,t,\mu}$, we get

$$\lim_{N \rightarrow \infty} T(N) = \mathbb{E}\left[\int_s^T |\sigma^\top(\hat{X}_t^{u,s,\nu}, \mathcal{L}(\hat{X}_t^{u,s,\nu}))\partial_\mu\phi_1(t, \mathcal{L}(\hat{X}_t^{u,s,\nu}))[\hat{X}_t^{u,s,\nu}]|^2dt\right]$$

where ϕ_1 is as in (38) with $k = 1$ and $u(t) = -\sigma^\top(\hat{X}_t^{u,s,\nu}, \mathcal{L}(\hat{X}_t^{u,s,\nu}))\partial_\mu\phi_0(t, \mathcal{L}(\hat{X}_t^{u,s,\nu}))[\hat{X}_t^{u,s,\nu}]$, $t \in [s, T]$.

6. CONCLUSIONS AND FUTURE WORK

We have derived an importance sampling scheme for exponential functionals of the empirical measure of weakly interacting diffusions. Using the connection between the large deviations rate function of [16] and mean-field optimal control, the asymptotic performance of the proposed scheme is characterized in terms of subsolutions of the Hamilton-Jacobi-Bellman equation on Wasserstein space. We provide both numerical and analytical evidence that sufficient smoothness of such a solution can yield relative error which vanishes

as the number of particles becomes large. We also numerically explore the impact of lack of smoothness of the solution of the HJB equation on the performance of the proposed importance sampling scheme.

In future work, a major hurdle to overcome will be adapting the scheme to situations where subsolutions of the HJB equation cannot be constructed analytically. In the standard small noise setting, this is also a major issue, and when the dimension of the system becomes large advanced techniques such as the use of machine learning, optimization software, and neural networks to identify the solution of the HJB equation and/or the optimal control are used [11, 71, 77, 79]. In the mean-field setting, methods of stochastic control are already being used to address the issue of numerically constructing solutions to the HJB equation (23), see e.g. [20, 22, 47, 54, 55, 60, 66]. Moreover, using the calculus of variations form of the rate function of Dawson-Gärtner [29], there are some examples in the literature where perturbation expansions have been made to approximate the optimal path of the controlled McKean-Vlasov equation corresponding to certain types of rare events [12, 49]. The marriage of such techniques with our proposed scheme would allow for applications beyond the linear quadratic regime and perturbations thereof, and perhaps even allow for extensions to finite-time probabilities and problems of metastability as discussed in Remark 2.6.

It may also prove useful for some target statistics to design an importance sampling scheme using the moderate deviations principle for the empirical measure [10, 17]. This is known to aid with the problems discussed above in the small noise setting due to the linearization of the HJB equation under the moderate deviations scaling [51, 69]. As discussed in Remark 3.4, this can likely also be supplemented via use of an importance sampling scheme arising from large deviations of the empirical measure in the joint small noise and large N limit as derived in [14, 72].

Another interesting avenue for future research is to see how our methodology can be extended to the setting where the interacting particles also share a common driving noise [15, 31].

Lastly, it is of great interest to establish rigorously the asymptotic expansions for the prelimit HJB equation presumed in Theorem 2.5—see Remark 2.7.

APPENDIX A. DIFFERENTIATION ON SPACES OF MEASURES

Definition A.1. *Given a function $u : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$, we may define a lifting of u to $\tilde{u} : L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d) \rightarrow \mathbb{R}$ via $\tilde{u}(X) = u(\mathcal{L}(X))$ for $X \in L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d)$. Here we assume $\tilde{\Omega}$ is a Polish space, $\tilde{\mathcal{F}}$ its Borel σ -field, and $\tilde{\mathbb{P}}$ is an atomless probability measure (since $\tilde{\Omega}$ is Polish, this is equivalent to every singleton having zero measure).*

Here:

$$\mathcal{P}_2(\mathbb{R}^d) := \{\mu \in \mathcal{P}(\mathbb{R}^d) : \int_{\mathbb{R}^d} |x|^2 \mu(dx) < \infty\}.$$

$\mathcal{P}_2(\mathbb{R}^d)$ is a Polish space under the L^2 -Wasserstein distance

$$\mathbb{W}_2(\mu_1, \mu_2) := \inf_{\pi \in \mathcal{C}_{\mu_1, \mu_2}} \left[\int_{\mathbb{R}^d \times \mathbb{R}^d} |x - y|^2 \pi(dx, dy) \right]^{1/2},$$

where $\mathcal{C}_{\mu_1, \mu_2}$ denotes the set of all probability measures on $\mathbb{R}^d \times \mathbb{R}^d$ with first marginal μ_1 and second marginal μ_2 .

We say u is **L-differentiable** or **Lions-differentiable** at $\mu_0 \in \mathcal{P}_2(\mathbb{R}^d)$ if there exists a random variable X_0 on some $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ satisfying the above assumptions such that $\mathcal{L}(X_0) = \mu_0$ and $\tilde{u}$ is Fréchet differentiable at X_0 .

The Fréchet derivative of $\tilde{u}$ can be viewed as an element of $L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d)$ by identifying $L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d)$ and its dual. From this, one can find that if u is L-differentiable at $\mu_0 \in \mathcal{P}_2(\mathbb{R}^d)$, there is a deterministic measurable function $\xi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ such that $D\tilde{u}(X_0) = \xi(X_0)$, and that ξ is uniquely defined μ_0 -almost everywhere on $\mathbb{R}^d$. We denote this equivalence class of $\xi \in L^2(\mathbb{R}^d, \mu_0; \mathbb{R}^d)$ by $\partial_\mu u(\mu_0)$ and call $\partial_\mu u(\mu_0)[\cdot] : \mathbb{R}^d \rightarrow \mathbb{R}^d$ the **Lions derivative** of u at μ_0 . Note that this definition is independent of the choice of X_0 and $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$. See [19] Section 5.2.

To avoid confusion when u depends on more variables than just μ , if $\partial_\mu u(\mu_0)$ is differentiable at $z_0 \in \mathbb{R}^d$, we denote its derivative at z_0 by $\partial_z \partial_\mu u(\mu_0)[z_0]$.

Definition A.2. ([19] Definition 5.83) We say $u : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ is **Fully C^2** if the following conditions are satisfied:

- (1) u is C^1 in the sense of L -differentiation, and its first derivative has a jointly continuous version $\mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \ni (\mu, z) \mapsto \partial_\mu u(\mu)[z] \in \mathbb{R}^d$.
- (2) For each fixed $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, the version of $\mathbb{R}^d \ni z \mapsto \partial_\mu u(\mu)[z] \in \mathbb{R}^d$ from the first condition is differentiable on $\mathbb{R}^d$ in the classical sense and its derivative is given by a jointly continuous function $\mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \ni (\mu, z) \mapsto \partial_z \partial_\mu u(\mu)[z] \in \mathbb{R}^{d \times d}$.
- (3) For each fixed $z \in \mathbb{R}^d$, the version of $\mathcal{P}_2(\mathbb{R}^d) \ni \mu \mapsto \partial_\mu u(\mu)[z] \in \mathbb{R}^d$ in the first condition is continuously L -differentiable component-by-component, with a derivative given by a function $\mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \times \mathbb{R}^d \ni (\mu, z, z') \mapsto \partial_{\mu'}^2 u(\mu)[z][z'] \in \mathbb{R}^{d \times d}$ such that for any $\mu \in \mathcal{P}_2(\mathbb{R}^d)$ and $X \in L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d)$ with $\mathcal{L}(X) = \mu$, $\partial_\mu^2 u(\mu)[z][X]$ gives the Fréchet derivative at X of $L^2(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}}; \mathbb{R}^d) \ni X' \mapsto \partial_\mu u(\mathcal{L}(X'))[z]$ for every $z \in \mathbb{R}^d$. Denoting $\partial_{\mu'}^2 u(\mu)[z][z']$ by $\partial_\mu^2 u(\mu)[z, z']$, the map $\mathcal{P}_2(\mathbb{R}^d) \times \mathbb{R}^d \times \mathbb{R}^d \ni (\mu, z, z') \mapsto \partial_\mu^2 u(\mu)[z, z']$ is also assumed to be continuous in the product topology.

We recall now a useful connection between the Lions derivative as defined in A.1 and the empirical measure.

Proposition A.3. For $g : \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ which is fully C^2 in the sense of definition A.2, we can define the empirical projection of g , as $g^N : (\mathbb{R}^d)^N \rightarrow \mathbb{R}$ given by

$$g^N(x_1, \dots, x_N) := g\left(\frac{1}{N} \sum_{i=1}^N \delta_{x_i}\right).$$

Then g^N is twice differentiable on $(\mathbb{R}^d)^N$, and for each $x_1, \dots, x_N \in \mathbb{R}^d$, $(i, j) \in \{1, \dots, N\}^2$:

$$(64) \quad \partial_{x_i} g^N(x_1, \dots, x_N) = \frac{1}{N} \partial_\mu g\left(\frac{1}{N} \sum_{i=1}^N \delta_{x_i}\right)[x_i]$$

and

$$(65) \quad \partial_{x_i} \partial_{x_j} g^N(x_1, \dots, x_N) = \frac{1}{N} \partial_z \partial_\mu g\left(\frac{1}{N} \sum_{i=1}^N \delta_{x_i}\right)[x_i] \mathbb{1}_{i=j} + \frac{1}{N^2} \partial_\mu^2 g\left(\frac{1}{N} \sum_{i=1}^N \delta_{x_i}\right)[x_i, x_j].$$

Proof. This follows from Propositions 5.35 and 5.91 of [19]. $\square$

DECLARATIONS

Bezemek was partially supported by NSF-DMS 2107856. Heldman was partially supported by NSF-DMS 1902854, ARO W911NF-20-1-0244, and a subgrant of NSF-OAC 2139536. The funding agencies are not expected to gain or lose financially through publication of this manuscript. The authors have no relevant financial interests to disclose.

REFERENCES

- [1] S. Adams, N. Dirr, M. A. Peletier, and J. Zimmer, *From a Large-Deviations Principle to the Wasserstein Gradient Flow: A New Micro-Macro Passage*, Communications in Mathematical Physics **307** (2011), no. 3.
- [2] ———, *Large Deviations and Gradient Flows*, Philosophical Transactions of the Royal Society A **371** (2013), no. 2005.
- [3] V. Babin, C. Roland, and C. Sagui, *Adaptively Biased Molecular Dynamics for Free Energy Calculations*, The Journal of Chemical Physics **128** (2008), no. 13, 128–135.
- [4] J. Backhoff, G. Conforti, I. Gentil, and C. Léonard, *The Mean Field Schrödinger Problem: Ergodic Behavior, Entropy Estimates and Functional Inequalities*, Probability Theory and Related Fields **178** (2020), no. 2, 475–530.
- [5] M. Basei and H. Pham, *A Weak Martingale Approach to Linear–Quadratic McKean–Vlasov Stochastic Control Problems*, Journal of Optimization Theory and Application **181** (2019), no. 2, 347–382.
- [6] K. Bashiri and G. Menz, *Metastability in a Continuous Mean-Field Model at Low Temperature and Strong Interaction*, Stochastic Processes and their Applications **134** (2021), 132–173.
- [7] N. Ben Rached, A.-L. Haji-Ali, S.M.S. Pillai, and R. Tempone, *Multilevel Importance Sampling for McKean-Vlasov Stochastic Differential Equation* (2022), available at [arXiv:2208.03225](https://arxiv.org/abs/2208.03225) [math.NA].

- [8] ———, *Single Level Importance Sampling for McKean-Vlasov Stochastic Differential Equations* (2022), available at [arXiv: 2207.06926 \[math.NA\]](#).
- [9] Z.W. Bezemek and K. Spiliopoulos, *Large Deviations for Interacting Multiscale Particle Systems*, *Stochastic Processes and their Applications* **155** (2023), 27–108.
- [10] ———, *Moderate Deviations for Fully Coupled Multiscale Weakly Interacting Particle Systems*, *Stochastics and Partial Differential Equations: Analysis and Computations* (2023).
- [11] E.R. Borrell, J. Quer, L. Richter, and C. Schütte, *Improving Control Based Importance Sampling Strategies for Metastable Diffusions via Adapted Metadynamics*, *SIAM Journal on Scientific Computing* (2023), S298–S323.
- [12] F. Bouchet, K. Gawędzki, and C. Nardini, *Perturbative Calculation of Quasi-Potential in Non-equilibrium Diffusions: A Mean-Field Example*, *Journal of Statistical Physics* **163** (2016), 1157–1210.
- [13] M. Boué and P. Dupuis, *A Variational Representation for Certain Functionals of Brownian Motion*, *The Annals of Probability* **26** (1998), no. 4, 1641–1659.
- [14] A. Budhiraja and M. Conroy, *Asymptotic behavior of stochastic currents under large deviation scaling with mean field interaction and vanishing noise*, *Annali della Scuola Normale di Pisa - Classe di Scienze* **23** (2022), no. 5, 1749–1805.
- [15] ———, *Empirical Measure and Small Noise Asymptotics Under Large Deviation Scaling for Interacting Diffusions*, *Journal of Theoretical Probability* **35** (2022), 295–349.
- [16] A. Budhiraja, P. Dupuis, and M. Fischer, *Large Deviation Properties of Weakly Interacting Particles via Weak Convergence Methods*, *The Annals of Probability* **40** (2012), no. 1, 74–100.
- [17] A. Budhiraja and R. Wu, *Moderate Deviation Principles for Weakly Interacting Particle Systems*, *Probability Theory and Related Fields* **168** (2017), 721–771.
- [18] M. Burzoni, V. Ignazio, A.M. Reppen, and H.M. Soner, *Viscosity Solutions for Controlled McKean-Vlasov Jump-Diffusions*, *SIAM Journal on Control and Optimization* **58** (2020), no. 3, 1676–1699.
- [19] R. Carmona and F. Delarue, *Probabilistic Theory of Mean Field Games with Applications I*, Springer, NY, 2018.
- [20] R. Carmona and M. Laurière, *Convergence Analysis of Machine Learning Algorithms for the Numerical Solution of Mean Field Control and Games: II – The Finite Horizon Case*, *The Annals of Applied Probability* **32** (2022), no. 6, 4065–4105.
- [21] J.F. Chassagneux, D. Crisan, and F. Delarue, *A Probabilistic Approach to Classical Solutions of the Master Equation for Large Population Equilibria*, *Memoirs of the American Mathematical Society*, 2022.
- [22] ———, *Numerical Method for FBSDEs of McKean-Vlasov Type*, *The Annals of Applied Probability* **29** (2019), no. 3.
- [23] P. E. Chaudru de Raynal and N. Frikha, *Well-Posedness for Some Non-Linear Diffusion Processes and Related PDE on the Wasserstein Space*, *Journal de Mathématiques Pures et Appliquées* **159** (2022), 1–167.
- [24] L.P. Chaintron and A. Diez, *Propagation of Chaos: A Review of Models, Methods and Applications. I. Models and Methods*, *Kinetic and Related Models* **15** (2022), no. 6, 895–1015.
- [25] J. Comer, J.C. Gumbart, J. Hénin, T. Lelièvre, A. Pohorille, and C. Chipot, *The Adaptive Biasing Force Method: Everything You Always Wanted To Know but Were Afraid To Ask*, *The Journal of Physical Chemistry B* **118** (2015), no. 3, 1129–1151.
- [26] A. Cossio, F. Gozzi, I. Kharroubi, H. Pham, and M. Rosestolato, *Master Bellman Equation in the Wasserstein Space: Uniqueness of Viscosity Solutions*, *Transactions of the American Mathematical Society* (2023).
- [27] ———, *Optimal Control of Path-Dependent McKean-Vlasov SDEs in Infinite Dimension*, *The Annals of Applied Probability* **33** (2023), no. 4, 2863–2918.
- [28] D.A. Dawson, *Critical Dynamics and Fluctuations for a Mean-Field Model of Cooperative Behavior*, *Journal of Statistical Physics* **31** (1983), 29–85.
- [29] D. A. Dawson and J. Gärtner, *Large Deviations from the McKean-Vlasov Limit for Weakly Interacting Diffusions*, *Stochastics* **20** (1987), no. 4, 247–308.
- [30] P. Del Moral and A. Guionnet, *Large Deviations for Interacting Particle Systems: Applications to Non-Linear Filtering*, *Stochastic Processes and their Applications* **78** (1998), no. 1, 69–95.
- [31] F. Delarue, D. Lacker, and K. Ramanan, *From the Master Equation to Mean Field Game Limit Theory: Large Deviations and Concentration of Measure*, *The Annals of Probability* **48** (2020), no. 1, 211–263.
- [32] G. dos Reis, W. Salkeld, and J. Tugaut, *Freidlin-Wentzell LDP in Path Space for McKean-Vlasov Equations and the Functional Iterated Logarithm Law*, *The Annals of Probability* **29** (2019), no. 3, 1487–1540.
- [33] G. dos Reis, G. Smith, and P. Tankov, *Importance sampling for McKean-Vlasov SDEs*, *Applied Mathematics and Computation* **453** (2023).
- [34] P. Dupuis and R. S. Ellis, *A Weak Convergence Approach to the Theory of Large Deviations*, Wiley, NY, 1997.
- [35] P. Dupuis, K. Spiliopoulos, and X. Zhou, *Escaping from an Attractor: Importance Sampling and Rest Points I*, *The Annals of Applied Probability* **25** (2015), no. 5, 2909–2958.
- [36] P. Dupuis, K. Spiliopoulos, and H. Wang, *Importance Sampling for Multiscale Diffusions*, *Multiscale Modeling & Simulation* **10** (2012), no. 1, 1–27.
- [37] P. Dupuis and H. Wang, *Importance Sampling, Large Deviations, and Differential Games*, *Stochastics and Stochastic Reports* **76** (2004), no. 6, 481–508.
- [38] ———, *Subsolutions of an Isaacs Equation and Efficient Schemes for Importance Sampling*, *Mathematics of Operations Research* **32** (2007), no. 3, 723–757.
- [39] W. E, W. Ren, and E. Vanden-Eijnden, *Minimum Action Method for the Study of Rare Events*, *Communications on Pure and Applied Mathematics* **57** (2004), no. 5, 637–656.
- [40] L. Ebener, G. Margazoglou, J. Friedrich, L. Biferale, and R. Grauer, *Instanton Based Importance Sampling for Rare Events in Stochastic PDEs*, *Chaos: An Interdisciplinary Journal of Nonlinear Science* **29** (2019), no. 6.
- [41] J. Feng and M. Katsoulakis, *A Comparison Principle for Hamilton-Jacobi Equations Related to Controlled Gradient Flows in Infinite Dimensions*, *Archive for Rational Mechanics and Analysis* **192** (2009), 275–310.

- [42] J. Feng and G. Kurtz, *Large Deviations for Stochastic Processes*, American Mathematical Society, Providence, 2006.
- [43] J. Feng, T. Mikami, and J. Zimmer, *A Hamilton–Jacobi PDE Associated with Hydrodynamic Fluctuations from a Nonlinear Diffusion Equation*, *Communications in Mathematical Physics* **385** (2021), 1–54.
- [44] J. Feng and T. Nguyen, *Hamilton–Jacobi Equations in Space of Measures Associated with a System of Conservation Laws*, *Journal de Mathématiques Pures et Appliquées* **97** (2012), no. 4, 318–390.
- [45] M. Fischer, *On the Form of the Large Deviation Rate Function for the Empirical Measures of Weakly Interacting Systems*, *Bernoulli* **20** (2014), no. 4, 1765–1801.
- [46] W.H. Fleming and M.R. James, *Asymptotic Series and Exit Time Probabilities*, *The Annals of Probability* **20** (1992), no. 3, 1369–1384.
- [47] J.P. Fouque and Z. Zhang, *Deep Learning Methods for Mean Field Control Problems With Delay*, *Frontiers in Applied Mathematics and Statistics* **12** (2020), 1369–1384.
- [48] M.I. Freidlin and A.D. Wentzell, *Random Perturbations of Dynamical Systems*, Springer, Heidelberg, 2012.
- [49] J. Garnier, G. Papanicolaou, and T.W. Yang, *Large Deviations for a Mean Field Model of Systemic Risk*, *SIAM Journal of Financial Mathematics* **4** (2013), no. 1, 151–184.
- [50] ———, *Consensus Convergence with Stochastic Effects*, *Vietnam Journal of Mathematics* **45** (2017), 51–75.
- [51] I. Gasteratos, M. Salins, and K. Spiliopoulos, *Importance Sampling for Stochastic Reaction-Diffusion Equations in the Moderate Deviation Regime* (2022), available at [arXiv:2206.00646 \[math.PR\]](#).
- [52] N. Gavish, P. Nyquist, and M. Peletier, *Large Deviations and Gradient Flows for the Brownian One-Dimensional Hard-Rod System*, *Potential Analysis* **58** (2021), 71–121.
- [53] M. Geha, A. Jacquier, and Ž. Žurič, *Large and Moderate Deviations for Importance Sampling in the Heston Model* (2021).
- [54] M. Germain, M. Laurière, H. Pham, and X. Warin, *DeepSets and their Derivative Networks for Solving Symmetric PDEs*, *Journal of Scientific Computing* **91** (2022), no. 63.
- [55] M. Germain, J. Mikael, and X. Warin, *Numerical Resolution of McKean-Vlasov FBSDEs Using Neural Networks*, *Methodology and Computing in Applied Probability* **24** (2022), 2557–2586.
- [56] M. Germain, H. Pham, and X. Warin, *Rate of Convergence for Particle Approximation of PDEs in Wasserstein Space*, *Journal of Applied Probability* **59** (2022), no. 4, 992–1008.
- [57] 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*, *Communications on Pure and Applied Mathematics* (2023).
- [58] F.A. Grünbaum, *Propagation of Chaos for the Boltzmann Equation*, *Archive for Rational Mechanics and Analysis* **42** (1971), 323–345.
- [59] R. S. Gvalani and A. Schlichting, *Barriers of the McKean–Vlasov Energy via a Mountain Pass Theorem in the Space of Probability Measures*, *Journal of Functional Analysis* **279** (2020), no. 11, 437–461.
- [60] J. Han, R. Hu, and J. Long, *Learning High-Dimensional McKean-Vlasov Forward-Backward Stochastic Differential Equations with General Distribution Dependence* (2022), available at [arXiv:2204.11924 \[math.OC\]](#).
- [61] C. Hartmann, R. Banisch, M. Sarich, T. Badowski, and C. Schütte, *Characterization of Rare Events in Molecular Dynamics*, *Entropy* **16** (2014), no. 1, 350–376.
- [62] C. Hartmann and L. Richter, *Nonasymptotic Bounds for Suboptimal Importance Sampling* (2021), available at [arXiv:2102.09606 \[math.ST\]](#).
- [63] M. Kac, *Probability and Related Topics in Physical Sciences*, American Mathematical Society, Providence, 1957.
- [64] I. Karatzas and S. Shreve, *Brownian Motion and Stochastic Calculus*, Springer, NY, 1998.
- [65] D. Lacker, *Limit Theory for Controlled McKean-Vlasov Dynamics*, *SIAM Journal on Control and Optimization* **55** (2017), no. 3, 1641–1672.
- [66] M. Laurière, *Numerical Methods for Mean Field Games and Mean Field Type Control*, *Proceedings of Symposia in Applied Mathematics* **78** (2021), 221–282.
- [67] M. Laurière and O. Pironneau, *Dynamic Programming for Mean-Field Type Control*, *Journal of Optimization Theory and Applications* **169** (2016), 902–924.
- [68] W. Liu, Y. Song, J. Zhai, and T. Zhang, *Large and Moderate Deviation Principles for McKean-Vlasov SDEs with Jumps*, *Potential Analysis* **59** (2023), no. 3, 1141–1190.
- [69] M. R. Morse and K. Spiliopoulos, *Importance Sampling for Slow-Fast Diffusions Based on Moderate Deviations*, *SIAM Journal on Multiscale Modeling and Simulation* **18** (2020), no. 1, 315–350.
- [70] V. Nilsson and P. Nyquist, *A Note on Large Deviations for Interacting Particle Dynamics for Finding Mixed Equilibria in Zero-Sum Games* (2022), available at [arXiv:2206.15177 \[stat.ML\]](#).
- [71] 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*, *Partial Differential Equations and Applications* **2** (2021), no. 48.
- [72] C. Orrieri, *Large Deviations for Interacting Particle Systems: Joint Mean-Field and Small-Noise Limit*, *Electronic Journal of Probability* **25** (2020), 1–44.
- [73] H. Pham and X. Wei, *Bellman Equation and Viscosity Solutions for Mean-Field Stochastic Control Problem*, *ESAIM: Control, Optimisation and Calculus of Variations* **24** (2018), no. 1, 437–461.
- [74] M. Salins and K. Spiliopoulos, *Rare Event Simulation via Importance Sampling for Linear SPDE’s*, *Stochastics and Partial Differential Equations: Analysis and Computations* **5** (2017), no. 4, 652–690.
- [75] K. Spiliopoulos, *Nonasymptotic Performance Analysis of Importance Sampling Schemes for Small Noise Diffusions*, *Journal of Applied Probability* **53** (2015), no. 3, 797–810.
- [76] M. Talbi, N. Touzi, and J. Zhang, *Dynamic Programming Equation for the Mean Field Optimal Stopping Problem*, *SIAM Journal on Control and Optimization* **61** (2023), no. 4, 2140–2164.

- [77] S. Tong and G. Stadler, *Large Deviation Theory-Based Adaptive Importance Sampling for Rare Events in High Dimensions*, SIAM/ASA Journal on Uncertainty Quantification **11** (2023), no. 3, 788–813.
- [78] O. Valsson, P. Tiwary, and M. Parrinello, *Enhancing Important Fluctuations: Rare Events and Metadynamics from a Conceptual Viewpoint*, Annual Review of Physical Chemistry **67** (2016), 159–184.
- [79] E. Vanden-Eijnden and J. Weare, *Rare Event Simulation of Small Noise Diffusions*, Communications on Pure and Applied Mathematics **65** (2012), no. 12, 1770–1803.
- [80] J. Yong, *Linear-Quadratic Optimal Control Problems for Mean-Field Stochastic Differential Equations*, SIAM Journal on Control and Optimization **51** (2013), no. 4, 2809–2838.

Email address, Zachary William Bezemek: `zwb@duke.edu`

(Zachary William Bezemek) DUKE UNIVERSITY, DEPARTMENT OF MATHEMATICS, 120 SCIENCE DRIVE, DURHAM, NC 27708, USA

Email address, Max Heldman: `maxh@vt.edu`

(Max Heldman) VIRGINIA TECH, DEPARTMENT OF MATHEMATICS, 225 STANGER STREET, BLACKSBURG, VA 24061, USA