

Theoretical guarantees for sampling and inference in generative models with latent diffusions

Belinda Tzen
University of Illinois

BTZEN2@ILLINOIS.EDU

Maxim Raginsky
University of Illinois

MAXIM@ILLINOIS.EDU

Editors: Alina Beygelzimer and Daniel Hsu

Abstract

We introduce and study a class of probabilistic generative models, where the latent object is a finite-dimensional diffusion process on a finite time interval and the observed variable is drawn conditionally on the terminal point of the diffusion. We make the following contributions:

We provide a unified viewpoint on both sampling and variational inference in such generative models through the lens of stochastic control.

We quantify the expressiveness of diffusion-based generative models. Specifically, we show that one can efficiently sample from a wide class of terminal target distributions by choosing the drift of the latent diffusion from the class of multilayer feedforward neural nets, with the accuracy of sampling measured by the Kullback–Leibler divergence to the target distribution.

Finally, we present and analyze a scheme for unbiased simulation of generative models with latent diffusions and provide bounds on the variance of the resulting estimators. This scheme can be implemented as a deep generative model with a random number of layers.

1. Introduction and informal summary of results

Recently there has been much interest in using continuous-time processes to analyze discrete-time algorithms and probabilistic models (Wibisono et al., 2016; Li et al., 2017; Mandt et al., 2017; Chen et al., 2018; Yang et al., 2018). In particular, diffusion processes have been examined as a way towards a better understanding of first- and second-order optimization methods, as they afford an analysis of behavior over non-convex landscapes using a rich array of techniques from the mathematical physics literature (Li et al., 2017; Raginsky et al., 2017; Zhang et al., 2017). Gradient flows and diffusions have also found a role in the analysis of deep neural nets, where they are interpreted as describing the limiting case of infinitely many layers, with each layer being ‘infinitesimally thin’ (e.g., Chen et al. (2018); Li et al. (2018)). As in the case of optimization, continuous-time frameworks enable the use of a different set of tools for studying standard questions of relevance, such as sampling and inference, i.e., forward and backward passes through the network.

In this work, we consider a class of generative models where the latent object $X = \{X_t\}_{t \in [0,1]}$ is a d -dimensional diffusion and the observable object Y is a random element of some space \mathcal{Y} :

$$dX_t = b(X_t, t; \theta) dt + dW_t, \quad X_0 = x \tag{1.1a}$$

$$Y \sim q(\cdot | X_1) \tag{1.1b}$$

where (1.1a) is a d -dimensional Itô diffusion process whose drift $b(\cdot, \cdot; \theta)$ is a member of some parametric function class, such as multilayer feedforward neural nets, and (1.1b) prescribes an observation model for generating Y conditionally on X_1 . To the best of our knowledge, generative models of this form were first considered by Movellan et al. (2002) as a noisy continuous-time counterpart of recurrent neural nets. More recently, Hashimoto et al. (2016) and Ryder et al. (2018) investigated the use of discrete-time recurrent neural nets to approximate the population dynamics of biological systems that are classically modeled by diffusions. It is natural to view (1.1) as a continuum limit of deep generative models introduced by Rezende et al. (2014) — in fact, as we explain in Section 4, one can simulate a model of the above form using a deep generative model with a *random* number of layers. Alternatively, one can think of (1.1) as a *neural stochastic differential equation*, in analogy to the neural ODE framework of Chen et al. (2018).

There are three main questions that are natural to ask concerning the usefulness of such models: How expressive can they be? How might one sample from such a diffusion process? How might one perform inference on it? As our first contribution, we provide a unified view of sampling and inference through the lens of stochastic control. In particular, by adding a *control* u_t to the drift of some reference diffusion, one can obtain a desired distribution at $t = 1$, and the minimal-cost control that yields exact sampling is given by the so-called Föllmer drift (Föllmer, 1985; Dai Pra, 1991; Lehec, 2013; Eldan and Lee, 2018). Complementarily, we show that any control u_t added to the drift $b(\cdot, t; \theta)$ in (1.1a) leads to a variational upper bound on the log-likelihood of a given tuple of observations (y_1, \dots, y_n) . Variational inference then reduces to minimizing the expected control cost over a tractable class of controls. While we provide a unifying viewpoint that captures both sampling and inference, we emphasize that this is a synthesis of a number of existing results, and serves as a conceptual underpinning and motivation for our subsequent analysis. Specifically, after establishing that diffusion-based generative models can be effectively worked with, we explore their expressive power vis-à-vis neural nets: We show that, if the target density of X_1 can be efficiently approximated using a neural net, then the corresponding Föllmer drift can also be efficiently approximated by a neural net, such that the terminal law of the diffusion with this approximate drift is ε -close to the target density in Kullback–Leibler divergence. Finally, we investigate unbiased simulation methods for generative models with underlying diffusion processes and provide bounds on the variance of the resulting estimators.

1.1. Method of analysis: an overview

To arrive at the unified perspective of sampling and inference, we begin by formulating a stochastic control problem that captures all of our desiderata: sampling from a target probability law μ at terminal time $t = 1$; a set of tractable controls that might be used to take it there; and an appropriate notion of cost with that captures both the ‘control effort’ and the terminal cost that quantifies the discrepancy between the final probability law and the target measure μ .

Our first result, stated in Theorem 1, is an explicit characterization of the value function of this control problem, which has a free-energy interpretation and can be understood from an information-theoretic viewpoint: the Kullback–Leibler divergence between the law of the path of the uncontrolled diffusion and that of the path of the controlled diffusion is the expected total work done by the control. The negative free energy with respect to the uncontrolled process is a lower bound on that of the controlled process after accounting for the work done, and equality is achieved by the

optimal control. As pointed out above, this result is a synthesis of a number of existing results, and its main purpose is to motivate the use of controlled diffusions in probabilistic generative modeling.

We next examine the *expressiveness* of these generative models, which refers to their ability to generate samples from a given target distribution for X_1 when the observation model $q(\cdot|\cdot)$ in (1.1b) is fixed. In Theorem 4, we provide quantitative guarantees for obtaining approximate samples from a given target distribution μ for X_1 when the drift b in (1.1a) is restricted to be a multilayer feedforward neural net. Specifically, we show that, if the density f of μ with respect to the standard Gaussian measure on \mathbb{R}^d can be efficiently approximated by a feedforward neural net, then the corresponding Föllmer drift can also be approximated efficiently by a neural net. Moreover, this approximate Föllmer drift yields a diffusion $\{\widehat{X}_t\}$, such that $\widehat{\mu} = \text{Law}(\widehat{X}_1)$ satisfies $D(\mu||\widehat{\mu}) \leq \varepsilon$ for a given accuracy $\varepsilon > 0$. Under some assumptions on the smoothness of f and ∇f and on their uniform approximability by neural nets, the proof proceeds as follows: First, we show that the Föllmer drift can be approximated by a neural net uniformly over a given compact subset of \mathbb{R}^d and for all $t \in [0, 1]$. Then, to show that the terminal distribution resulting from this approximation is ε -close to μ in KL-divergence, we use Girsanov’s theorem to relate $D(\mu||\widehat{\mu})$ to the expected squared error between the Föllmer drift and its neural-net approximation.

Finally, we discuss the issue of unbiased simulation with the goal of estimating expected values of functions of X_1 . The standard Euler–Maruyama scheme (Graham and Talay, 2013, Chap. 7) is straightforward, but produces a biased estimator. Typically, one uses Monte Carlo sampling to reduce the variance; if the estimator is biased, then the variance will be reduced by a factor of $N^{1-\delta}$ for some $\delta \in (0, 1)$, instead of the optimal reduction by the factor of N , for N Monte Carlo runs. One way to obtain an *unbiased* estimator is to employ a *random* discretization of the time interval $[0, 1]$, where the sampling times are generated by a point process on the real line. Unbiased simulation schemes of this type have been proposed and analyzed by Bally and Kohatsu-Higa (2015), Andersson and Kohatsu-Higa (2017), and Henry-Labordère et al. (2017). Our final result, Theorem 6, builds on the latter work and presents an unbiased, finite-variance simulation scheme. Conceptually, the simulation scheme can be thought of as a deep latent Gaussian model in the sense of Rezende et al. (2014), but with a random number of layers. Unfortunately, the variance of the resulting estimator can exhibit exponential dependence on dimension. We show why this is the case via an analysis of the moment-generating function of the point process used to generate the random mesh and propose alternatives to reduce the variance.

1.2. Notation

The Euclidean norm of a vector $x \in \mathbb{R}^d$ will be denoted by $\|x\|$, the transpose of a vector or a matrix will be indicated by $(\cdot)^T$. The d -dimensional Euclidean ball of radius R centered at the origin will be denoted by $B^d(R)$. The standard Gaussian measure on \mathbb{R}^d will be denoted by γ_d . The Euclidean heat semigroup Q_t , $t \geq 0$, acts on measurable functions $f : \mathbb{R}^d \rightarrow \mathbb{R}$ as follows:

$$Q_t f(x) := \int_{\mathbb{R}^d} f(x + \sqrt{t}z) \gamma_d(dz) = \mathbf{E}[f(x + \sqrt{t}Z)], \quad Z \sim \gamma_d. \quad (1.2)$$

A function $g : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}$ is of class $C^{2,1}$ if it is twice continuously differentiable in the space variable $x \in \mathbb{R}^d$ and once continuously differentiable in the time variable $t \in [0, 1]$.

2. Exact sampling and variational inference: a unified stochastic control viewpoint

Before addressing the specific questions posed in the Introduction, we aim to demonstrate that both sampling and variational inference in generative models of the form (1.1) can be viewed through the lens of stochastic control. We give a brief description of the relevant ideas in Appendix A; the book by Fleming and Rishel (1975) is an excellent and readable reference.

2.1. A stochastic control problem

Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbf{P})$ be a probability space with a complete and right-continuous filtration $\{\mathcal{F}_t\}$, and let $W = \{W_t\}$ be a standard d -dimensional Brownian motion adapted to $\{\mathcal{F}_t\}$. Consider the Itô diffusion process

$$dX_t = b(X_t, t) dt + dW_t, \quad t \in [0, 1]; X_0 = x_0 \quad (2.1)$$

where the drift $b : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ is sufficiently well-behaved (say, bounded and Lipschitz). Then the process $\{X_t\}$ admits a *transition density*, i.e., a family of functions $p_{s,t} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ for all $0 \leq s < t \leq 1$, such that, for all points $x, y \in \mathbb{R}^d$ and all Borel sets $A \subset \mathbb{R}^d$,

$$\mathbf{P}[X_t \in A | X_s = x] = \int_A p_{s,t}(x, y) dy \quad (2.2)$$

(see, e.g., Protter (2005, Chap. V)).

Consider the following stochastic control problem: Let \mathcal{U} be the set of *controls*, i.e., measurable functions $u : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$. Any $u \in \mathcal{U}$ defines a diffusion process $X^u = \{X_t^u\}_{t \in [0,1]}$ by

$$dX_t^u = (b(X_t^u, t) + u(X_t^u, t)) dt + dW_t, \quad t \in [0, 1]; X_0^u = x_0. \quad (2.3)$$

We say that X^u is a diffusion controlled by u . Let a function $g : \mathbb{R}^d \rightarrow (0, \infty)$ be given. For each $u \in \mathcal{U}$, we define the family of *cost-to-go functions*

$$J^u(x, t) := \mathbf{E} \left[\frac{1}{2} \int_t^1 \|u_s\|^2 ds - \log g(X_1^u) \middle| X_t^u = x \right], \quad x \in \mathbb{R}^d, t \in [0, 1] \quad (2.4)$$

where u_s is shorthand for $u(X_s^u, s)$. The *value functions* $v : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}_+$ are defined by

$$v(x, t) := \inf_{u \in \mathcal{U}} J^u(x, t), \quad (2.5)$$

and we say that a control $u^* \in \mathcal{U}$ is *optimal* if $J^{u^*}(x, t) = v(x, t)$ for all x and t . The following theorem is, essentially, a synthesis of the results of Pavon (1989) and Dai Pra (1991):

Theorem 1 Consider the control problem (2.4). The value function v is given by

$$v(x, t) = -\log \mathbf{E}[g(X_1) | X_t = x], \quad (2.6)$$

where the conditional expectation is with respect to the uncontrolled diffusion process (2.1). Moreover, the optimal control u^* is given by $u^*(x, t) = -\nabla v(x, t)$, where the gradient is taken with respect to the space variable $x \in \mathbb{R}^d$, and the corresponding controlled diffusion $\{X_t^*\} = \{X_t^{u^*}\}$ has the transition density

$$p_{s,t}^*(x, y) = p_{s,t}(x, y) \exp(v(x, s) - v(y, t)), \quad (2.7)$$

where $p_{s,t}(\cdot)$ is the transition density (2.2) of the uncontrolled process.

This result, proved in Appendix A, also admits an information-theoretic interpretation. Let \mathbf{P}^0 denote the probability law of the path $X_{[0,1]}$ of the uncontrolled diffusion process (2.1) and let \mathbf{P}^u denote the corresponding object for the controlled diffusion (2.3). Since X and X^u differ from each other by a change of drift, the probability measures \mathbf{P}^u and \mathbf{P}^0 are mutually absolutely continuous, and the Radon–Nikodym derivative $d\mathbf{P}^u/d\mathbf{P}^0$ is given by the Girsanov formula (Protter, 2005)

$$\frac{d\mathbf{P}^u}{d\mathbf{P}^0} = \exp\left(-\int_0^1 u_t^T dW_t + \frac{1}{2}\int_0^1 \|u_t\|^2 dt\right), \quad (2.8)$$

where $u_t^T dW_t := \sum_{i=1}^d u_{i,t} dW_{i,t}$, with $u_{i,\cdot}$ and $dW_{i,\cdot}$ denoting the i th coordinates of u and W respectively. From (2.8), we can calculate the Kullback–Leibler divergence between \mathbf{P}^u and \mathbf{P}^0 :

$$D(\mathbf{P}^u\|\mathbf{P}^0) = \mathbf{E}_{\mathbf{P}^u}\left[\log\frac{d\mathbf{P}^u}{d\mathbf{P}^0}\right] = \mathbf{E}\left[\frac{1}{2}\int_0^1 \|u_t\|^2 dt\right]. \quad (2.9)$$

Therefore, by Theorem 1, for any control $u \in \mathcal{U}$, we can write

$$-\log \mathbf{E}[g(X_1)|X_0 = x] \leq D(\mathbf{P}^u\|\mathbf{P}^0) - \mathbf{E}[\log g(X_1^u)|X_0^u = x], \quad (2.10)$$

with equality if and only if $u = u^*$. An inequality of this form holds more generally for real-valued measurable functions of the entire path $X_{[0,1]}$ (Boué and Dupuis, 1998).

We will now demonstrate how both the problem of sampling and the problem of variational inference can be addressed via the above theorem.

2.2. Exact sampling: the Föllmer drift

Recall that, in the context of exact sampling, the objective is to construct a diffusion process $\{X_t\}_{t \in [0,1]}$, such that X_1 has a given target distribution μ . We will consider the case when μ is absolutely continuous with respect to the standard Gaussian measure γ_d and let f denote the Radon–Nikodym derivative $d\mu/d\gamma_d$. This problem goes back to a paper of Schrödinger (1931); for rigorous treatments, see, e.g., Jamison (1975), Föllmer (1985), Dai Pra (1991), Lehec (2013), Eldan and Lee (2018). The derivation we give below is not new (see, e.g., Dai Pra (1991, Thm. 3.1)), but the route we take is somewhat different in that we make the stochastic control aspect more explicit.

We take $b(x, t) \equiv 0$ and $X_0 = 0$ in (2.1). Then the diffusion process $\{X_t\}$ is simply the standard d -dimensional Brownian motion $\{W_t\}$, which has the Gaussian transition density

$$p_{s,t}(x, y) = \frac{1}{(2\pi(t-s))^{d/2}} \exp\left(-\frac{1}{2(t-s)}\|x-y\|^2\right).$$

Now consider the control problem (2.4) with $g = f$. By Theorem 1, the value function v is given by $v(x, t) = -\log \mathbf{E}[f(W_1)|W_t = x]$, and can be computed explicitly. For $0 \leq t < 1$, we have

$$\begin{aligned} e^{-v(x,t)} &= \mathbf{E}[f(W_1)|W_t = x] \\ &= \frac{1}{(2\pi(1-t))^{d/2}} \int_{\mathbb{R}^d} f(y) \exp\left(-\frac{1}{2(1-t)}\|x-y\|^2\right) dy \\ &= Q_{1-t}f(x), \end{aligned}$$

where Q denotes the Euclidean heat semigroup (1.2). Hence, $v(x, t) = -\log Q_{1-t}f(x)$, and the optimal diffusion process $\{X_t^*\}$ has the drift $u^*(x, t) = -\nabla v(x, t) = \nabla \log Q_{1-t}f(x)$. Following Lehec (2013) and Eldan and Lee (2018), we will refer to u^* as the *Föllmer drift* in the sequel.

It remains to show that $X_1^* \sim \mu$. Using the formula (2.7) for the transition density of X^* together with the fact that $e^{-v(y,1)} = f(y)$ and $e^{-v(0,0)} = \mathbf{E}[f(W_1)] = \int f d\gamma_d = 1$, we see that $p_{0,1}^*(0, y) dy = f(y)\gamma_d(dy)$. Then, for any Borel set $A \subseteq \mathbb{R}^d$,

$$\mathbf{P}[X_1^* \in A] = \int_A p_{0,1}^*(0, y) dy = \int_A f(y)\gamma_d(dy) = \mu(A).$$

Moreover, using the entropy inequality (2.10), we can show that the Föllmer drift is optimal in the following strong sense: Consider any control $u \in \mathcal{U}$ with $X_0^u = 0$ and with the property that $\text{Law}(X_1^u) = \mu$. For any such control,

$$\mathbf{E}[\log f(X_1^u) | X_0^u = 0] = \int_{\mathbb{R}^d} d\mu \log f = \int_{\mathbb{R}^d} d\mu \log \frac{d\mu}{d\gamma_d} = D(\mu \| \gamma_d),$$

while clearly $\log \mathbf{E}[f(W_1)] = 0$. Therefore, it follows from (2.10) that, for any such control u ,

$$D(\mathbf{P}^u \| \mathbf{P}^0) = \frac{1}{2} \mathbf{E} \left[\int_0^1 \|u_t\|^2 dt \right] \geq D(\mu \| \gamma_d),$$

with equality if and only if $u = u^*$. Thus, the Föllmer drift has the minimal ‘energy’ among all admissible controls that induce the distribution μ at $t = 1$, and this energy is precisely the Kullback–Leibler divergence between μ and the standard Gaussian measure γ_d (Dai Pra, 1991; Lehec, 2013; Eldan and Lee, 2018).

2.3. Variational inference

We now turn to the problem of variational inference. We are given an n -tuple of observations $\mathbf{y} = (y_1, \dots, y_n) \in \mathcal{Y}^n$, and wish to upper-bound the negative log-likelihood

$$L_n(\mathbf{y}; \theta) := \frac{1}{n} \sum_{i=1}^n L(y_i; \theta),$$

where $L(y; \theta) := -\log \mathbf{E}[q(y|X_1)]$ and $\{X_t\}$ is the diffusion process (1.1).

We take $b = b(\cdot, \cdot; \theta)$ in (2.1) and consider the control problem (2.4) with $g(x) = q(y|x)$ for some fixed $y \in \mathcal{Y}$. Then, by Theorem 1, any control $u \in \mathcal{U}$ gives rise to an upper bound on $L(y; \theta)$:

$$L(y; \theta) \leq \mathbf{E} \left[\frac{1}{2} \int_0^1 \|u_t\|^2 dt - \log q(y|X_1^u) \middle| X_0^u = x \right] =: F^u(y; \theta),$$

where the quantity on the right-hand side can be thought of as the *variational free energy* that depends on the choice of the control u , and equality is achieved when $u = u^*$. While the structure of the optimal control u^* is described in Theorem 1, it may not be possible to derive it in closed form. However, we can fix a class $\tilde{\mathcal{U}} \subset \mathcal{U}$ of tractable suboptimal controls and upper-bound $L(y; \theta)$ by $\inf_{u \in \tilde{\mathcal{U}}} F^u(y; \theta)$. For example, we can take $\tilde{\mathcal{U}}$ to consist of all controls of the form $u(x, t) =$

$\phi - b(x, t; \theta)$ for some $\phi \in \mathbb{R}^d$. In that case, X_t^u is the sum of the Brownian motion W_t and the affine drift $x + t\phi$, and consequently

$$F^u(y; \theta) = \mathbf{E} \left[\frac{1}{2} \int_0^1 \|\phi - b(x + t\phi + W_t; \theta)\|^2 dt - \log q(y|x + \phi + W_1) \right],$$

where the expectation is taken with respect to the standard Brownian motion W . Another possibility is to consider controls of the form $u(x, t) = Ax - b(x, t; \theta)$, for some $A \in \mathbb{R}^{d \times d}$. The corresponding controlled diffusion is the Ornstein–Uhlenbeck process $X_t^u = e^{At}x + \int_0^t e^{A(t-s)} dW_s$, and the variational free energy can be minimized over $A \in \mathbb{R}^{d \times d}$.

3. Expressiveness

Now that we have shown that generative models of the form (1.1) allow for both sampling and variational inference, we turn to the analysis of their *expressiveness*. Specifically, our objective is to show that, by working with a suitable structured class of drifts $b(\cdot, \cdot; \theta)$, we can achieve approximate sampling from a rich class of distributions at the terminal time $t = 1$.

Let μ be the target probability measure for X_1 . We assume that μ is absolutely continuous with respect to γ_d and let f denote the Radon–Nikodym derivative $d\mu/d\gamma_d$. From Section 2.2 we know that the diffusion process governed by the Itô SDE

$$dX_t = b(X_t, t) dt + dW_t, \quad X_0 = 0 \tag{3.1}$$

with the Föllmer drift $b(x, t) = \nabla \log Q_{1-t} f(x)$ has the property that $\mu = \text{Law}(X_1)$, and, moreover, it is optimal in the sense that it minimizes the ‘energy’ $\frac{1}{2} \int_0^1 \mathbf{E} \|u_t\|^2 dt$ among all adapted drifts $\{u_t\}$ that result in distribution μ at time $t = 1$. The main result of this section is as follows: If the Radon–Nikodym derivative f can be approximated efficiently by multilayer feedforward neural nets, then, for any $\varepsilon > 0$, there exists a drift $b(x, t) = \widehat{b}(x, t; \theta)$ that can be implemented *exactly* by a neural net whose parameters θ do not depend on time or space, and the terminal law $\widehat{\mu} := \text{Law}(\widehat{X}_1)$ of the diffusion process

$$d\widehat{X}_t = \widehat{b}(\widehat{X}_t, t) dt + dW_t, \quad \widehat{X}_0 = 0 \tag{3.2}$$

is an ε -approximation to μ in the KL-divergence: $D(\mu|\widehat{\mu}) \leq \varepsilon$. Moreover, the size of the neural net that implements the approximate Föllmer drift \widehat{b} can be estimated explicitly in terms of the size of a suitable approximating neural net for f .

We begin by imposing some assumptions on f . The first assumption is needed to guarantee enough regularity for the Föllmer drift:

Assumption 1 *The function f is differentiable, both f and ∇f are L -Lipschitz, and there exists a constant $c \in (0, 1]$, such that $f \geq c$ everywhere.*

This assumption is satisfied, for example, by Gibbs measures of the form $\mu(dx) = Z^{-1} e^{-\frac{1}{2}\|x\|^2 - F(x)} dx$ with a differentiable potential $F : \mathbb{R}^d \rightarrow \mathbb{R}_+$, such that both F and ∇F are Lipschitz, and F is bounded from above; see Appendix B for details.

Next, we introduce the assumptions pertaining to the approximability of f by neural nets. Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be a fixed nonlinearity. Given a vector $w \in \mathbb{R}^n$ and scalars α, β , define the function

$$N_{w,\alpha,\beta}^\sigma : \mathbb{R}^n \rightarrow \mathbb{R}, \quad N_{w,\alpha,\beta}^\sigma(x) := \alpha \cdot \sigma(w^T x + \beta).$$

For $\ell \geq 2$, we define the class \mathcal{N}_ℓ^σ of ℓ -layer feedforward neural nets with activation function σ recursively as follows: \mathcal{N}_2^σ consists of all functions of the form $x \mapsto \sum_{i=1}^m N_{w_i,\alpha_i,\beta_i}^\sigma(x)$ for all $m \in \mathbb{N}$, $w_1, \dots, w_m \in \mathbb{R}^d$, $\alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_m \in \mathbb{R}$, and, for each $\ell \geq 2$,

$$\mathcal{N}_{\ell+1}^\sigma := \bigcup_{k \geq 1} \bigcup_{m \geq 1} \left\{ x \mapsto \sum_{i=1}^m N_{w_i,\alpha_i,\beta_i}^\sigma(h_1(x), \dots, h_k(x)) : \right. \\ \left. \alpha_1, \dots, \alpha_m, \beta_1, \dots, \beta_m \in \mathbb{R}, w_1, \dots, w_m \in \mathbb{R}^k, h_1, \dots, h_k \in \mathcal{N}_\ell^\sigma \right\}.$$

Thus, each element of \mathcal{N}_ℓ^σ is a function that represents computation by a directed acyclic graph, where each node receives inputs u_1, \dots, u_k , performs a computation of the form $(u_1, \dots, u_k) \mapsto \sigma(w_1 u_1 + \dots + w_k u_k + \beta)$, and communicates the outcome of the computation to all the nodes in the next layer. We refer to ℓ as the *depth* of the neural net, and define the *size* of the neural net as the total number of nodes in its computation graph. We will denote by $\mathcal{N}_{\ell,s}^\sigma$ the collection of all neural nets with depth ℓ and size s . All these definitions extend straightforwardly to the case of neural nets with vector-valued output and to the case where each node may have a different activation function.

We assume that the activation function σ is differentiable and *universal*, in the sense that any univariate Lipschitz function which is nonconstant on a bounded interval can be approximated arbitrarily well by an element of \mathcal{N}_2^σ :

Assumption 2 *The activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is differentiable. Moreover, there exists a constant $c_\sigma > 0$ depending only on σ , such that the following holds: For any L -Lipschitz function $h : \mathbb{R} \rightarrow \mathbb{R}$ which is constant outside the interval $[-R, R]$ and for any $\delta > 0$, there exist real numbers $a, \{(\alpha_i, \beta_i, \gamma_i)\}_{i=1}^m$, where $m \leq c_\sigma \frac{RL}{\delta}$, such that the function*

$$\tilde{h}(x) = a + \sum_{i=1}^m \alpha_i \sigma(\beta_i x + \gamma_i) \tag{3.3}$$

satisfies $\sup_{x \in \mathbb{R}} |\tilde{h}(x) - h(x)| \leq \delta$.

Remark 2 Apart from differentiability, this is the same assumption made by [Eldan and Shamir \(2016\)](#). For example, it holds for differentiable sigmoidal activation functions, i.e., monotonic functions that satisfy $\lim_{u \rightarrow -\infty} \sigma(u) = a$ and $\lim_{u \rightarrow +\infty} \sigma(u) = b$ for some $a \neq b$. The popular rectified linear unit (or ReLU) activation function $u \mapsto u \vee 0$ is universal in the above sense but not differentiable. However, we can replace it by the differentiable softplus function $u \mapsto \log(1 + e^{cu})$, where increasing the value of $c > 0$ results in finer approximations to the ReLU. Also, note that the function \tilde{h} differs from the elements of \mathcal{N}_2^σ by the presence of the constant term a . However, the constant function $x \mapsto a$ can be implemented by $N_{0,a/\sigma(z),z}^\sigma$, for any $z \in \mathbb{R}$ such that $\sigma(z) \neq 0$. Thus, we will refer to functions of the form (3.3) as 2-layer neural networks of size $m + 1$.

We also make the following assumption regarding approximability of f by neural nets:

Assumption 3 For any $R > 0$ and $\varepsilon > 0$, there exists a neural net $\widehat{f} \in \mathcal{N}_{\ell,s}^\sigma$ with $\ell, s \leq \text{poly}(1/\varepsilon, d, L, R)$, such that

$$\sup_{x \in \mathbb{B}^d(R)} |f(x) - \widehat{f}(x)| \leq \varepsilon \quad \text{and} \quad \sup_{x \in \mathbb{B}^d(R)} \|\nabla f(x) - \nabla \widehat{f}(x)\| \leq \varepsilon. \quad (3.4)$$

Remark 3 Typical results on neural net approximation are concerned with approximating a given function uniformly on a given compact set. By contrast, Assumption 3 requires uniform approximability of both f and its gradient ∇f on a compact set by some neural net \widehat{f} and its gradient $\nabla \widehat{f}$. Such simultaneous approximation guarantees can also be found in the literature, see, e.g., [Hornik et al. \(1990\)](#); [Yukich et al. \(1995\)](#); [Li \(1996\)](#). See [Safran and Shamir \(2017\)](#) for a discussion of various trade-offs between depth and width (maximum number of neurons per layer) in neural net approximation.

We are now in a position to state the main result of this section:

Theorem 4 Suppose Assumptions 1–3 are in force. Let L denote the maximum of the Lipschitz constants of f and ∇f . Then, for any $0 < \varepsilon < 16L^2/c^2$, there exists a neural net $\widehat{v} : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ with size polynomial in $1/\varepsilon, d, L, c, 1/c$, such that the activation function of each neuron is an element of the set $\{\sigma, \sigma', \text{ReLU}\}$, and the following holds: If $\{\widehat{X}_t\}_{t \in [0,1]}$ is the diffusion process governed by the Itô SDE

$$d\widehat{X}_t = \widehat{b}(\widehat{X}_t, t) dt + dW_t, \quad \widehat{X}_0 = 0 \quad (3.5)$$

with the drift $\widehat{b}(x, t) = \widehat{v}(x, \sqrt{1-t})$, then $\widehat{\mu} := \text{Law}(\widehat{X}_1)$ satisfies $D(\mu \| \widehat{\mu}) \leq \varepsilon$.

3.1. The proof of Theorem 4

The proof relies on three key steps: First, we show that the heat semigroup $Q_t f(x)$ can be approximated by a finite sum of the form $\frac{1}{N} \sum_{n \leq N} f(x + \sqrt{t} z_n)$ uniformly for all $x \in \mathbb{B}^d(R)$ and all $t \in [0, 1]$, where $z_1, \dots, z_N \in \mathbb{R}^d$ lie in a ball of radius $\mathcal{O}(\sqrt{d \log N})$. This result is stated in [Appendix C](#) and proved using empirical process methods. Next, replacing f with a suitable neural net approximation \widehat{f} , we build on this result to show that the Föllmer drift $\nabla \log Q_{1-t} f(x)$ can be approximated by a neural net using σ, σ' , and ReLU as activation functions. This is the content of [Theorem 5](#) below (the proof is given in [Appendix D](#)). The third step uses Girsanov theory to upper-bound the approximation error that results from replacing the Föllmer drift by this neural net.

Theorem 5 Let $0 < \varepsilon < 4L/c$ and $R > 0$ be given. Then there exists a neural net $\widehat{v} : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ of size polynomial in $1/\varepsilon, d, L, R, c, 1/c$, such that the activation function of each neuron is an element of the set $\{\sigma, \sigma', \text{ReLU}\}$, and the following holds:

$$\sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \left\| \widehat{v}(x, \sqrt{t}) - \nabla \log Q_t f(x) \right\| \leq \varepsilon$$

and

$$\max_{i \in [d]} \sup_{x \in \mathbb{R}^d} \sup_{t \in [0,1]} |\widehat{v}_i(x, \sqrt{t})| \leq \frac{2L}{c}.$$

We now complete the proof of Theorem 4. For any $R > 0$, Theorem 5 guarantees the existence of a neural net $\widehat{v} : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ that satisfies

$$\sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0, 1]} \left\| \widehat{v}(x, \sqrt{t}) - \nabla \log Q_t f(x) \right\| \leq \sqrt{\varepsilon} \quad (3.6)$$

and

$$\max_{i \in [d]} \sup_{x \in \mathbb{R}^d} \sup_{t \in [0, 1]} |\widehat{v}_i(x, \sqrt{t})| \leq \frac{2L}{c}. \quad (3.7)$$

Let $\mu := \text{Law}(X_{[0, 1]})$ and $\widehat{\mu} := \text{Law}(\widehat{X}_{[0, 1]})$. The Girsanov formula gives

$$D(\mu \| \widehat{\mu}) = \frac{1}{2} \int_0^1 \mathbf{E} \|b(X_t, t) - \widehat{b}(X_t, t)\|^2 dt,$$

where the interchange of the integral and the expectation follows from Fubini's theorem because both b and \widehat{b} are bounded by Lemma 9 in Appendix B and (3.7). We now proceed to estimate the integrand. For each $t \in [0, 1]$,

$$\begin{aligned} & \mathbf{E} \|b(X_t, t) - \widehat{b}(X_t, t)\|^2 \\ &= \mathbf{E} \left[\|b(X_t, t) - \widehat{b}(X_t, t)\|^2 \cdot \mathbf{1}\{X_t \in \mathbb{B}^d(R)\} \right] + \mathbf{E} \left[\|b(X_t, t) - \widehat{b}(X_t, t)\|^2 \cdot \mathbf{1}\{X_t \notin \mathbb{B}^d(R)\} \right] \\ &=: T_1 + T_2, \end{aligned}$$

where $T_1 \leq \varepsilon$ by (3.6). To estimate T_2 , we first observe that, since the Föllmer drift is bounded in norm by L/c by Lemma 9, we have

$$\mathbf{P} \left\{ \sup_{t \in [0, 1]} \|X_t\| \geq R \right\} \leq \frac{\sqrt{d} + L/c}{R}$$

(Bubeck et al., 2018, Lemma 3.8). Therefore,

$$T_2 \leq \frac{9dL^2}{c^2} \cdot \frac{\sqrt{d} + L/c}{R}.$$

Choosing R large enough to guarantee $T_2 \leq \varepsilon$ and putting everything together, we obtain $D(\mu \| \widehat{\mu}) \leq \varepsilon$. Therefore, $D(\mu \| \widehat{\mu}) \leq D(\mu \| \widehat{\mu}) \leq \varepsilon$ by the data processing inequality.

4. Unbiased simulation

Now that we have shown that generative models with latent diffusions are capable of expressing a rich class of probability distributions, we turn to the problem of unbiased simulation. Specifically, given a function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, we wish to estimate the expectation $\mathbf{E}[g(X_1) | X_0 = x]$, where $X = \{X_t\}_{t \in [0, 1]}$ with $X_0 = x$ is a diffusion process of the form (1.1). The simplest approach is to use the Euler–Maruyama scheme: Fix a partition $0 = t_0 < t_1 < \dots < t_n < t_{n+1} = 1$ of $[0, 1]$ and define the Itô process $\{\tilde{X}_t\}_{t \in [0, 1]}$ by $\tilde{X}_0 = x$ and

$$\tilde{X}_t = \tilde{X}_{t_i} + \int_{t_i}^t b(\tilde{X}_s, s; \theta) ds + \int_{t_i}^t dW_s, \quad t \in (t_i, t_{i+1}], i = 0, \dots, n. \quad (4.1)$$

In particular, for each $1 \leq i \leq n + 1$,

$$\tilde{X}_{t_i} = \tilde{X}_{t_{i-1}} + b(\tilde{X}_{t_{i-1}}, t_{i-1}; \theta)(t_i - t_{i-1}) + W_{t_i} - W_{t_{i-1}}.$$

We can then estimate the expectation $\mathbf{E}[g(X_1)]$ by $g(\tilde{X}_{t_{n+1}}) \equiv g(\tilde{X}_1)$, but this estimate is biased: if g is, say, bounded, then

$$|\mathbf{E}[g(X_1)] - \mathbf{E}[g(\tilde{X}_1)]| \leq C_g(x) \cdot \max_{0 \leq i \leq n} (t_{i+1} - t_i),$$

where $C_g(x) > 0$ is some constant that depends on g and on the starting point x (Graham and Talay, 2013). Recently, several authors (Bally and Kohatsu-Higa, 2015; Andersson and Kohatsu-Higa, 2017; Henry-Labordère et al., 2017) have studied unbiased simulation of SDEs using Euler–Maruyama schemes with *random* partitions, where the partition breakpoints are generated by a Poisson point process on the real line. In this section, we build on this line of work and present a scheme for unbiased simulation in the context of generative models of the form (1.1) that uses random partitions generated by arbitrary renewal processes (Kallenberg, 2002, Chap. 9) with sufficiently well-behaved densities of interrenewal times. Our analysis closely follows that of Henry-Labordère et al. (2017), but we provide a more refined analysis of the variance of the resulting estimators.

We first describe the simulation procedure. In what follows, we will drop the index θ from the drift to keep the notation clean. Let τ_1, τ_2, \dots be i.i.d. nonnegative random variables with an absolutely continuous distribution whose support contains the interval $[0, 1 + \varepsilon]$ for some $\varepsilon > 0$. Let F_τ and f_τ denote the cdf and the pdf of τ_1 . Let $T_0 = 0$ and

$$T_k := \left(\sum_{i=1}^k \tau_i \right) \wedge 1, \quad k \geq 1 \quad \text{and} \quad N := \max\{k : T_k < 1\}.$$

Define a process $\hat{X} = \{\hat{X}_t\}_{t \in [0,1]}$ with $\hat{X}_0 = x$ as the Euler–Maruyama scheme (4.1) on the random partition $0 = T_0 < T_1 < \dots < T_N < T_{N+1} \equiv 1$ of $[0, 1]$, and let

$$\hat{\psi} := \frac{1}{1 - F_\tau(1 - T_N)} \cdot \left(g(\hat{X}_1) - g(\hat{X}_{T_N}) \mathbf{1}_{\{N > 0\}} \right) \cdot \prod_{k=1}^N \frac{1}{f_\tau(T_k - T_{k-1})} \hat{W}_k, \quad (4.2)$$

where

$$\hat{W}_k := \frac{(b(\hat{X}_{T_k}, T_k) - b(\hat{X}_{T_{k-1}}, T_{k-1}))^T (W_{T_{k+1}} - W_{T_k})}{T_{k+1} - T_k}.$$

This process can be interpreted as a deep generative model in the sense of Rezende et al. (2014), but with a *random* number of layers. Specifically, let $\xi_1, \xi_2, \dots \stackrel{\text{i.i.d.}}{\sim} \gamma_d$ be independent of $\{\tau_i\}$, and define $\hat{X}^{(0)}, \hat{X}^{(1)}, \dots, \hat{X}^{(N+1)}$ recursively by taking $\hat{X}^{(0)} = x$ and

$$\hat{X}^{(k+1)} = \hat{X}^{(k)} + b(\hat{X}^{(k)}, T_k) \cdot (T_{k+1} - T_k) + (T_{k+1} - T_k)^{1/2} \xi_{k+1}, \quad k = 0, 1, \dots, N.$$

Then

$$\hat{\psi} \stackrel{d}{=} g(\hat{X}^{(N+1)}) \cdot \frac{1}{1 - F_\tau(1 - T_N)} \cdot \prod_{k=1}^N \frac{(b(\hat{X}^{(k)}, T_k) - b(\hat{X}^{(k-1)}, T_{k-1}))^T \xi_{k+1}}{f_\tau(T_k - T_{k-1}) \cdot (T_{k+1} - T_k)^{1/2}},$$

where $\stackrel{d}{=}$ denotes equality of probability distributions. We are now ready to state our main result on unbiased simulation (see Appendix E for the proof):

Theorem 6 *Suppose that the drift $b(x, t)$ is uniformly bounded, Lipschitz in x , and $\frac{1}{2}$ -Hölder in t , i.e., for some constants $b_\infty > 0$ and $L_b > 0$,*

$$\|b(x, t)\| \leq b_\infty \quad \text{and} \quad \|b(x, s) - b(y, t)\| \leq L_b \left(\|x - y\| + |s - t|^{1/2} \right). \quad (4.3)$$

for all $x, y \in \mathbb{R}^d$ and all $s, t \in [0, 1]$. Suppose also that

$$\frac{1}{f_\tau(s)} \leq C e^{as}, \quad s \in (0, 1) \quad (4.4)$$

for some constants $C > 0$ and $a \geq 0$. Then, for any Lipschitz-continuous $g : \mathbb{R}^d \rightarrow \mathbb{R}$ with Lipschitz constant L_g , $\hat{\psi}$ is an unbiased estimator of $\mathbf{E}[g(X_1)|X_0 = x]$ with

$$\text{Var}[\hat{\psi}] \leq \left(\frac{e^a}{1 - F_\tau(1)} \right)^2 K M_N(\kappa), \quad (4.5)$$

where $K = \text{poly}(|g(x)|, L_b, L_g, b_\infty, d)$, $\kappa = \log \text{poly}(C, L_b, L_g, b_\infty, d)$, and $M_N(\theta) := \mathbf{E}[\exp(\theta N)]$ is the moment-generating function of N .

For example, the type of drift used in the construction of Section 3 has the property (4.3). The key implication of Theorem 6 is that the variance of the estimator $\hat{\psi}$ is controlled by the moment-generating function of N , and is therefore related to the tail behavior of the sums $S_k := \sum_{i=1}^k \tau_i$. In some cases, one can calculate M_N in closed form. For instance, if we take $\tau_1, \tau_2, \dots \stackrel{\text{i.i.d.}}{\sim} \text{Exp}(\lambda)$ for some $\lambda > 0$, then the estimator (4.2) reduces to the one introduced by Henry-Labordère et al. (2017). Since $F_\tau(s) = 1 - e^{-\lambda s}$ and $f_\tau(s) = \lambda e^{-\lambda s}$ for $s \geq 0$, (4.4) holds with $C = 1/\lambda$ and $a = \lambda$; moreover, $N \sim \text{Pois}(\lambda)$ with

$$M_N(\theta) = \exp(\lambda(e^\theta - 1)).$$

Thus, $\text{Var}[\hat{\psi}]$ grows like $\exp(d^2)$, as already observed by Henry-Labordère et al. (2017). One way to reduce the variance is to choose the τ_i 's with lighter tails. To see this, we need estimates of Λ_N ; the following lemma provides a computable upper bound:

Lemma 7 *Let M_τ denote the moment-generating function of τ . Then*

$$M_N(\theta) \leq 1 + e^\theta \inf_{\beta > 0} \left\{ (\beta + 1)e^{\theta\beta} + \sum_{k=0}^{\infty} (e^{\theta+1} M_\tau(-\beta))^k \right\}. \quad (4.6)$$

As an example, suppose τ_1, τ_2, \dots are i.i.d. samples from the uniform distribution on $[0, T]$ for some $T > 1$. Then

$$M_\tau(-\beta) = \frac{1}{\beta T} (1 - e^{-\beta T}),$$

and it is a matter of straightforward but lengthy algebra to show that $M_\tau(-\beta) \leq e^{-2(\theta+1)}$ for all β satisfying

$$e^{-\beta T} \geq 2 \left(1 + e^{-2(\theta+1)} \log 2 - (2\theta + 3)e^{-2(\theta+1)} \right).$$

Using this in (4.6) yields the estimate $M_N(\theta) \lesssim e^{\text{poly}(\theta)}$. The density of a $\text{Uniform}(0, T)$ random variable clearly satisfies (4.4). Thus, applying Theorem 6 to the estimator (4.2) with $\tau_i \stackrel{\text{i.i.d.}}{\sim} \text{Uniform}(0, T)$, we see that its variance scales *quasipolynomially* in d , i.e., $\text{Var}[\hat{\psi}] \lesssim e^{\text{polylog}(d)}$. However, choosing τ_i 's with lighter tails will generally lead to larger values of N , i.e., a deeper generative model will be needed.

Appendix A. The proof of Theorem 1

We first need some background on controlled diffusion processes, see, e.g., [Fleming and Rishel \(1975\)](#). As in Section 2, let \mathcal{U} be the set of controls, where each $u \in \mathcal{U}$ defines a controlled diffusion governed by the Itô SDE

$$dX_t^u = (b(X_t^u, t) + u(X_t^u, t)) dt + dW_t, \quad t \in [0, 1]; \quad X_0^u = x_0.$$

Let $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ and $\tilde{c} : \mathbb{R}^d \rightarrow \mathbb{R}_+$ be given. For each $u \in \mathcal{U}$, we define the cost-to-go functions

$$J^u(x, t) := \mathbf{E} \left[\int_t^1 c(X_s^u, u_s) ds + \tilde{c}(X_1^u) \middle| X_t^u = x \right], \quad x \in \mathbb{R}^d, t \in [0, 1] \quad (\text{A.1})$$

where u_s is shorthand for $u(X_s^u, s)$. The value functions $v : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}_+$ are defined in (2.5). In general, finding an optimal control is difficult. However, a sufficient condition for optimality is given by the so-called *verification theorem* from the theory of controlled diffusions (see, e.g., Chap. VI of [Fleming and Rishel \(1975\)](#)): Suppose that there exists a function $v \in C^{2,1}(\mathbb{R}^d \times [0, 1])$ that solves the Cauchy problem

$$\frac{\partial v(x, t)}{\partial t} + \mathcal{L}_t v(x, t) = - \min_{\alpha \in \mathbb{R}^d} \{ \alpha^T \nabla v(x, t) + c(x, \alpha) \} \text{ on } \mathbb{R}^d \times [0, 1]; \quad g(\cdot, 1) = \tilde{c}(\cdot) \quad (\text{A.2})$$

where \mathcal{L}_t is the (time-varying) generator of the diffusion (2.1):

$$\mathcal{L}_t h(x, t) := b(x, t)^T \nabla h(x, t) + \frac{1}{2} \text{tr} \nabla^2 h(x, t) \quad (\text{A.3})$$

for any $h \in C^{2,1}(\mathbb{R}^d \times [0, 1])$, and where the gradient and the Hessian are taken with respect to the ‘space variable’ $x \in \mathbb{R}^d$. Then v is the value function for (A.1), and the optimal control u^* is given by

$$u^*(x, t) = \arg \min_{\alpha \in \mathbb{R}^d} \{ \alpha^T \nabla v(x, t) + c(x, \alpha) \}. \quad (\text{A.4})$$

The PDE (A.2) is called the *Bellman equation* associated to the control problem (A.1).

Remark 8 In fact, the control (A.4) is optimal among a much wider class of *adapted controls*, i.e., all stochastic processes $\{u_t\}_{t \in [0, 1]}$ adapted to the filtration $\{\mathcal{F}_t\}$. The class \mathcal{U} defined above consists of so-called *Markov controls*, where u_t is a deterministic function of X_t^u and t . In that case, the controlled diffusion X^u is a Markov process.

We now turn to the proof of Theorem 1. The first step is to use the logarithmic transformation due to [Fleming \(1978\)](#); see also [Fleming and Sheu \(1985\)](#); [Sheu \(1991\)](#). Consider the function $h(x, t) := \mathbf{E}[g(X_1) | X_t = x]$. By the Feynman–Kac formula ([Kallenberg, 2002](#), Thm. 24.1), this function is a $C^{2,1}$ solution of the Cauchy problem

$$\frac{\partial h}{\partial t} + \mathcal{L}_t h = 0 \text{ on } \mathbb{R}^d \times [0, 1]; \quad h(\cdot, 1) = g(\cdot). \quad (\text{A.5})$$

It is a matter of simple calculus to verify that $v(x, t) = -\log h(x, t)$ solves the Cauchy problem

$$\frac{\partial v}{\partial t} + \mathcal{L}_t v = \frac{1}{2} \|\nabla v\|^2 \text{ on } \mathbb{R}^d \times [0, 1]; \quad v(\cdot, 1) = -\log g(\cdot). \quad (\text{A.6})$$

Moreover, using the variational representation

$$\frac{1}{2} \|\nabla v\|^2 = - \min_{\alpha \in \mathbb{R}^d} \left\{ \alpha^T \nabla v + \frac{1}{2} \|\alpha\|^2 \right\},$$

where the optimizer is given by $\alpha^* = -\nabla v$, it is readily verified that (A.6) is the Bellman equation (A.2) associated to the control problem (2.4). Hence, by the verification theorem, $v(x, t) = -\log h(x, t)$ is the value function we seek, and the optimal control is given by $u^*(x, t) = -\nabla v(x, t)$.

Now consider the diffusion process

$$dX_t^* = (b(X_t^*, t) + \nabla \log h(X_t^*, t)) dt + dW_t,$$

which satisfies

$$\begin{aligned} -\log \mathbf{E}[g(X_1) | X_t = x] &= \mathbf{E} \left[\frac{1}{2} \int_T^1 \|\nabla \log h(X_s^*, s)\|^2 ds - \log g(X_1^*) \middle| X_t^* = x \right] \\ &= \min_{u \in \mathcal{U}} \mathbf{E} \left[\frac{1}{2} \int_t^1 \|u_s\|^2 ds - \log g(X_1^u) \middle| X_t^u = x \right]. \end{aligned}$$

Since h solves (A.5), the transition density of $\{X_t^*\}$ is given by (2.7) by a result of Jamison (1975) and Dai Pra (1991).

Appendix B. Regularity properties of f and the Föllmer drift

We first show that Assumption 1 holds for Gibbs measures

$$\mu(dx) = Z^{-1} e^{-\frac{1}{2}\|x\|^2 - F(x)} dx$$

with sufficiently well-behaved potentials F . Suppose that $F : \mathbb{R}^d \rightarrow \mathbb{R}_+$ is differentiable, and both F and ∇F are L -Lipschitz. Then $f = d\mu/d\gamma_d = \text{const} \cdot e^{-F}$, and the Lipschitz continuity of f follows from the Lipschitz continuity of $u \mapsto e^{-u}$ on $[0, \infty)$:

$$|e^{-F(x)} - e^{-F(y)}| \leq |F(x) - F(y)| \leq L\|x - y\|.$$

Likewise, the Lipschitz continuity of ∇f follows from the Lipschitz continuity of ∇F : since $\nabla e^{-F} = -e^{-F} \nabla F$, we have

$$\begin{aligned} \|\nabla e^{-F(x)} - \nabla e^{-F(y)}\| &\leq e^{-F(x)} \|\nabla F(x) - \nabla F(y)\| + \|\nabla F(y)\| |e^{-F(x)} - e^{-F(y)}| \\ &\leq (L + L^2) \|x - y\|. \end{aligned}$$

Finally, suppose that F is also bounded from above, $F \leq a$ for some $a > 0$. Then $f \geq c$ everywhere, where $0 < c \leq 1$ because both μ and γ_d are probability measures.

We will also need the following simple lemma:

Lemma 9 (Regularity of the Föllmer drift) *Under Assumption 1, the Föllmer drift $b(x, t) = \nabla \log Q_{1-t}f(x)$ is bounded in norm by L/c and is Lipschitz with Lipschitz constant $L/c + L^2/c^2$, where L is the maximum of the Lipschitz constants of f and ∇f .*

Proof The heat semigroup $Q_t f(x) = \mathbf{E}[f(x + \sqrt{t}Z)]$, $Z \sim \gamma_d$, commutes with the gradient operator: for any differentiable and Lipschitz $f : \mathbb{R}^d \rightarrow \mathbb{R}$, $\partial_i Q_t f = Q_t \partial_i f$ for all $i \in [d]$ (Stroock, 2008, Corollary 2.2.8). Therefore, since $f(x) \geq c$ and $\|\nabla f(x)\| \leq L$ for all x , we have $Q_t f(x) \geq c$ and $\|\nabla Q_t f(x)\| \leq L$ for all $x \in \mathbb{R}^d$ and all $t \geq 0$. Consequently, for any $x \in \mathbb{R}^d$ and $t \in [0, 1]$

$$\|b(x, t)\| = \left\| \frac{\nabla Q_{1-t}f(x)}{Q_{1-t}f(x)} \right\| \leq \frac{L}{c}.$$

Also, since ∇f is Lipschitz, $\|\nabla Q_t f(x) - \nabla Q_t f(x')\| \leq L\|x - x'\|$ for any $x, x' \in \mathbb{R}^d$ and $t \in [0, 1]$, and thus

$$\begin{aligned} \|b(x, t) - b(x', t)\| &= \left\| \frac{\nabla Q_{1-t}f(x)}{Q_{1-t}f(x)} - \frac{\nabla Q_{1-t}f(x')}{Q_{1-t}f(x')} \right\| \\ &\leq \frac{\|\nabla Q_{1-t}f(x) - \nabla Q_{1-t}f(x')\|}{Q_{1-t}f(x')} + \|b(x, t)\| \cdot \frac{|Q_{1-t}f(x) - Q_{1-t}f(x')|}{Q_{1-t}f(x')} \\ &\leq \left(\frac{L}{c} + \frac{L^2}{c^2} \right) \|x - x'\|, \end{aligned}$$

and the proof is complete. \blacksquare

Appendix C. Uniform approximation of the heat semigroup by a finite sum

In this appendix, we prove the following result, which is used in the proof of Theorem 5:

Theorem 10 *For any $\varepsilon > 0$ and any $R > 0$, there exist $N = \text{poly}(1/\varepsilon, d, L, R)$ points $z_1, \dots, z_N \in \mathbb{R}^d$, for which the following holds:*

$$\begin{aligned} \max_{n \leq N} \|z_n\| &\leq 8\sqrt{(d+6)\log N} \\ \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \left| \frac{1}{N} \sum_{n=1}^N f(x + \sqrt{t}z_n) - Q_t f(x) \right| &\leq \varepsilon \\ \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \left\| \frac{1}{N} \sum_{n=1}^N \nabla f(x + \sqrt{t}z_n) - \nabla Q_t f(x) \right\| &\leq \varepsilon \end{aligned}$$

We gather some preliminaries first. We recall the definition of the Orlicz exponential norm of order 2 (Giné and Nickl, 2016, Sec. 2.3): for a real-valued random variable U ,

$$\|U\|_{\psi_2} := \inf \left\{ c > 0 : \mathbf{E} \exp \left(\frac{|U|}{c} \right)^2 \leq 2 \right\}.$$

The ψ_2 norm dominates the L^2 norm $\|U\|_2 := (\mathbf{E}|U|^2)^{1/2}$: $\|U\|_2 \leq \|U\|_{\psi_2}$. A simple application of Markov's inequality leads to the following tail bound:

$$\mathbf{P} [|U| \geq t\|U\|_{\psi_2}] \leq \frac{1}{e^{t^2} - 1}. \quad (\text{C.2})$$

Lemma 11 *Let $U = \|Z\|$, where $Z \sim \gamma_d$. Then $\|U\|_{\psi_2} \leq \sqrt{d} + \sqrt{6}$.*

Proof If $F : \mathbb{R}^d \rightarrow \mathbb{R}$ is 1-Lipschitz, then the centered random variable $\xi = F(Z) - \mathbf{E}F(Z)$ has subgaussian tails (Boucheron et al., 2013, Theorem 5.6):

$$\mathbf{P} \{|\xi| \geq t\} \leq 2e^{-t^2/2} \text{ for all } t > 0.$$

This implies that $\|\xi\|_{\psi_2} \leq \sqrt{6}$ (Giné and Nickl, 2016, Eq. (2.25)). Taking $F(Z) = U$ and using the triangle inequality, we obtain

$$\|U\|_{\psi_2} \leq \mathbf{E}U + \|U - \mathbf{E}U\|_{\psi_2} \leq \sqrt{d} + \sqrt{6},$$

where $\mathbf{E}U \leq \|U\|_2 = \sqrt{d}$ by Jensen's inequality. ■

Let U_1, \dots, U_N , $N \geq 2$, be a collection of (possibly dependent) random variables with finite ψ_2 norms. Then we have the following maximal inequality:

$$\left\| \max_{j \leq N} |U_j| \right\|_{\psi_2} \leq 4\sqrt{\log N} \max_{j \leq N} \|U_j\|_{\psi_2}, \quad (\text{C.3})$$

(Lemma 2.3.3 in Giné and Nickl (2016)).

We also need some results on suprema of empirical processes. Let \mathcal{G} be a class of real-valued functions on some measurable space Z . We say that a positive function $F : Z \rightarrow \mathbb{R}_+$ is an *envelope* of \mathcal{G} if $|g(z)| \leq F(z)$ for all $g \in \mathcal{G}$ and $z \in Z$. Let Z_1, \dots, Z_N be i.i.d. random elements of Z with probability law P and denote by P_N the corresponding empirical distribution, i.e., $P_N(A) = N^{-1} \sum_{n \leq N} \mathbf{1}_{\{Z_n \in A\}}$ for all measurable sets $A \subset Z$. We will use the linear functional notation for expectations, i.e., $Pg := \mathbf{E}_P[g(Z)]$ and $P_N g := \mathbf{E}_{P_N}[g(Z)] = N^{-1} \sum_{n \leq N} g(Z_n)$. We are interested in the quantity

$$\|P_N - P\|_{\mathcal{G}} := \sup_{g \in \mathcal{G}} |P_N g - Pg|,$$

which is a random variable under standard regularity assumptions on \mathcal{G} , such as separability. The expected supremum $\mathbf{E}\|P_N - P\|_{\mathcal{G}}$ is controlled by the covering numbers of \mathcal{G} . The $L^2(Q)$ covering numbers of \mathcal{G} with respect to a probability measure Q on Z are defined by

$$N(\mathcal{G}, L^2(Q), \varepsilon) := \min \left\{ K : \text{there exist } f_1, \dots, f_K \in L^2(Q) \right. \\ \left. \text{such that } \sup_{g \in \mathcal{G}} \min_{k \leq K} \|g - f_k\|_{L^2(Q)} \leq \varepsilon \right\}.$$

The *Koltchinskii–Pollard* ε -entropy of \mathcal{G} is given by

$$H(\mathcal{G}, F, \varepsilon) := \sup_Q \sqrt{\log 2N(\mathcal{G}, L^2(Q), \varepsilon \|F\|_{L^2(Q)})},$$

where the supremum is over all probability measures Q supported on finitely many points of Z . Then we have the following bound on the expectation of $\|P_N - P\|_{\mathcal{G}}$ (Theorem 3.54 and Eq. (3.177) in Giné and Nickl (2016)):

Lemma 12 *Let \mathcal{G} be a class of functions containing 0, such that*

$$J(\mathcal{G}, F) := \int_0^\infty H(\mathcal{G}, F, \varepsilon) d\varepsilon < \infty.$$

Let Z_1, \dots, Z_N be i.i.d. copies of a random element Z of Z with probability law P , such that $F \in L^2(P)$. Then

$$\mathbf{E} \|P_N - P\|_{\mathcal{G}} \leq \frac{8\sqrt{2}J(\mathcal{G}, F)\|F\|_{L^2(P)}}{\sqrt{N}}.$$

We also have the following generalization of Talagrand's concentration inequality to unbounded classes of functions, due to Adamczak (2008) (see also Sec. 2.3 in Koltchinskii (2011)):

Lemma 13 *Let \mathcal{G} be a class of real-valued functions on Z with envelope F . Then there exists an absolute constant $C > 0$, such that, for any $\gamma > 0$,*

$$\mathbf{P} \left\{ \|P_N - P\|_{\mathcal{G}} \geq C \left[\mathbf{E} \|P_N - P\|_{\mathcal{G}} + \sigma_P(\mathcal{G}) \sqrt{\frac{\gamma}{N}} + \left\| \max_{n \leq N} F(Z_n) \right\|_{\psi_2} \frac{\sqrt{\gamma}}{N} \right] \right\} \leq e^{-\gamma},$$

where

$$\sigma_P^2(\mathcal{G}) := \sup_{g \in \mathcal{G}} \left(P g^2 - (P g)^2 \right).$$

With these preliminaries out of the way, we have the following result:

Lemma 14 *Let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be L -Lipschitz with respect to the Euclidean norm. Let Z_1, \dots, Z_N be i.i.d. copies of a d -dimensional random vector Z , such that $U := \|Z\|$ has finite ψ_2 norm. Then there exists an absolute constant $C > 0$, such that, for any $\gamma > 0$,*

$$\begin{aligned} & \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \left| \frac{1}{N} \sum_{n=1}^N g(x + \sqrt{t}Z_n) - \mathbf{E}[g(x + \sqrt{t}Z)] \right| \\ & \leq C \left[\frac{16L\sqrt{6\pi R d}((R \vee 1) + \|U\|_{\psi_2})}{\sqrt{N}} + 5L((R \vee 1) + \|U\|_{\psi_2}) \sqrt{\frac{\gamma}{N}} \right] \end{aligned} \quad (\text{C.4})$$

with probability at least $1 - e^{-\gamma}$.

Proof For each $x \in \mathbb{R}^d$ and $t \geq 0$ let $g_{x,t}(z) := g(x + \sqrt{t}z)$. Let P denote the probability law of Z . Since $P_N g_{x,t} - P g_{x,t} = P_N(g_{x,t} - g_{0,0}) - P(g_{x,t} - g_{0,0})$ for all x, t , where $g_{0,0}(\cdot) = g(0)$ is a constant, we can replace each $g_{x,t}$ with $\bar{g}_{x,t} := g_{x,t} - g_{0,0}$, introduce the function class $\mathcal{G} := \{\bar{g}_{x,t} : x \in \mathbb{B}^d(R), t \in [0, 1]\}$, and analyze the empirical process supremum

$$\|P_N - P\|_{\mathcal{G}} = \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} |P_N \bar{g}_{x,t} - P \bar{g}_{x,t}|.$$

Define the function $F(z) := L((R \vee 1) + \|z\|)$. Since $\|\cdot\|_2 \leq \|\cdot\|_{\psi_2}$, $F \in L^2(P)$. By Lipschitz continuity, for all $z \in \mathbb{R}^d$, $x \in \mathbb{B}^d(R)$, $t \in [0, 1]$, we have

$$|\bar{g}_{x,t}(z)| \leq |g(x + \sqrt{t}z) - g(0)| \leq L\|x + \sqrt{t}z\| \leq F(z),$$

so F is a square-integrable envelope of \mathcal{G} . Moreover, for any probability measure Q supported on finitely many points in \mathbb{R}^d and for all $x, x' \in \mathbb{B}^d(R)$ and $t, t' \in [0, 1]$,

$$\|\bar{g}_{x,t} - \bar{g}_{x',t'}\|_{L^2(Q)} \leq \|F\|_{L^2(Q)} \cdot (\|x - x'\| + |t - t'|^{1/2}).$$

Thus we can estimate the $L^2(Q)$ covering numbers of \mathcal{G} by

$$N(\mathcal{G}, L^2(Q), \varepsilon \|F\|_{L^2(Q)}) \leq N(\mathbb{B}^d(R), \|\cdot\|, \varepsilon/2) \cdot N([0, 1], |\cdot|, \varepsilon^2/4).$$

Using standard volumetric estimates on the covering numbers of ℓ_2 balls, we obtain the following bound on the Koltchinskii–Pollard entropy of \mathcal{G} :

$$H(\mathcal{G}, F, \varepsilon) \leq \left(4d \log \frac{2\sqrt{3R}}{\varepsilon} \right)_+$$

where $(u)_+ := u \vee 0$, and therefore

$$J(\mathcal{G}, F) = \int_0^\infty H(\mathcal{G}, F, \varepsilon) d\varepsilon \leq 2\sqrt{3\pi R}d.$$

Lemma 12 then gives

$$\begin{aligned} \mathbf{E}\|P_N - P\|_{\mathcal{G}} &\leq \frac{8\sqrt{2}J(\mathcal{G})\|F\|_{L^2(P)}}{\sqrt{N}} \\ &\leq \frac{16\sqrt{6\pi R}d\|F\|_{L^2(P)}}{\sqrt{N}} \\ &= \frac{16L\sqrt{6\pi R}d((R \vee 1) + \|U\|_2)}{\sqrt{N}} \\ &\leq \frac{16L\sqrt{6\pi R}d((R \vee 1) + \|U\|_{\psi_2})}{\sqrt{N}} \end{aligned} \tag{C.5}$$

Furthermore, we estimate

$$\begin{aligned} \sigma_P(\mathcal{G}) &\leq \|F\|_{L^2(P)} \\ &\leq \|F(Z)\|_{\psi_2} \\ &= \|L((R \vee 1) + U)\|_{\psi_2} \\ &\leq L((R \vee 1) + \|U\|_{\psi_2}) \end{aligned} \tag{C.6}$$

and

$$\begin{aligned} \left\| \max_{j \leq N} F(Z_j) \right\|_{\psi_2} &= L \left\| (R \vee 1) + \max_{j \leq N} U_j \right\|_{\psi_2} \\ &\leq L(R \vee 1) + 4L\sqrt{\log N}\|U\|_{\psi_2}, \end{aligned} \tag{C.7}$$

where we have used the triangle inequality for $\|\cdot\|_{\psi_2}$, as well as the maximal inequality (C.3). Using the estimates (C.5), (C.6), and (C.7) in Adamczak's inequality, we obtain (C.4). \blacksquare

We are now ready to prove Theorem 10. The proof is via the probabilistic method. Let $\varepsilon > 0$ and $R > 0$ be given, and choose

$$N = \left\lceil \left(\frac{C\sqrt{d}}{\varepsilon} \cdot L \left((R \vee 1) + \sqrt{d} + \sqrt{6} \right) \cdot \left(16\sqrt{6\pi R d} + 5\sqrt{\log 4(d+1)} \right) \right)^2 \right\rceil,$$

where $C > 0$ is the absolute constant in the bound of Lemma 14. Let Z_1, \dots, Z_N be i.i.d. copies of $Z \sim \gamma_d$, and observe that $\mathbf{E}[f(x + \sqrt{t}Z)] = Q_t f(x)$ and $\mathbf{E}[\partial_i f(x + \sqrt{t}Z)] = \partial_i Q_t f(x) = Q_t \partial_i f(x)$ for all $x \in \mathbb{R}^d$, $t \geq 0$, and $i \in [d]$. Define the events

$$\begin{aligned} E_0 &:= \left\{ \max_{n \leq N} \|Z_n\| \geq 8\sqrt{(d+6)\log N} \right\} \\ E_1 &:= \left\{ \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \left| \frac{1}{N} \sum_{n=1}^N f(x + \sqrt{t}Z_n) - Q_t f(x) \right| \geq \varepsilon \right\} \\ E_2 &:= \left\{ \max_{i \in [d]} \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \left| \frac{1}{N} \sum_{n=1}^N \partial_i f(x + \sqrt{t}Z_n) - \partial_i Q_t f(x) \right| \geq \frac{\varepsilon}{\sqrt{d}} \right\}. \end{aligned}$$

We will show that $\mathbf{P}\{E_0 \cup E_1 \cup E_2\} < 1$, which will imply that there exists at least one realization of Z_1, \dots, Z_N verifying the statement of the theorem.

By Lemma 11, $U = \|Z\|$ satisfies $\|U\|_{\psi_2} \leq \sqrt{d} + \sqrt{6}$, and therefore $U_N^* := \max_{n \leq N} U_n$ satisfies $\|U_N^*\|_{\psi_2} \leq \sqrt{32(d+6)\log N}$ by the maximal inequality (C.3). Consequently, it follows from (C.2) that

$$\mathbf{P}\{E_0\} \leq \mathbf{P}\{U_N^* \geq \sqrt{2}\|U_N^*\|_{\psi_2}\} \leq \frac{1}{e^2 - 1} \leq \frac{1}{4}.$$

Moreover, since the function f and all of its partial derivatives are L -Lipschitz, Lemma 14 (with $\gamma = \log 4(d+1)$) and the union bound give $\mathbf{P}\{E_1 \cup E_2\} \leq 1/4$. Therefore, $\mathbf{P}\{E_0 \cup E_1 \cup E_2\} \leq 1/2$.

Appendix D. The proof of Theorem 5: uniform approximation of the Föllmer drift by a neural net

We first collect a few preliminaries.

Lemma 15 (cheap gradient principle, Griewank and Walther (2008)) *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be implementable by a neural net with differentiable activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$, where the neural net has size (number of nodes) m and depth (number of layers) ℓ . Then each coordinate of the gradient ∇f can be computed by a neural net that has size $\mathcal{O}(m + \ell)$, and where the activation function of each neuron is an element of the set $\{\sigma, \sigma'\}$.*

Lemma 16 (approximating multiplication and reciprocals) *Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be an activation function satisfying Assumption 2. Then:*

1. For any $M > 0$ and any $\delta > 0$, there exists a 2-layer neural net $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ of size $m \leq 8c_\sigma \frac{M^2}{\delta} + 1$, such that

$$\sup_{x,y \in [-M,M]} |g(x,y) - xy| \leq \delta. \quad (\text{D.1})$$

2. For any $0 < a \leq b < \infty$ and any $\delta > 0$, there exists a 2-layer neural net $q : \mathbb{R} \rightarrow \mathbb{R}$ of size $m \leq c_\sigma \frac{b}{a^2\delta} + 1$, such that

$$\sup_{x \in [a,b]} \left| q(x) - \frac{1}{x} \right| \leq \delta. \quad (\text{D.2})$$

Remark 17 These approximations suffice for our purposes. However, if one uses the ReLU activation function $x \mapsto x \vee 0$, then both multiplication and reciprocals can be ε -approximated by neural nets with size and depth polylogarithmic in $1/\varepsilon$ (Yarotsky, 2017; Telgarsky, 2017).

Proof For multiplication, we first consider the function $x \mapsto x^2 \wedge (4M^2)$, which is $4M$ -Lipschitz and constant outside the interval $[-2M, 2M]$. Assumption 2 then grants the existence of a univariate function $g_0 : \mathbb{R} \rightarrow \mathbb{R}$ of the form (3.3) with $m \leq 4c_\sigma \frac{M^2}{\delta}$ satisfying $|g_0(x) - x^2| \leq 2\delta$ for all $x \in [-M, M]$. The desired approximation $g : \mathbb{R}^2 \rightarrow \mathbb{R}$ is given by

$$g(x,y) = \frac{1}{4} (g_0(x+y) - g_0(x-y)),$$

which is a 2-layer neural net with size $m \leq 8c_\sigma \frac{M^2}{\delta} + 1$. Indeed, using the polarization identity $4xy = (x+y)^2 - (x-y)^2$, we have

$$\begin{aligned} & \sup_{x,y \in [-M,M]} |g(x,y) - xy| \\ & \leq \frac{1}{4} \sup_{x,y \in [-M,M]} \left| g_0(x+y) - (x+y)^2 \right| + \frac{1}{4} \sup_{x,y \in [-M,M]} \left| g_0(x-y) - (x-y)^2 \right| \\ & \leq \delta. \end{aligned}$$

For approximating the reciprocal, consider the univariate function

$$x \mapsto \frac{1}{a} \mathbf{1}\{x < a\} + \frac{1}{x} \mathbf{1}\{a \leq x \leq b\} + \frac{1}{b} \mathbf{1}\{x > b\},$$

which is $(1/a^2)$ -Lipschitz and constant outside of the interval $[-b, b]$. The existence of the function q with the stated properties follows immediately from Assumption 2. \blacksquare

We now prove Theorem 5. Let $\delta = \frac{c^2\varepsilon}{16L}$. By Theorem 10, there exist points $z_1, \dots, z_N \in \mathbb{R}^d$ with $N = \text{poly}(1/\delta, d, L, R)$, such that $R_{N,d} := \max_{n \leq N} \|z_n\| \leq 8\sqrt{(d+6)\log N}$, and the function $\varphi : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}$ defined by

$$\varphi(x, t) := \frac{1}{N} \sum_{n=1}^N f(x + tz_n)$$

satisfies

$$\sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} |\varphi(x, \sqrt{t}) - Q_t f(x)| \leq \delta \quad \text{and} \quad \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \|\nabla \varphi(x, \sqrt{t}) - \nabla Q_t f(x)\| \leq \delta.$$

By Assumption 3, there exists a neural net $\widehat{f} : \mathbb{R}^d \rightarrow \mathbb{R}$ be that approximates f and the gradient of f to accuracy δ on the blown-up ball $\mathbb{B}^d(R + R_{N,d})$. Then the function

$$\widehat{\varphi} : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}, \quad \widehat{\varphi}(x, t) := \frac{1}{N} \sum_{n=1}^N \widehat{f}(x + tz_n)$$

can be computed by a neural net of size $N \cdot \text{poly}(1/\delta, d, L, R)$, such that

$$\begin{aligned} & \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} |\widehat{\varphi}(x, \sqrt{t}) - Q_t f(x)| \\ & \leq \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} |\widehat{\varphi}(x, \sqrt{t}) - \varphi(x, \sqrt{t})| + \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} |\varphi(x, \sqrt{t}) - Q_t f(x)| \\ & \leq \sup_{x \in \mathbb{B}^d(R+R_{N,d})} |\widehat{f}(x) - f(x)| + \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} |\varphi(x, \sqrt{t}) - Q_t f(x)| \leq 2\delta \end{aligned}$$

and

$$\begin{aligned} & \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \|\nabla \widehat{\varphi}(x, \sqrt{t}) - \nabla Q_t f(x)\| \\ & \leq \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \|\nabla \widehat{\varphi}(x, \sqrt{t}) - \nabla \varphi(x, \sqrt{t})\| + \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \|\nabla \varphi(x, \sqrt{t}) - \nabla Q_t f(x)\| \\ & \leq \sup_{x \in \mathbb{B}^d(R+R_{N,d})} \|\nabla \widehat{f}(x) - \nabla f(x)\| + \sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0,1]} \|\nabla \varphi(x, \sqrt{t}) - \nabla Q_t f(x)\| \leq 2\delta. \end{aligned}$$

Since f is L -Lipschitz and bounded below by c , we have $c \leq Q_t f(x) \leq L(\|x\| + \sqrt{d}) + f(0)$ for any $x \in \mathbb{R}^d$ and $t \in [0, 1]$. Therefore, on $\mathbb{B}^d(R) \times [0, 1]$,

$$\frac{c}{2} \leq \widehat{\varphi}(x, \sqrt{t}) \leq L(R + \sqrt{d}) + f(0) + \frac{c}{2}$$

where we have used the fact that $\delta \leq c/4$. Without loss of generality, we may assume that $L \geq 1$. Then, for any $x \in \mathbb{B}^d(R)$ and $t \in [0, 1]$,

$$\begin{aligned} & \left\| \nabla \log \widehat{\varphi}(x, \sqrt{t}) - \nabla \log Q_t f(x) \right\| \\ & = \left\| \frac{\nabla \widehat{\varphi}(x, \sqrt{t})}{\widehat{\varphi}(x, \sqrt{t})} - \frac{\nabla Q_t f(x)}{Q_t f(x)} \right\| \\ & \leq \frac{1}{\widehat{\varphi}(x, \sqrt{t})} \|\nabla \widehat{\varphi}(x, \sqrt{t}) - \nabla Q_t f(x)\| + \left\| \frac{\nabla Q_t f(x)}{Q_t f(x)} \right\| \frac{|\widehat{\varphi}(x, \sqrt{t}) - Q_t f(x)|}{\widehat{\varphi}(x, \sqrt{t})} \\ & \leq \frac{2L}{c} \cdot 2\delta + \frac{L}{c} \cdot \frac{2}{c} \cdot 2\delta \\ & \leq \frac{\varepsilon}{2}, \end{aligned}$$

where we have used Lemma 9 to bound $\|\frac{\nabla Q_t f}{Q_t f}\| \leq L/c$. In other words, $\nabla \log \widehat{\varphi}(x, \sqrt{t})$ approximates $\nabla \log Q_t f(x)$ to accuracy $\varepsilon/2$ uniformly on $\mathbb{B}^d(R) \times [0, 1]$. It remains to approximate $\nabla \log \widehat{\varphi}(x, \sqrt{t})$ by a neural net to accuracy $\varepsilon/2$.

To that end, we first represent $\nabla \log \widehat{\varphi}(x, \sqrt{t})$ as a composition of several elementary operations and then approximate each step by a neural net. Specifically, the computation of $v_i = \partial_i \log \widehat{\varphi}(x, \sqrt{t})$ can be represented as a computation graph with the following structure:

1. Compute $a = \widehat{\varphi}(x, \sqrt{t})$.
2. Compute $b_i = \partial_i \widehat{\varphi}(x, \sqrt{t})$.
3. Compute $r = 1/a$.
4. Compute $v_i = r b_i$.

Given x and \sqrt{t} , a is computed by a neural net with activation function σ , of size $\text{poly}(1/\delta, d, L, R)$ and depth $\text{poly}(1/\delta, d, L, R)$. Therefore, by the cheap gradient principle (Lemma 15), b_i can be computed by a neural net of size $\text{poly}(1/\delta, d, L, R)$, where the activation function of each neuron is an element of the set $\{\sigma, \sigma'\}$. Next, since a takes values in $[c/2, L(R + \sqrt{d}) + f(0) + c/2]$, by Lemma 16 the reciprocal $r = 1/a$ can be computed to accuracy $\varepsilon/(4L\sqrt{d})$ by a 2-layer neural net with activation function σ and of size

$$\mathcal{O}\left(\frac{4}{c^2} \cdot (L(R + \sqrt{d}) + f(0) + c/2) \cdot \frac{4L\sqrt{d}}{\varepsilon}\right) \leq \text{poly}(1/\varepsilon, d, L, R, c, 1/c)$$

Let \widehat{r} denote the resulting approximation. Then, since $|b_i| \leq 2L$ and $|\widehat{r}| \leq 2/c + \varepsilon/(4L\sqrt{d}) \leq 4/c$, by Lemma 16 the product $\widehat{r}b_i$ can be approximated to accuracy $\varepsilon/4\sqrt{d}$ by a 2-layer neural net with activation function σ and with at most

$$\mathcal{O}\left((4/c \vee 2L)^2 \cdot \frac{4\sqrt{d}}{\varepsilon}\right) \leq \text{poly}(1/\varepsilon, d, L, 1/c)$$

neurons. The overall accuracy of approximation is

$$|\widehat{v}_i - v_i| \leq |\widehat{v}_i - \widehat{r}b_i| + |\widehat{r}b_i - r b_i| \leq \frac{\varepsilon}{2\sqrt{d}}.$$

Thus, the vector $v = (v_1, \dots, v_d)$ can be $\varepsilon/2$ -approximated by $\tilde{v}(x, \sqrt{t})$, where $\tilde{v} : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}^d$ is a neural net with vector-valued output that has the size $\text{poly}(1/\varepsilon, d, L, R, c, 1/c)$. Finally, since $\sup_{x \in \mathbb{B}^d(R)} \sup_{t \in [0, 1]} |\tilde{v}_i(x, \sqrt{t})| \leq 2L/c$, the function

$$\widehat{v}_i(x, \sqrt{t}) := \min\{\max\{\tilde{v}_i(x, \sqrt{t}), -2L/c\}, 2L/c\}$$

is continuous, takes values in $[-2L/c, 2L/c]$ and coincides with \tilde{v}_i on $\mathbb{B}^d(R) \times [0, 1]$. Moreover, the min and max operations can each be implemented exactly using $\mathcal{O}(1)$ ReLU neurons.

Appendix E. Proof of Theorem 6

E.1. Unbiasedness

We follow the strategy of [Henry-Labordère et al. \(2017\)](#) and construct a sequence $\{\psi_n\}_{n \geq 0}$ of unbiased estimators, such that $\mathbf{E}[\psi_n] \xrightarrow{n \rightarrow \infty} \mathbf{E}[\psi]$, where $\psi := \lim_{n \rightarrow \infty} \psi_n$. By a standard approximation argument, we can assume that g is bounded and Lipschitz.

Let $\Delta_k^T := T_k - T_{k-1}$ and $\Delta_k^W := W_{T_k} - W_{T_{k-1}}$, for $k \geq 1$. For each $n \geq 0$, let

$$\begin{aligned} \psi_n := & g(\widehat{X}_1) \cdot \frac{1}{1 - F_\tau(\Delta_{n+1}^T)} \prod_{k=1}^{N \wedge n} \frac{\left(b(\widehat{X}_{T_k}, T_k) - b(\widehat{X}_{T_{k-1}}, T_{k-1})\right)^T \Delta_{k+1}^W}{f_\tau(\Delta_k^T) \Delta_{k+1}^T} \cdot \mathbf{1}_{\{N \leq n\}} \\ & + \prod_{k=1}^{n+1} \frac{1}{f_\tau(\Delta_k^T)} \cdot \left(b(\widehat{X}_{T_{n+1}}, T_{n+1}) - b(\widehat{X}_{T_n}, T_n)\right)^T \nabla h(\widehat{X}_{T_{n+1}}, T_{n+1}) \cdot \frac{\Delta_{n+1}^W}{\Delta_{n+1}^T} \cdot \mathbf{1}_{\{N > n\}}, \end{aligned} \quad (\text{E.1})$$

where $h(x, t) := \mathbf{E}[g(X_1) | X_t = x]$. We will show that $\mathbf{E}[\psi_n] = \mathbf{E}[g(X_1)]$ for all n and that the sequence $\{\psi_n\}_{n \geq 0}$ is uniformly integrable. Then it will follow from the dominated convergence theorem that

$$\psi = \lim_{n \rightarrow \infty} \psi_n = \frac{1}{1 - F_\tau(1 - T_N)} \cdot g(\widehat{X}_1) \cdot \prod_{k=1}^N \frac{1}{f_\tau(T_k - T_{k-1})} \widehat{W}_k \quad (\text{E.2})$$

is also an unbiased estimator. Observe that the estimator $\widehat{\psi}$ defined in (C.2) differs from ψ : instead of $g(\widehat{X}_1)$, we have $g(\widehat{X}_1) - g(\widehat{X}_N) \mathbf{1}_{\{N > 0\}}$. Just as in [Henry-Labordère et al. \(2017\)](#), the term proportional to $g(\widehat{X}_N) \mathbf{1}_{\{N > 0\}}$ serves as a *control variate* to ensure that $\widehat{\psi}$ has finite variance. Indeed, since $\mathbf{E}[\Delta_{N+1}^W | T_N] = 0$, it is easy to see that

$$\mathbf{E} \left[\frac{1}{1 - F_\tau(1 - T_N)} \cdot g(\widehat{X}_N) \mathbf{1}_{\{N > 0\}} \cdot \prod_{k=1}^N \frac{1}{f_\tau(T_k - T_{k-1})} \widehat{W}_k \right] = 0,$$

and therefore $\mathbf{E}[\widehat{\psi} - \psi] = 0$.

Given $x, v \in \mathbb{R}^d$ and $t \in [0, 1]$, consider the constant-drift diffusion process $\{\tilde{X}_s^{t,x,v}\}_{s \in [t, 1]}$ with $\tilde{X}_t^{t,x,v} = x$ and

$$d\tilde{X}_s^{t,x,v} = v ds + dW_s, \quad s \in [t, 1].$$

This process has the infinitesimal generator

$$\mathcal{L}^v h(x, t) := v^T \nabla h(x, t) + \frac{1}{2} \text{tr} \nabla^2 h(x, t), \quad \forall h \in C^{2,1}(\mathbb{R}^d \times [0, 1]).$$

Then, by Dynkin's formula ([Kallenberg, 2002](#), Lemma 19.21), for any $t \leq s \leq 1$,

$$h(\tilde{X}_s^{t,x,v}, s) = h(\tilde{X}_t^{t,x,v}, t) + \int_t^s \left\{ \frac{\partial}{\partial r} + \mathcal{L}^v \right\} h(\tilde{X}_r^{t,x,v}, r) dr + M_s^t, \quad (\text{E.3})$$

where $\{M_s^t\}_{s \in [t, 1]}$ is a martingale. In particular, let $h \in C^{2,1}(\mathbb{R}^d \times [0, 1])$ be a bounded solution of the Cauchy problem

$$\frac{\partial h}{\partial t} + \mathcal{L}_t h = 0, \quad h(\cdot, 1) = g(\cdot) \quad (\text{E.4})$$

where

$$\mathcal{L}_t h(x, t) := b(x, t)^T \nabla h(x, t) + \frac{1}{2} \text{tr} \nabla^2 h(x, t).$$

Rewriting (E.4) as

$$\frac{\partial h}{\partial t} + \mathcal{L}^v h = (v - b)^T h, \quad h(\cdot, 1) = g(\cdot)$$

and using this in (E.3), we obtain the formula

$$\begin{aligned} & h(\tilde{X}_s^{t,x,v}, s) \\ &= g(\tilde{X}_1^{t,x,v}) + \int_s^1 (b(\tilde{X}_r^{t,x,v}, r) - v)^T \nabla h(\tilde{X}_r^{t,x,v}, r) dr + M_s^t - M_1^t, \quad t \leq s \leq 1. \end{aligned}$$

In particular, since $h(x, t) = \mathbf{E}[g(X_1) | X_t = x]$ by the Feynman–Kac formula, we have

$$h(x, t) = \mathbf{E} \left[g(\tilde{X}_1^{t,x,v}) + \int_t^1 (b(\tilde{X}_s^{t,x,v}, s) - v)^T \nabla h(\tilde{X}_s^{t,x,v}, s) ds \right], \quad (\text{E.5})$$

where $\mathbf{E}[M_t^t - M_1^t] = 0$ since $M^{h,t}$ is a martingale.

Using Eq. (E.5) with $t = 0$ and $v = v_0 := b(x, 0)$, we have

$$h(x, 0) = \mathbf{E} \left[g(\tilde{X}_1^{0,x,v_0}) + \int_0^1 (b(\tilde{X}_s^{0,x,v_0}, s) - b(x, 0))^T \nabla h(\tilde{X}_s^{0,x,v_0}, s) ds \right].$$

Recalling that $T_1 = \tau_1 \wedge 1$ is independent of the Brownian motion $\{W_t\}$ and $\mathbf{P}[T_1 \geq 1] = \mathbf{P}[\tau_1 \geq 1] = 1 - F_\tau(1)$, we have

$$\mathbf{E}[g(\tilde{X}_1^{0,x,v_0})] = \frac{1}{1 - F_\tau(1)} \mathbf{E}[g(\tilde{X}_1^{0,x,v_0}) \mathbf{1}_{\{T_1 \geq 1\}}], \quad (\text{E.6})$$

and

$$\begin{aligned} & \mathbf{E} \left[\int_0^1 (b(\tilde{X}_s^{0,x,v_0}, s) - b(x, 0))^T \nabla h(\tilde{X}_s^{0,x,v_0}, s) ds \right] \\ &= \mathbf{E} \left[\frac{1}{f_\tau(T_1)} (b(\tilde{X}_{T_1}^{0,x,v_0}, T_1) - b(x, 0))^T \nabla h(\tilde{X}_{T_1}^{0,x,v_0}, T_1) \mathbf{1}_{\{T_1 < 1\}} \right]. \end{aligned} \quad (\text{E.7})$$

Since the process \tilde{X}^{0,x,v_0} coincides with \hat{X} on $[0, T_1]$, it follows from (E.6) and (E.7) that

$$\begin{aligned} & h(x, 0) \\ &= \mathbf{E} \left[\frac{1}{1 - F_\tau(\Delta_1^T)} g(\hat{X}_1) \mathbf{1}_{\{T_1 \geq 1\}} + \frac{1}{f_\tau(\Delta_1^T)} (b(\hat{X}_{T_1}, T_1) - b(\hat{X}_{T_0}, T_0))^T \nabla h(\hat{X}_{T_1}, T_1) \mathbf{1}_{\{T_1 < 1\}} \right] \\ &= \mathbf{E}[\psi_0], \end{aligned} \quad (\text{E.8})$$

where the last equality follows from the fact that $T_1 = \Delta_1^T \geq 1$ if and only if $N = 0$.

By Lemma 18 in Section E.3,

$$\begin{aligned}
 & \nabla h(x, 0) \\
 &= \mathbf{E} \left[g(\tilde{X}_1^{0,x,v_0}) W_1 + \int_0^1 \left((b(\tilde{X}_s^{t,x,v}, s) - b(x, 0))^T \nabla h(\tilde{X}_s^{t,x,v}, s) \right) \frac{W_s}{s} ds \right] \\
 &= \mathbf{E} \left[\frac{1}{1 - F_\tau(1)} g(\hat{X}_1) \frac{\Delta_1^W}{\Delta_1^T} \mathbf{1}_{\{T_1 \geq 1\}} \right. \\
 & \quad \left. + \frac{1}{f_\tau(\Delta_1^T)} \left((b(\hat{X}_{T_1}, T_1) - b(\hat{X}_{T_0}, T_0))^T \nabla h(\hat{X}_{T_1}, T_1) \right) \frac{\Delta_1^W}{\Delta_1^T} \mathbf{1}_{\{T_1 < 1\}} \right] \quad (\text{E.9})
 \end{aligned}$$

Moreover, if we change the initial condition from $t = 0, v = v_0$ to $t = T_1, v = v_1 := b(\hat{X}_{T_1}, T_1)$, then it follows from (E.9) that, conditionally on (\hat{X}_{T_1}, T_1) , whenever $T_1 < 1$,

$$\begin{aligned}
 \nabla h(\hat{X}_{T_1}, T_1) &= \mathbf{E} \left[\frac{1}{1 - F_\tau(\Delta_2^T)} g(\hat{X}_1) \frac{\Delta_2^W}{\Delta_2^T} \mathbf{1}_{\{T_2 \geq 1\}} \right. \\
 & \quad \left. + \frac{1}{f_\tau(\Delta_2^T)} \left((b(\hat{X}_{T_2}, T_2) - b(\hat{X}_{T_1}, T_1))^T \nabla h(\hat{X}_{T_2}, T_2) \right) \frac{\Delta_2^W}{\Delta_2^T} \mathbf{1}_{\{T_2 < 1\}} \middle| \hat{X}_{T_1}, T_1 \right]. \quad (\text{E.10})
 \end{aligned}$$

Substituting (E.10) into (E.8) and using the fact that the event $\{T_1 < 1 \leq T_2\}$ is equivalent to $\{N = 1\}$, we have $h(x, 0) = \mathbf{E}[\psi_1]$. Repeating this procedure, we have

$$\mathbf{E}[g(X_1) | X_0 = x] = h(x, 0) = \mathbf{E}[\psi_n], \quad n \geq 0.$$

We claim that the sequence $\{\psi_n\}_{n \geq 0}$ is uniformly integrable. To see this, first observe that, for each k , $\mathbf{E}[\|\Delta_{k+1}^W\| | T_{k+1}] \leq (\Delta_{k+1}^T d)^{1/2}$. Then the uniform integrability follows from the boundedness of $b, g, \nabla h$, and from Lemma 19 in Section E.3. Therefore, taking the limit as $n \rightarrow \infty$, we obtain

$$\mathbf{E}[g(X_1) | X_0 = x] = \lim_{n \rightarrow \infty} \mathbf{E}[\psi_n] = \mathbf{E} \left[\lim_{n \rightarrow \infty} \psi_n \right] = \mathbf{E}[\psi],$$

where the second equality follows from the dominated convergence theorem.

E.2. Variance

Let $L := L_b \vee L_g$. For $1 \leq k \leq N + 1$, let $\Delta_k^{\hat{X}} := \hat{X}_{T_{k+1}} - \hat{X}_{T_k}$ denote the increments of \hat{X} . Since $T_{N+1} = 1$, we have

$$\left| g(\hat{X}_1) - g(\hat{X}_{T_N}) \mathbf{1}_{\{N > 0\}} \right| \leq \begin{cases} |g(x)| + L |\Delta_1^{\hat{X}}|, & N = 0 \\ L |\Delta_{N+1}^{\hat{X}}|, & N > 0 \end{cases}$$

which gives

$$\left| g(\hat{X}_1) - g(\hat{X}_{T_N}) \mathbf{1}_{\{N > 0\}} \right| \leq |g(x)| \mathbf{1}_{\{N=0\}} + L \left(\sqrt{\Delta_{N+1}^T} + \|\Delta_{N+1}^{\hat{X}}\| \right).$$

Using this and (4.4), we can upper-bound $\widehat{\psi}$ as follows:

$$|\widehat{\psi}| \leq \frac{e^a}{1 - F_\tau(1)} \cdot \left(|g(x)| + L \left(\sqrt{\Delta_1^T} + \|\Delta_1^{\widehat{X}}\| \right) \right) \cdot \prod_{k=1}^N \frac{CL \left(\sqrt{\Delta_{k+1}^T} + \|\Delta_{k+1}^{\widehat{X}}\| \right)}{\Delta_{k+1}^T} \cdot \|\Delta_{k+1}^W\|,$$

where, for $k \geq 0$,

$$\begin{aligned} \|\Delta_{k+1}^{\widehat{X}}\| &= \|b(\widehat{X}_{T_k}, T_k) \cdot \Delta_{k+1}^T + \Delta_{k+1}^W\| \\ &\leq b_\infty \Delta_{k+1}^T + \|\Delta_{k+1}^W\|. \end{aligned}$$

Let $\mathcal{F}_k := \sigma(T_j, \widehat{X}_j : 1 \leq j \leq k)$. Then, since $\text{Law}(\Delta_{k+1}^W | \mathcal{F}_k) = \text{Law}((\Delta_{k+1}^T)^{1/2} Z | \mathcal{F}_k)$, where $Z \sim \gamma_d$ is independent of $\mathcal{F}_k \vee \sigma(T_{k+1})$, we have

$$\begin{aligned} &\mathbf{E} \left[\left(\frac{(\Delta_{k+1}^T)^{1/2} + \|\Delta_{k+1}^{\widehat{X}}\|}{\Delta_{k+1}^T} \cdot \|\Delta_{k+1}^W\| \right)^2 \middle| \mathcal{F}_k \right] \\ &\leq \mathbf{E}_k \left[\left(\frac{b_\infty \Delta_{k+1}^T + (\Delta_{k+1}^T)^{1/2} (1 + \|Z\|)}{\Delta_{k+1}^T} \cdot \sqrt{\Delta_{k+1}^T} \|Z\| \right)^2 \middle| \mathcal{F}_k \right] \\ &\leq \mathbf{E} \left[(1 + b_\infty + \|Z\|)^2 \|Z\|^2 \right] \\ &=: \kappa. \end{aligned}$$

Therefore, we can estimate

$$\mathbf{E}[\widehat{\psi}^2] \leq \left(\frac{e^a}{1 - F_\tau(1)} \right)^2 \cdot \mathbf{E} \left[\left(|g(x)| + L(1 + \sqrt{d}) \right)^2 \right] \cdot \mathbf{E} [\exp(\kappa N)].$$

E.3. Auxiliary lemmas

The following lemma is a straightforward consequence of the Gaussian integration-by-parts formula $\nabla_x \mathbf{E}[f(x + Z)] = \mathbf{E}[f(x + Z)Z]$, $Z \sim \gamma_d$, for any C^1 function $f : \mathbb{R}^d \rightarrow \mathbb{R}$:

Lemma 18 (Henry-Labordère et al. (2017)) *Let ν be a positive measure on $[0, 1]$. Let $\varphi : \mathbb{R}^d \times [0, 1] \rightarrow \mathbb{R}$ be a continuous function, such that*

$$\int_0^1 \mathbf{E} \left[\left\| \varphi(x + vt + W_t) \frac{W_t}{t} \right\| \right] \nu(dt) < \infty.$$

Then

$$\nabla_x \left(\int_0^1 \mathbf{E}[\varphi(x + vt + W_t)] \nu(dt) \right) = \int_0^1 \mathbf{E} \left[\varphi(x + vt + W_t) \frac{W_t}{t} \right] \nu(dt)$$

The next lemma is used to show that the sequence $\{\psi_n\}$ is uniformly integrable:

Lemma 19 *For any $C > 0$,*

$$\mathbf{E} \left[\frac{C^N}{1 - F_\tau(\Delta_{N+1}^T)} \prod_{k=1}^N \frac{1}{f_\tau(\Delta_k^T) (\Delta_{k+1}^T)^{1/2}} \right] < \infty. \quad (\text{E.11})$$

Proof For each $n \geq 0$, define the n -simplex

$$\mathcal{S}^n := \{(s_1, s_2, \dots, s_n) \in [0, 1]^n : 0 < s_1 < \dots < s_n < 1\}$$

with $s_0 \equiv 0$ and $s_{n+1} \equiv 1$. Consider the partial sums $S_k := \sum_{i=1}^k \tau_i$. Since the τ_i 's are i.i.d., the conditional joint density of (S_1, S_2, \dots, S_n) given $N = n$ is equal to

$$q_n(s_1, s_2, \dots, s_n) = \frac{1}{\mathbf{P}[N = n]} \cdot (1 - F_\tau(1 - s_n)) \cdot \prod_{k=1}^n f_\tau(s_k - s_{k-1}), \quad (s_1, \dots, s_n) \in \mathcal{S}^n$$

where we have set $s_0 \equiv 0$. Then a calculation similar to the one in Appendix B of [Andersson and Kohatsu-Higa \(2017\)](#) leads to

$$\begin{aligned} & \mathbf{E} \left[\frac{C^N}{1 - F_\tau(\Delta_{N+1}^T)} \prod_{k=1}^N \frac{1}{f_\tau(\Delta_k^T)(\Delta_{k+1}^T)^{1/2}} \right] \\ &= \sum_{n \geq 0} \mathbf{P}[N = n] \cdot C^n \int_{\mathcal{S}^n} \frac{1}{1 - F_\tau(s_n)} \prod_{k=1}^n \frac{1}{f_\tau(s_k - s_{k-1})(s_{k+1} - s_k)^{1/2}} q_n(s_1, \dots, s_n) \, ds \\ &\leq \sum_{n \geq 0} C^n \int_{\mathcal{S}^n} \prod_{k=0}^n \frac{1}{(s_{k+1} - s_k)^{1/2}} \, ds \\ &= \sqrt{\pi} \cdot E_{1/2, 1/2}(C\sqrt{\pi}), \end{aligned}$$

where ds is the Lebesgue measure on \mathcal{S}^n and

$$E_{\alpha, \beta}(z) := \sum_{k=0}^{\infty} \frac{z^k}{\Gamma(\beta + \alpha k)}, \quad z \in \mathbb{C}, \alpha, \beta > 0 \quad (\text{E.12})$$

is the Mittag–Leffler function ([Erdélyi et al., 1955](#)). When α and β are both real and positive, the series in (E.12) converges for all values of $z \in \mathbb{C}$, which completes the proof. \blacksquare

Appendix F. Proof of Lemma 7

For each $t \geq 0$, let $N_t := \max\{k : S_k < t \leq S_{k+1}\}$. Then $N_1 = N$ and $T_n = S_n$ for $n \leq N$. Moreover, $\{N_t\}_{t \geq 0}$ is a renewal process with renewal times $\{S_k\}_{k \geq 0}$ and i.i.d. interrenewal times with pdf f_τ . The moment-generating function of M_t can be upper-bounded as follows ([Glynn and Whitt, 1994](#)):

$$\mathbf{E}[e^{\theta N_t}] \leq 1 + e^\theta \sum_{k=0}^{\infty} e^{\theta k} \mathbf{P}[S_k < t]. \quad (\text{F.1})$$

Let $t = 1$ and fix some $\beta > 0$. Then

$$\begin{aligned} \sum_{k=0}^{\infty} e^{\theta k} \mathbf{P}[S_k < 1] &= \sum_{k \leq \beta} e^{\theta k} \mathbf{P}[S_k < 1] + \sum_{k > \beta} e^{\theta k} \mathbf{P}[S_k < 1] \\ &\leq (\beta + 1)e^{\theta \beta} + \sum_{k=0}^{\infty} e^{\theta k} \mathbf{P}[S_k < k\beta^{-1}]. \end{aligned}$$

Using Markov’s inequality and the fact that the τ_i ’s are i.i.d., we can further estimate

$$\mathbf{P}[S_k < k\beta^{-1}] = \mathbf{P}[k - \beta S_k > 0] \leq e^k \mathbf{E} \left[e^{-\beta S_k} \right] = \left(e M_\tau(-\beta) \right)^k.$$

Substituting these estimates into (F.1) and optimizing over β , we get (4.6).

Acknowledgments

The authors would like to thank Matus Telgarsky for many enlightening discussions. This work was supported in part by the NSF CAREER award CCF-1254041, in part by the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370, in part by the Center for Advanced Electronics through Machine Learning (CAEML) I/UCRC award no. CNS-16-24811, and in part by the Office of Naval Research under grant no. N00014-12-1-0998.

References

- Radosław Adamczak. A tail inequality for suprema of unbounded empirical processes with applications to Markov chains. *Electronic Journal of Probability*, 13:1000–1034, 2008.
- Patrik Andersson and Arturo Kohatsu-Higa. Unbiased simulation of stochastic differential equations using parametrix expansions. *Bernoulli*, 23(3):2028–2057, 2017.
- Vlad Bally and Arturo Kohatsu-Higa. A probabilistic interpretation of the parametrix method. *Annals of Applied Probability*, 25(6):3095–3138, 2015.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013.
- Michelle Boué and Paul Dupuis. A variational representation for certain functionals of Brownian motion. *Annals of Probability*, 26(4):1641–1659, 1998.
- Sébastien Bubeck, Ronen Eldan, and Joseph Lehec. Sampling from a log-concave distribution with projected Langevin Monte Carlo. *Discrete and Computational Geometry*, 57:757–783, 2018.
- Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K. Duvenaud. Neural ordinary differential equations. In *Neural Information Processing Systems*, 2018.
- Paolo Dai Pra. A stochastic control approach to reciprocal diffusion processes. *Applied Mathematics and Optimization*, 23:313–329, 1991.
- Ronen Eldan and James R. Lee. Regularization under diffusion and anticoncentration of the information content. *Duke Mathematical Journal*, 167(5):969–993, 2018.
- Ronen Eldan and Ohad Shamir. The power of depth for feedforward neural networks. In *Proceedings of the 29th Annual Conference on Learning Theory*, pages 907–940, 2016.
- Arthur Erdélyi, Wilhelm Magnus, Fritz Oberhettinger, and Francesco G. Tricomi. *Higher Transcendental Functions*, volume III. McGraw-Hill, New York, 1955.

- Wendell H. Fleming. Exit probabilities and optimal stochastic control. *Applied Mathematics and Optimization*, 4:329–346, 1978.
- Wendell H. Fleming and Raymond W. Rishel. *Deterministic and Stochastic Optimal Control*. Springer, 1975.
- Wendell H. Fleming and Sheunn-Jyi Sheu. Stochastic variational formula for fundamental solutions of parabolic PDE. *Applied Mathematics and Optimization*, 13:193–204, 1985.
- Hans Föllmer. An entropy approach to time reversal of diffusion processes. In *Stochastic Differential Systems (Marseille-Luminy, 1984)*, volume 69 of *Lecture Notes in Control and Information Sciences*. Springer, 1985.
- Evarist Giné and Richard Nickl. *Mathematical Foundations of Infinite-Dimensional Statistical Models*. Cambridge University Press, 2016.
- Peter W. Glynn and Ward Whitt. Large deviations behavior of counting processes and their inverses. *Queueing Systems*, 17:107–128, 1994.
- Carl Graham and Denis Talay. *Stochastic Simulation and Monte Carlo Methods: Mathematical Foundations of Stochastic Simulation*, volume 68 of *Stochastic Modeling and Applied Probability*. Springer, 2013.
- Andreas Griewank and Andrea Walther. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. SIAM, 2nd edition, 2008.
- Tatsunori Hashimoto, David Gifford, and Tommi Jaakkola. Learning population-level diffusions with generative RNNs. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 2417–2426, 2016.
- Pierre Henry-Labordère, Xiaolu Tian, and Nizar Touzi. Unbiased simulation of stochastic differential equations. *Annals of Applied Probability*, 27(6):3305–3341, 2017.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural Networks*, 3(5):551–560, 1990.
- Benton Jamison. The Markov processes of Schrödinger. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 32(4):323–331, 1975.
- Olav Kallenberg. *Foundations of Modern Probability*. Springer, 2nd edition, 2002.
- Vladimir Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*, volume 2033 of *Lecture Notes in Mathematics*. Springer, 2011.
- Joseph Lehec. Representation formula for the entropy and functional inequalities. *Annales de l’Institut Henri Poincaré – Probabilités et Statistiques*, 49(3):885–899, 2013.
- Qianxiao Li, Cheng Tai, and Weinan E. Stochastic modified equations and adaptive stochastic algorithms. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2101–2110, 2017.

- Qianxiao Li, Long Chen, Cheng Tai, and Weinan E. Maximum principle based algorithms for deep learning. *Journal of Machine Learning Research*, 18:1–29, 2018.
- Xin Li. Simultaneous approximation of multivariate functions and their derivatives by neural networks with one hidden layer. *Neurocomputing*, 12:327–343, 1996.
- Stephan Mandt, Matthew D. Hoffman, and David M. Blei. Stochastic gradient descent as approximate Bayesian inference. *Journal of Machine Learning Research*, 18:1–35, 2017.
- Javier R. Movellan, Paul Mineiro, and Ruth J. Williams. A Monte Carlo EM approach for partially observable diffusion processes: theory and applications to neural networks. *Neural Computation*, 14:1507–1544, 2002.
- Michele Pavon. Stochastic control and nonequilibrium thermodynamical systems. *Applied Mathematics and Optimization*, 19:187–202, 1989.
- Philip E. Protter. *Stochastic Integration and Differential Equations*. Springer, 2nd edition, 2005.
- Maxim Raginsky, Alexander Rakhlin, and Matus Telgarsky. Non-convex learning via Stochastic Gradient Langevin Dynamics: a nonasymptotic analysis. In *Proceedings of the 2017 Conference on Learning Theory*, 2017.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 2014 International Conference on Machine Learning*, pages 1278–1286, 2014.
- Tom Ryder, Andrew Golightly, A. Steven McGough, and Dennis Prangle. Black-box variational inference for stochastic differential equations. In *Proceedings of the 35th International Conference on Machine Learning*, pages 4423–4432, 2018.
- Itay Safran and Ohad Shamir. Depth-width tradeoffs in approximating natural functions with neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2979–2987, 2017.
- Erwin Schrödinger. Über die Umkehrung der Naturgesetze. *Sitzung ber Preuss. Akad. Wissen., Berlin Phys. Math.*, 144, 1931.
- Sheunn-Jyi Sheu. Some estimates of the transition density of a nondegenerate diffusion Markov process. *Annals of Probability*, 19(2):538–561, 1991.
- Daniel W. Stroock. *An Introduction to Partial Differential Equations for Probabilists*. Cambridge University Press, 2008.
- Matus Telgarsky. Neural networks and rational functions. In *Proceedings of the 34th International Conference on Machine Learning*, pages 3387–3393, 2017.
- Andre Wibisono, Ashia C. Wilson, and Michael I. Jordan. A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences (U.S.)*, 113(47):E7351–E7358, 2016.

Lin Yang, Raman Arora, Vladimir Braverman, and Tuo Zhao. The physical systems behind optimization algorithms. In *Neural Information Processing Systems*, 2018.

Dmitry Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114, 2017.

Joseph E. Yukich, Maxwell B. Stinchcombe, and Halbert White. Sup-norm approximation bounds for neural networks through probabilistic methods. *IEEE Transactions on Information Theory*, 41(4):1021–1027, July 1995.

Yuchen Zhang, Percy Liang, and Moses Charikar. A hitting time analysis of stochastic gradient Langevin dynamics. In *Proceedings of the 2017 Conference on Learning Theory*, pages 1980–2022, 02 2017.