# SCORE-BASED PARAMETER ESTIMATION FOR A CLASS OF CONTINUOUS-TIME STATE SPACE MODELS 

ALEXANDROS BESKOS* ${ }^{*}$ DAN CRISAN ${ }^{\dagger}$, AJAY JASRA ${ }^{\ddagger}$, NIKOLAS KANTAS ${ }^{\dagger}$, AND HAMZA<br>RUZAYQAT ${ }^{\ddagger}$


#### Abstract

We consider the problem of parameter estimation for a class of continuous-time state space models (SSMs). In particular, we explore the case of a partially observed diffusion, with data also arriving according to a diffusion process. Based upon a standard identity of the score function, we consider two particle filter based methodologies to estimate the score function. Both methods rely on an online estimation algorithm for the score function, as described, e.g., in [13], of $\mathcal{O}\left(N^{2}\right)$ cost, with $N \in \mathbb{N}$ the number of particles. The first approach employs a simple Euler discretization and standard particle smoothers and is of cost $\mathcal{O}\left(N^{2}+N \Delta_{l}^{-1}\right)$ per unit time, where $\Delta_{l}=2^{-l}, l \in \mathbb{N}_{0}$, is the time-discretization step. The second approach is new and based upon a novel diffusion bridge construction. It yields a new backward type Feynman-Kac formula in continuous-time for the score function and is presented along with a particle method for its approximation. Considering a time-discretization, the cost is $\mathcal{O}\left(N^{2} \Delta_{l}^{-1}\right)$ per unit time. To improve computational costs, we then consider multilevel methodologies for the score function. We illustrate our parameter estimation method via stochastic gradient approaches in several numerical examples.


Key words. Score Function, Parameter Estimation, Particle Filter, Diffusion Bridges.
AMS subject classifications. 65C05, 65C35, 60G35, 60J60, 60J65, 60H10, 60H35, 91G60

1. Introduction. We consider the problem of parameter estimation for continuous-time SSMs. These are models comprising stochastic differential equations (SDEs) describing a hidden dynamic state and their observations. Such models are ubiquitous in a large number of practical applications in science, engineering, finance, and economics, see [11] for an overview. Inference in SSMs, also known as hidden Markov models, hinges upon computing conditional probability distributions of the dynamic hidden state given the acquired observations and unknown static parameters. This is referred to as the stochastic filtering problem, which is in general intractable, but reliable numerical approximations are routinely available [1, 11]. The problem of inferring the unknown parameters is more challenging. In this paper we focus on maximum likelihood inference and gradient methods that are performed in an online manner. The offline or iterative case given a fixed batch of observations can also be treated using our proposed methods. The approach considered in this article is to make use of the gradient of the log-likelihood, commonly referred to as the score function, within a stochastic gradient algorithm (see e.g. [13, 14]). Intrinsically, there are several challenges arising with such an approach. Firstly, when one adopts a continuous-time model and assumes access to arbitrarily high frequency observations, one does not observe in practice truly in continuous-time, therefore some sort of time-discretization is required. Secondly, in both discrete-time and continuous-time formulations, there are very few cases when the score function is analytically available. Both these issues imply that numerical approximations are required.

We consider two different approaches for the numerical approximation of the score function. The first is to simply time-discretize a representation of the score function and then apply discretetime numerical approximation schemes [13, 14, 23]. The second, is to develop a numerical approx-

[^0]imation scheme directly on continuous-time path-space to estimate the score function and then (necessarily) discretize the algorithm in time. The order of designing the estimation method and time-discretization can be rather important. Often the second approach is preferable in terms of both performance and robustness as the discretization mesh vanishes, see e.g. [25, 24]. We then use the score estimate for implementing recursive maximum likelihood, where the parameters are updated at unit time intervals $([21,27])$. The particular choice of time interval length is without loss of generality, and allows the score to accumulate sufficient information from the observations before updating the parameters.

In the first approach, we consider a well-known expression for the score function, for instance as given in [8]. Given this formula, one can produce an Euler discretized version of the score and work in discrete-time but with high frequency. The score is an expectation of an additive functional of the hidden state path conditioned on the available observations, which is commonly referred to as the smoothing distribution. In this context, many well-known particle smoothing schemes can now be adopted, such as the ones described in [13, 23]. These latter approaches are simulation-based schemes whose convergence is based upon the number of samples $N \in \mathbb{N}$. We prove some technical results for the discretized problem, which together with the work in [13] allow us conjecture that to obtain a mean square error (MSE) of $\mathcal{O}\left(\epsilon^{2}\right)$, for given $\epsilon>0$, we require a computational cost of $\mathcal{O}\left(\epsilon^{-4}\right)$ per unit time. The latter derives from an algorithmic cost of $\mathcal{O}\left(N^{2}+N \Delta_{l}^{-1}\right)$ per unit time, where $\Delta_{l}=2^{-l}$ is the time-discretization. As we explain later in the article, this corresponds to a best case scenario, due to the intrinsic nature of the algorithm. In particular, we start with a continuous-time formula and time-discretize it, but the deduced numerical algorithm can be problematic in terms of computational complexity as $l \rightarrow \infty$. Whilst in some scenarios one does not observe any issues, examples can be found where the variance of the method can explode as $l$ grows ([36]), thus putting into question the validity of the conjecture on the cost to achieve an MSE of $\mathcal{O}\left(\epsilon^{2}\right)$.

This motivates the introduction of our second approach, where we build upon a change of measure technique proposed in [29]. This latter approach has so far been used in very different contexts than the present paper, namely related to discretely observed diffusions for Bayesian inference and Markov chain Monte Carlo (MCMC) [34, 35] or smoothing for potentially non-linear observation functions [22]. Our approach is a data augmentation scheme, whereby, at unit time intervals, the end points for the hidden state are sampled and the path is connected using diffusion bridges. Then, starting again from the formula for the score function in [8], we will use this change of measure associated to a diffusion bridge and its driving Brownian noise; see also [36] where a related approach is used for a different class of problems. Based upon this change of measure we develop a new backward type Feynman-Kac formula in continuous-time. This new formula facilitates an adaptation of the method in [13] in true continuous-time, albeit one cannot apply it in practice. We time-discretize the algorithm and conjecture that to obtain a MSE of $\mathcal{O}\left(\epsilon^{2}\right)$, for given $\epsilon>0$, we need a cost of $\mathcal{O}\left(\epsilon^{-6}\right)$ per unit time, which derives from an algorithmic cost of $\mathcal{O}\left(N^{2} \Delta_{l}^{-1}\right)$ per unit time. We note however, that this computational complexity will not explode with increasing $l$ as may be the case in the first approach. To improve the cost required for a given MSE, we develop a novel multilevel Monte Carlo extension that can, in some cases, achieve a MSE of $\mathcal{O}\left(\epsilon^{2}\right)$ at a cost per unit time of $\mathcal{O}\left(\epsilon^{-4}\right)$. We remark that although our MSE-cost statements are based upon conjectures, they are verified numerically. Direct proofs of these require substantial technical results that will be the topic of future work.
1.1. Contributions and Organization. We conclude this introduction by emphasizing that our contributions are aimed to deal with both continuous-time observations and hidden states.

As mentioned earlier this poses very particular challenges relative to earlier works that deal with filtering and smoothing when discrete-time observations/models are used as in [7, 15, 16, 22, 35, 28]. None of these works look at continuous-time observations. Similarly online likelihood estimation of the parameters using the score function is considered in $[15,16]$ only for the case of discrete-time observations of hidden diffusions.

The contributions of this paper can be summarized as follows:

- We investigate the efficiency and accuracy of two fundamentally different numerical approximations of the score function on its own and when used for the purpose of recursive maximum likelihood. Both methods rely on fairly standard tools such as changes of measure, particle smoothing and Euler time-discretization.
- We provide a detailed discussion on the computational complexity of each method. We illustrate the performance and computational cost for several models in numerical examples that consider estimation of the score function and parameter estimation.
- The second approach is a novel method that operates directly on the path-space. The approach improves performance and is robust to arbitrarily small time-discretization at the expense of additional computational cost. The latter is reduced using a new Multilevel Particle Filter; see [17, 20] for some existing approaches.
This article is structured as follows. In Section 2 the basic problem is formulated in continuoustime. In Section 3 we consider our first method for online score estimation and explain the various features associated to it. In Section 4 our second method for online score estimation is developed. In Section 5 our numerical results are presented.
1.2. Notation. Let $(X, \mathcal{X})$ be a measurable space. We write $\mathcal{B}_{b}(X)$ for the set of bounded measurable functions, $\varphi: \mathrm{X} \rightarrow \mathbb{R}^{d}, d \in \mathbb{N}$, and $\mathcal{C}(\mathrm{X})$ for the continuous ones. Let $\varphi: \mathbb{R}^{d} \rightarrow \mathbb{R}$; $\operatorname{Lip}_{\|\cdot\|_{2}}\left(\mathbb{R}^{d}\right)$ denotes the collection of real-valued functions that are Lipschitz w.r.t. the Euclidean distance $\|\cdot\|_{2}$. That is, $\varphi \in \operatorname{Lip}_{\|\cdot\|_{2}}\left(\mathbb{R}^{d}\right)$ if there exists a $C<+\infty$ such that for any $(x, y) \in \mathbb{R}^{2 d}$,

$$
|\varphi(x)-\varphi(y)| \leq C\|x-y\|_{2}
$$

$\mathcal{N}_{s}(\mu, \Sigma)$ denotes an $s$-dimensional Gaussian law of mean $\mu$ and covariance $\Sigma$; if $s=1$ we omit subscript $s$. For a vector/matrix $X, X^{*}$ denotes the transpose of $X$. For $A \in \mathcal{X}, \delta_{A}(d u)$ denotes the Dirac measure of $A$, and if $A=\{x\}$ with $x \in X$, we write $\delta_{x}(d u)$. For a vector-valued function in $d$-dimensions $\varphi(x)$ (resp. $d$-dimensional vector $x$ ) we write the $i^{\text {th }}$-component, $1 \leq i \leq d$, as $\varphi^{(i)}(x)$ (resp. $x^{i}$ ). For a $d \times q$ matrix $x$ we write the $(i, j)^{\text {th }}$-entry as $x^{(i j)} . \mathbb{N}=\{1,2, \ldots\}$ and $\mathbb{N}_{0}=\mathbb{N} \cup\{0\}$.

## 2. Problem Formulation.

2.1. Preliminaries. We consider the parameter space $\theta \in \Theta \subset \mathbb{R}^{d_{\theta}}$, with $\Theta$ being compact, $d_{\theta} \in \mathbb{N}$. The stochastic processes $\left\{Y_{t}\right\}_{t \geq 0},\left\{X_{t}\right\}_{t \geq 0}$ of interest are defined upon the probability triple $\left(\Omega, \mathcal{F}, \mathbb{P}_{\theta}\right)$, with $Y_{t} \in \mathbb{R}^{d_{y}}, X_{t} \in \mathbb{R}^{d_{x}}, d_{y}, d_{x} \in \mathbb{N}$, initial conditions $X_{0}=x_{*} \in \mathbb{R}^{d_{x}}, Y_{0}=y_{*} \in \mathbb{R}^{d_{y}}$, and are determined as the solution of the system of SDEs:

$$
\begin{align*}
d Y_{t} & =h_{\theta}\left(X_{t}\right) d t+d B_{t}  \tag{2.1}\\
d X_{t} & =b_{\theta}\left(X_{t}\right) d t+\sigma\left(X_{t}\right) d W_{t} \tag{2.2}
\end{align*}
$$

Here, for each $\theta \in \Theta, h_{\theta}: \mathbb{R}^{d_{x}} \rightarrow \mathbb{R}^{d_{y}}, b_{\theta}: \mathbb{R}^{d_{x}} \rightarrow \mathbb{R}^{d_{x}}, \sigma: \mathbb{R}^{d_{x}} \rightarrow \mathbb{R}^{d_{x} \times d_{x}}$, with $\sigma$ being of full rank, and $\left\{B_{t}\right\}_{t \geq 0},\left\{W_{t}\right\}_{t \geq 0}$ are independent standard Brownian motions of dimension $d_{y}, d_{x}$ respectively.

To minimize technical difficulties, the following assumptions are made throughout the paper:
(D1) (i) $\sigma$ is continuous, bounded; $a(x):=\sigma(x) \sigma(x)^{*}$ is uniformly elliptic;
(ii) for each $\theta, h_{\theta}$ and $b_{\theta}$ are bounded, measurable; $h_{\theta}^{(i)} \in \operatorname{Lip}_{\|\cdot\|_{2}}\left(\mathbb{R}^{d_{x}}\right), 1 \leq i \leq d_{y}$;
(iii) the gradients $\nabla_{\theta} h_{\theta}: \mathbb{R}^{d_{x}} \rightarrow \mathbb{R}^{d_{y} \times d_{\theta}}$ and $\nabla_{\theta} b_{\theta}: \mathbb{R}^{d_{x}} \rightarrow \mathbb{R}^{d_{x} \times d_{\theta}}$ exist, and are continuous, bounded, measurable; $\nabla_{\theta} h_{\theta}^{(i j)} \in \operatorname{Lip}_{\|\cdot\|_{2}}\left(\mathbb{R}^{d_{x}}\right), 1 \leq i \leq d_{y}, 1 \leq j \leq d_{\theta}$;
(iv) let $\phi_{\theta}(x)=\left(\nabla_{\theta} b_{\theta}(x)\right)^{*} a(x)^{-1} \sigma(x)$; for any $\theta, \phi_{\theta}^{(i j)} \in \operatorname{Lip}_{\|\cdot\|_{2}}\left(\mathbb{R}^{d_{x}}\right), 1 \leq i \leq d_{\theta}$, $1 \leq j \leq d_{x}$.
We introduce the probability measure $\overline{\mathbb{P}}_{\theta}$, defined via the Radon-Nikodym derivative:

$$
\begin{equation*}
Z_{t, \theta}:=\left.\frac{d \mathbb{P}_{\theta}}{d \overline{\mathbb{P}}_{\theta}}\right|_{\mathcal{F}_{t}}=\exp \left\{\int_{0}^{t} h_{\theta}\left(X_{s}\right)^{*} d Y_{s}-\frac{1}{2} \int_{0}^{t} h_{\theta}\left(X_{s}\right)^{*} h_{\theta}\left(X_{s}\right) d s\right\} \tag{2.3}
\end{equation*}
$$

with $\mathcal{F}_{t}=\sigma\left(\left\{X_{s}, Y_{s}\right\}_{0 \leq s \leq t}\right)$. Henceforth, $\overline{\mathbb{E}}_{\theta}$ denotes expectation w.r.t. $\overline{\mathbb{P}}_{\theta}$, so that under $\overline{\mathbb{P}}_{\theta}$, the process $\left\{X_{t}\right\}_{t \geq 0}$ follows the dynamics in (2.2), whereas $\left\{Y_{t}\right\}_{t \geq 0}$ is a Brownian motion independent of $\left\{X_{t}\right\}_{t \geq 0}$. We define, for $\varphi \in \mathcal{B}_{b}\left(\mathbb{R}^{d_{x}}\right)$ :

$$
\gamma_{t, \theta}(\varphi):=\overline{\mathbb{E}}_{\theta}\left[\varphi\left(X_{t}\right) Z_{t, \theta} \mid \mathcal{Y}_{t}\right]
$$

where $\mathcal{Y}_{t}$ is the filtration generated by the process $\left\{Y_{s}\right\}_{0 \leq s \leq t}$. Our objective is to produce estimates of the gradient of the score function $\nabla_{\theta} \log \left(\gamma_{T, \theta}(1)\right)$.

REmARK 2.1. To connect the changes of measures with standard likelihood derivations, notice that - via Girsanov's theorem $-Z_{t, \theta}$ is the density (w.r.t. to a Wiener measure) of the distribution of $\left\{Y_{s}\right\}_{0 \leq s \leq t}$ conditionally on $\left\{X_{s}\right\}_{0 \leq s \leq t}$. Then, $\gamma_{T, \theta}(1)$ integrates out $\left\{X_{t}\right\}_{0 \leq t \leq T}$, thus corresponds to the marginal density - i.e., the likelihood - of the observations $\left\{Y_{t}\right\}_{0 \leq t \leq T}$.

In our setting, the score function writes as (see e.g. [8]):

$$
\begin{equation*}
\nabla_{\theta} \log \left(\gamma_{T, \theta}(1)\right)=\frac{\overline{\mathbb{E}}_{\theta}\left[\lambda_{T, \theta} Z_{T, \theta} \mid \mathcal{Y}_{T}\right]}{\overline{\mathbb{E}}_{\theta}\left[Z_{T, \theta} \mid \mathcal{Y}_{T}\right]} \tag{2.4}
\end{equation*}
$$

where we have defined:

$$
\begin{aligned}
& \lambda_{T, \theta}:=\int_{0}^{T}\left(\nabla_{\theta} b_{\theta}\left(X_{t}\right)\right)^{*} a\left(X_{t}\right)^{-1} \sigma\left(X_{t}\right) d W_{t} \\
&+\int_{0}^{T}\left(\nabla_{\theta} h_{\theta}\left(X_{t}\right)\right)^{*} d Y_{t}-\int_{0}^{T}\left(\nabla_{\theta} h_{\theta}\left(X_{t}\right)\right)^{*} h_{\theta}\left(X_{t}\right) d t
\end{aligned}
$$

For completeness, a derivation of (2.4) can be found in [5, Appendix] and the supplementary material of this article. We remark that one can derive a formula for the score function when $\sigma$ depends upon $\theta$, which is given in Section 4. We will assume throughout that $T \in \mathbb{N}$. Note also that an application of Bayes' rule gives that, almost surely:

$$
\begin{equation*}
\frac{\overline{\mathbb{E}}_{\theta}\left[\lambda_{T, \theta} Z_{T, \theta} \mid \mathcal{Y}_{T}\right]}{\overline{\mathbb{E}}_{\theta}\left[Z_{T, \theta} \mid \mathcal{Y}_{T}\right]}=\mathbb{E}_{\theta}\left[\lambda_{T, \theta} \mid \mathcal{Y}_{T}\right] \tag{2.5}
\end{equation*}
$$

where $\mathbb{E}_{\theta}$ denotes expectation w.r.t. $\mathbb{P}_{\theta}$.
2.2. Parameter Estimation. In the offline case suppose one has obtained data $\left\{Y_{t}\right\}_{0 \leq t \leq T}$. Then it is possible to perform standard gradient descent using (2.5) and updating $\theta$ iteratively:

$$
\begin{equation*}
\theta^{m+1}=\theta^{m}+\alpha_{m} \mathbb{E}_{\theta^{m}}\left[\lambda_{T, \theta^{m}} \mid \mathcal{Y}_{T}\right], \tag{2.6}
\end{equation*}
$$

where $\alpha_{m} \in \mathbb{R}^{+}, m \in \mathbb{N}_{0}$, are decreasing step-sizes. Instead here we will mainly focus on an online gradient estimation procedure. To obtain this one can aim to maximise the following limiting average log-likelihood,

$$
\mathcal{L}(\theta)=\lim _{t \rightarrow \infty} \frac{1}{t} \int_{0}^{t} \log \gamma_{s, \theta}(1) d s
$$

Let the filter be denoted as $\pi_{t, \theta}(\varphi)=\mathbb{E}_{\theta}\left[\varphi\left(X_{t}\right) \mid \mathcal{Y}_{s}\right]$ and using standard arguments (e.g. Lemma 3.29 p. 67 [1]) one can re-write $\gamma_{t, \theta}(1)$ as

$$
\log \gamma_{t, \theta}(1)=\int_{0}^{t} \pi_{s, \theta}(h)^{T} d Y_{s}-\frac{1}{2} \int_{0}^{t} \pi_{s, \theta}(h)^{T} \pi_{s, \theta}(h) d s
$$

Under appropriate stability and regularity conditions for both $Y_{t}$ and $\pi_{t, \theta}$ (see [31] for more details), then both $\mathcal{L}(\theta)$ and $\nabla_{\theta} \mathcal{L}(\theta)$ are ergodic averages. This means one can implement stochastic gradient ascent using either $\nabla_{\theta} \log \gamma_{t, \theta}(1)$ or $\nabla_{\theta} \log \gamma_{t_{n}, \theta}(1)-\nabla_{\theta} \log \gamma_{t_{n-1}, \theta}(1)$ for any $t_{n-1}<t_{n}$ as estimates of $\mathcal{L}(\theta)$. Given an initial $\theta_{0} \in \Theta$, as we obtain the observation path continuously in time, we will update $\theta$ at times $T \in \mathbb{N}$ using the following recursion:

$$
\begin{equation*}
\theta_{T}=\theta_{T-1}+\alpha_{T}\left(\nabla_{\theta} \log \left(\gamma_{T, \theta_{T-1}}(1)\right)-\nabla_{\theta} \log \left(\gamma_{T-1, \theta_{T-1}}(1)\right)\right) \tag{2.7}
\end{equation*}
$$

where, for $T \in \mathbb{N}, \alpha_{T} \in \mathbb{R}^{+}$is a collection of step-sizes that satisfy $\sum_{T \in \mathbb{N}} \alpha_{T}=\infty, \sum_{T \in \mathbb{N}} \alpha_{T}^{2}<\infty$ to ensure convergence of the estimation as $T \rightarrow \infty$; see [3, 21] for details. This scheme can provide an online estimate for the parameter vector as data arrive. Steps are performed at $\mathcal{O}(1)$ times to ensure that enough information has accumulated to update the parameter. The adoption of unit times is made only for notational convenience. As both recursions (2.6) and (2.7) cannot be computed exactly, we focus upon methodologies that approximate the score function $\nabla_{\theta} \log \left(\gamma_{T, \theta}(1)\right)$.

## 3. Direct Feynman-Kac Formulation.

3.1. Discretized Model. In practice, we will have to work with a discretization of the model in (2.1)-(2.2). We assume access to path of the data $\left\{Y_{t}\right\}_{0 \leq t \leq T}$ which is available up-to an (almost) arbitrarily fine level of time discretization. This would be a very finely discretized path, as accessing the actual continuous path of observation is not possible; this point is discussed later on. One could focus on a time-discretization of either side of (2.5), however, as is conventional in the literature (e.g. $[1,20]$ ) we focus on the left hand side.

Let $l \in \mathbb{N}_{0}$ and consider an Euler-Maruyama time-discretization with step-size $\Delta_{l}=2^{-l}$. That is, for $k \in\left\{1,2, \ldots, T / \Delta_{l}\right\}$ :

$$
\begin{equation*}
\widetilde{X}_{k \Delta_{l}}=\widetilde{X}_{(k-1) \Delta_{l}}+b_{\theta}\left(\widetilde{X}_{(k-1) \Delta_{l}}\right) \Delta_{l}+\sigma\left(\widetilde{X}_{(k-1) \Delta_{l}}\right)\left(W_{k \Delta_{l}}-W_{(k-1) \Delta_{l}}\right), \quad \widetilde{X}_{0}=x_{*} . \tag{3.1}
\end{equation*}
$$

Note that the Brownian motion in (3.1) is the same as in (2.2) under both $\mathbb{P}_{\theta}$ and $\overline{\mathbb{P}}_{\theta}$. We set:

$$
\begin{gather*}
\lambda_{T, \theta}^{l}\left(x_{0}, x_{\Delta_{l}}, \ldots, x_{T}\right):=\sum_{k=0}^{T / \Delta_{l}-1}\left\{\left(\nabla_{\theta} b_{\theta}\left(x_{k \Delta_{l}}\right)\right)^{*} a\left(x_{k \Delta_{l}}\right)^{-1} \sigma\left(x_{k \Delta_{l}}\right)\left(W_{(k+1) \Delta_{l}}-W_{k \Delta_{l}}\right)\right. \\
\left.+\left(\nabla_{\theta} h_{\theta}\left(x_{k \Delta_{l}}\right)\right)^{*}\left(Y_{(k+1) \Delta_{l}}-Y_{k \Delta_{l}}\right)-\left(\nabla_{\theta} h_{\theta}\left(x_{k \Delta_{l}}\right)\right)^{*} h_{\theta}\left(x_{k \Delta_{l}}\right) \Delta_{l}\right\} . \tag{3.2}
\end{gather*}
$$

We remark that $\lambda_{T, \theta}^{l}$ is a function also of the observations, but this dependence is suppressed from the notation. For $k \in\left\{0,1, \ldots, T / \Delta_{l}-1\right\}$, we define:

$$
g_{k, \theta}^{l}\left(x_{k \Delta_{l}}\right):=\exp \left\{h_{\theta}\left(x_{k \Delta_{l}}\right)^{*}\left(y_{(k+1) \Delta_{l}}-y_{k \Delta_{l}}\right)-\frac{\Delta_{l}}{2} h_{\theta}\left(x_{k \Delta_{l}}\right)^{*} h_{\theta}\left(x_{k \Delta_{l}}\right)\right\} .
$$

Note that:

$$
\begin{aligned}
Z_{T, \theta}^{l}\left(x_{0}, x_{\Delta_{l}}, \ldots, x_{T}\right) & :=\prod_{k=0}^{T / \Delta_{l}-1} g_{k, \theta}^{l}\left(x_{k \Delta_{l}}\right) \\
& =\exp \left\{\sum_{k=0}^{T / \Delta_{l}-1}\left[h_{\theta}\left(x_{k \Delta_{l}}\right)^{*}\left(y_{(k+1) \Delta_{l}}-y_{k \Delta_{l}}\right)-\frac{\Delta_{l}}{2} h_{\theta}\left(x_{k \Delta_{l}}\right)^{*} h_{\theta}\left(x_{k \Delta_{l}}\right)\right]\right\}
\end{aligned}
$$

is a time-discretization of $Z_{T, \theta}$. We thus obtain the discretized approximation of the score function $\nabla_{\theta} \log \left(\gamma_{T, \theta}(1)\right):$

$$
\begin{equation*}
\nabla_{\theta} \log \left(\gamma_{T, \theta}^{l}(1)\right):=\frac{\overline{\mathbb{E}}_{\theta}\left[\lambda_{T, \theta}^{l}\left(\widetilde{X}_{0}, \widetilde{X}_{\Delta_{l}}, \ldots, \widetilde{X}_{T}\right) Z_{T, \theta}^{l}\left(\widetilde{X}_{0}, \widetilde{X}_{\Delta_{l}}, \ldots, \widetilde{X}_{T}\right) \mid \mathcal{Y}_{T}\right]}{\overline{\mathbb{E}}_{\theta}\left[Z_{T, \theta}^{l}\left(\widetilde{X}_{0}, \widetilde{X}_{\Delta_{l}}, \ldots, \widetilde{X}_{T}\right) \mid \mathcal{Y}_{T}\right]} \tag{3.3}
\end{equation*}
$$

We have the following result which establishes the convergence of our Euler approximation. Below $\|\cdot\|_{2}$ is the $L_{2}-$ norm for vectors. The proof is given in the supplementary material of the article.

Theorem 3.1. Assume (D1). Then for any $(r, T) \in[1, \infty) \times \mathbb{N}$ there exists a $C<+\infty$ such that for any $l \in \mathbb{N}_{0}$

$$
\mathbb{E}_{\theta}\left[\left\|\nabla_{\theta} \log \left(\gamma_{T, \theta}(1)\right)-\nabla_{\theta} \log \left(\gamma_{T, \theta}^{l}(1)\right)\right\|_{2}^{r}\right]^{1 / r} \leq C \Delta_{l}^{1 / 2}
$$

The result is fairly standard, but we note it is not a simple application of results on SDEs in filtering (e.g. [26, 32]) and this is reflected in the proof. The rate of convergence of the approximation will be very relevant for some of our subsequent discussions.
3.2. Backward Feynman-Kac Model and Particle Smoothing. From herein the $\widetilde{X}$ notation is dropped for simplicity. Consider the time interval $[k, k+1]$ and the $k$-th update of (2.7). We define the discrete-time approximation (at level $l$ ) as:

$$
u_{k, l}=\left(x_{k+\Delta_{l}}, \ldots, x_{k+1}\right) \in E_{l}:=\left(\mathbb{R}^{d_{x}}\right)^{\Delta_{l}^{-1}}, \quad k \in \mathbb{N}_{0}
$$

Recall (2.3). A discrete-time approximation of $p_{\theta}\left(\left\{Y_{t}\right\}_{k \leq t \leq k+1} \mid\left\{X_{t}\right\}_{0 \leq t \leq T}\right)$ is:

$$
G_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right):=\prod_{p=0}^{\Delta_{l}^{-1}-1} g_{k+p, \theta}^{l}\left(x_{k+p \Delta_{l}}\right)
$$

where we set $u_{-1, l}=x_{*}$, for each $l \in \mathbb{N}_{0}$. We denote by $m_{\theta}^{l}$ the Euler transition density induced by time-discretisation (3.1), and then write the initial distribution and Markov transition kernel for the discrete-time process with $k \in \mathbb{N}$ as follows:

$$
\begin{aligned}
\eta_{0, \theta}^{l}\left(d u_{0, l}\right) & =\prod_{p=1}^{\Delta_{l}^{-1}} m_{\theta}^{l}\left(x_{(p-1) \Delta_{l}}, x_{p \Delta_{l}}\right) d x_{p \Delta_{l}} \\
M_{\theta}^{l}\left(u_{k-1, l}, d u_{k, l}\right) & =\prod_{p=1}^{\Delta_{l}^{-1}} m_{\theta}^{l}\left(x_{k+(p-1) \Delta_{l}}, x_{k+p \Delta_{l}}\right) d x_{k+p \Delta_{l}}
\end{aligned}
$$

REmARK 3.1. The definition of $G_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right), M_{\theta}^{l}\left(u_{k-1, l}, d u_{k, l}\right)$ implies: i) $G_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right)$ involves $u_{k-1, l}$ only via its very last element, $x_{k}$; ii) the dynamics of $u_{k, l}$ conditionally on $u_{k-1, l}$ depend only on the very last element, $x_{k}$, of $u_{k-1, l}$.
We can now state the discrete-time filtering distribution for $k \in \mathbb{N}_{0}$ :

$$
\begin{equation*}
\pi_{k, \theta}^{l}\left(d\left(u_{0, l}, \ldots, u_{k, l}\right)\right):=\frac{\left(\prod_{p=0}^{k} G_{p, \theta}^{l}\left(u_{p-1, l}, u_{p, l}\right)\right) \eta_{0, \theta}^{l}\left(d u_{0, l}\right) \prod_{p=1}^{k} M_{\theta}^{l}\left(u_{p-1, l}, d u_{p, l}\right)}{\int_{E_{l}^{k+1}}\left(\prod_{p=0}^{k} G_{p, \theta}^{l}\left(u_{p-1, l}, u_{p, l}\right)\right) \eta_{0, \theta}^{l}\left(d u_{0, l}\right) \prod_{p=1}^{k} M_{\theta}^{l}\left(u_{p-1, l}, d u_{p, l}\right)} \tag{3.4}
\end{equation*}
$$

That is, $\pi_{k, \theta}^{l}\left(d\left(u_{0, l}, \ldots, u_{k, l}\right)\right)$ is a discrete-time approximation of the filtering distribution:

$$
\pi_{k, \theta}\left(d\left(\left\{X_{t}\right\}_{0 \leq t \leq k}\right)\right):=\mathbb{P}_{\theta}\left(d\left\{X_{t}\right\}_{0 \leq t \leq k} \mid\left\{Y_{t}\right\}_{0 \leq t \leq k}\right)
$$

Expression (3.4) corresponds to a standard Feynman-Kac model (see e.g. [12]), thus one can approximate the involved filtering distributions via the corresponding Monte Carlo methodology.

We develop a Monte Carlo method for the approximation of the discretised score function in (3.3). This is accomplished by presenting a backward formula for (3.3). We define for any $p \in \mathbb{N}_{0}$ :

$$
\begin{aligned}
& f_{\theta}^{l}\left(x_{p \Delta_{l}}, x_{(p+1) \Delta_{l}}\right):=\left(\nabla_{\theta} b_{\theta}\left(x_{p \Delta_{l}}\right)\right)^{*} a\left(x_{p \Delta_{l}}\right)^{-1} \sigma\left(x_{p \Delta_{l}}\right)\left(W_{(p+1) \Delta_{l}}-W_{p \Delta_{l}}\right) \\
&+\left(\nabla_{\theta} h_{\theta}\left(x_{p \Delta_{l}}\right)\right)^{*}\left(Y_{(p+1) \Delta_{l}}-Y_{p \Delta_{l}}\right)-\left(\nabla_{\theta} h_{\theta}\left(x_{p \Delta_{l}}\right)\right)^{*} h_{\theta}\left(x_{p \Delta_{l}}\right) \Delta_{l}
\end{aligned}
$$

and let:

$$
\begin{align*}
\Lambda_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) & :=\sum_{p=0}^{\Delta_{l}^{-1}-1} f_{\theta}^{l}\left(x_{k+p \Delta_{l}}, x_{k+(p+1) \Delta_{l}}\right) \\
F_{T, \theta}^{l}\left(u_{0, l}, \ldots, u_{T-1, l}\right) & :=\sum_{k=0}^{T-1} \Lambda_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right)\left(\equiv \lambda_{T, \theta}^{l}\left(x_{0}, x_{\Delta_{l}}, \ldots, x_{T}\right)\right) \tag{3.5}
\end{align*}
$$

for $\lambda_{T, \theta}^{l}\left(x_{0}, x_{\Delta_{l}}, \ldots, x_{T}\right)$ as defined in (3.2) and used in the score function approximation (3.3). Thus:

$$
\begin{equation*}
\nabla_{\theta} \log \left(\gamma_{T, \theta}^{l}(1)\right)=\int_{E_{l}^{T}} F_{T, \theta}^{l}\left(u_{0, l}, \ldots, u_{T-1, l}\right) \mathbb{Q}_{T-1, \theta}^{l}\left(d\left(u_{0, l}, \ldots, u_{T-1, l}\right)\right) \tag{3.6}
\end{equation*}
$$

where $\mathbb{Q}_{T-1, \theta}^{l}\left(d\left(u_{0, l}, \ldots, u_{T-1, l}\right)\right)$ is a time-discretisation of the smoothing law:

$$
\mathbb{Q}_{T-1, \theta}\left(d\left\{X_{t}\right\}_{0 \leq t \leq T}\right):=\mathbb{P}_{\theta}\left(d\left\{X_{t}\right\}_{0 \leq t \leq T} \mid\left\{Y_{t}\right\}_{0 \leq t \leq T}\right)
$$

Now, by the time-reversal formula for hidden Markov models (see e.g. [13, 14]) one has:

$$
\mathbb{Q}_{T-1, \theta}^{l}\left(d\left(u_{0, l}, \ldots, u_{T-1, l}\right)\right):=\pi_{T-1, \theta}^{l}\left(d u_{T-1, l}\right) \prod_{k=1}^{T-1} B_{k, \theta, \pi_{k-1, \theta}^{l}}^{l}\left(u_{k, l}, d u_{k-1, l}\right)
$$

for the backward Markov kernel:

$$
\begin{equation*}
B_{k, \theta, \pi_{k-1, \theta}^{l}}^{l}\left(u_{k, l}, d u_{k-1, l}\right):=\frac{\pi_{k-1, \theta}^{l}\left(d u_{k-1, l}\right) G_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) m_{\theta}^{l}\left(u_{k-1, l}, u_{k, l}\right)}{\pi_{k-1, \theta}^{l}\left(G_{k, \theta}^{l}\left(\cdot, u_{k, l}\right) m_{\theta}^{l}\left(\cdot, u_{k, l}\right)\right)} \tag{3.7}
\end{equation*}
$$

under the standard notation:

$$
\pi_{k-1, \theta}^{l}\left(G_{k, \theta}^{l}\left(\cdot, u_{k, l}\right) m_{\theta}^{l}\left(\cdot, u_{k, l}\right)\right)=\int_{E_{l}} \pi_{k-1, \theta}^{l}\left(d u_{k-1, l}\right) G_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) m_{\theta}^{l}\left(u_{k-1, l}, u_{k, l}\right)
$$

REmark 3.2. The model structure gives important cancellations in (3.7), so that:

$$
B_{k, \theta, \pi_{k-1, \theta}^{l}}^{l}\left(u_{k, l}, d u_{k-1, l}\right) \equiv \frac{\pi_{k-1, \theta}^{l}\left(d u_{k-1, l}\right) g_{k, \theta}^{l}\left(x_{k}\right) m_{\theta}^{l}\left(x_{k}, x_{k+\Delta_{l}}\right)}{\int_{E_{l}} \pi_{k-1, \theta}^{l}\left(d u_{k-1, l}\right) g_{k, \theta}^{l}\left(x_{k}\right) m_{\theta}^{l}\left(x_{k}, x_{k+\Delta_{l}}\right)} .
$$

The objective now is to approximate the right hand side of (3.6) using particle approximations. Our online particle approximation of the gradient of the log-likelihood in (3.6), for a given $l \in \mathbb{N}_{0}$ is presented in Algorithm 3.1. Our estimates are given in (3.8) and (3.10) in Algorithm 3.1. The approach is the method introduced in [13, 14].

## Algorithm 3.1 Online Score Function Estimation for a given $l \in \mathbb{N}_{0}$.

1. For $i \in\{1, \ldots, N\}$, sample $u_{0, l}^{i}$ i.i.d. from $\eta_{0, \theta}^{l}(\cdot)$. The estimate of $\nabla_{\theta} \log \left(\gamma_{1, \theta}^{l}(1)\right)$ is:

$$
\begin{equation*}
\left.\nabla_{\theta} \widehat{\log \left(\gamma_{1, \theta}^{l}\right.}(1)\right):=\frac{\sum_{i=1}^{N} G_{0, \theta}^{l}\left(x_{*}, u_{0, l}^{i}\right) \Lambda_{0, \theta}^{l}\left(x_{*}, u_{0, l}^{i}\right)}{\sum_{i=1}^{N} G_{0, \theta}^{l}\left(x_{*}, u_{0, l}^{i}\right)} \tag{3.8}
\end{equation*}
$$

Set $k=1$, and for $i \in\{1, \ldots, N\}, \breve{u}_{-1, l}^{i}=x_{*}$.
2. For $i \in\{1, \ldots, N\}$ sample $\breve{u}_{k-1, l}^{i}$ from:

$$
\sum_{i=1}^{N} \frac{G_{k-1, \theta}^{l}\left(\check{u}_{k-2, l}^{i}, u_{k-1, l}^{i}\right)}{\sum_{j=1}^{N} G_{k-1, \theta}^{l}\left(\check{u}_{k-2, l}^{j}, u_{k-1, l}^{j}\right)} \delta_{\left\{u_{k-1, l}^{i}\right\}}(\cdot)
$$

If $k=1$, for $i \in\{1, \ldots, N\}$, set $F_{k-1, \theta}^{l, N}\left(\breve{u}_{0, l}^{i}\right)=\Lambda_{0, \theta}^{l}\left(x_{*}, \check{u}_{0, l}^{i}\right)$.
3. For $i \in\{1, \ldots, N\}$, sample $u_{k, l}^{i}$ from $M_{\theta}^{l}\left(\check{u}_{k-1, l}^{i}, \cdot\right)$. For $i \in\{1, \ldots, N\}$, compute:

$$
\begin{equation*}
F_{k, \theta}^{l, N}\left(u_{k, l}^{i}\right)=\frac{\sum_{j=1}^{N} g_{k, \theta}^{l}\left(\check{x}_{k}^{j}\right) m_{\theta}^{l}\left(\check{x}_{k}^{j}, x_{k+\Delta_{l}}^{i}\right)\left\{F_{k-1, \theta}^{l, N}\left(\check{u}_{k-1, l}^{j}\right)+\Lambda_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{j}, u_{k, l}^{i}\right)\right\}}{\sum_{j=1}^{N} g_{k, \theta}^{l}\left(\check{x}_{k}^{j}\right) m_{\theta}^{l}\left(\check{x}_{k}^{j}, x_{k+\Delta_{l}}^{i}\right)} \tag{3.9}
\end{equation*}
$$

The estimate of $\nabla_{\theta} \log \left(\gamma_{k+1, \theta}^{l}(1)\right)$ is:

$$
\begin{equation*}
\left.\nabla_{\theta} \widehat{\log \left(\gamma_{k+1, \theta}^{l}\right.}(1)\right):=\frac{\sum_{i=1}^{N} G_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{i}, u_{k, l}^{i}\right) F_{k, \theta}^{l, N}\left(u_{k, l}^{i}\right)}{\sum_{i=1}^{N} G_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{i}, u_{k, l}^{i}\right)} \tag{3.10}
\end{equation*}
$$

Set $k=k+1$ and return to the start of $2 .$.
3.3. Discussion of Algorithm 3.1. There are several remarks worth making, before proceeding. Firstly, the cost of this algorithm per unit time is $\mathcal{O}\left(N \Delta_{l}^{-1}+N^{2}\right)$. In detail, the cost of the particle filter is $\mathcal{O}\left(N \Delta_{l}^{-1}\right)$. The cost of calculating $F_{k, \theta}^{l, N}\left(u_{k, l}^{i}\right), 1 \leq i \leq N$, in (3.9), is $\mathcal{O}\left(N^{2}\right)$;
$\Delta_{l}$ is not involved here due to cancellations - see Remark 3.2. The cost of (3.10) is $\mathcal{O}(N)$, given the particle filter has already been executed. There are several implications of this remark. Based upon the results in [13] and Theorem 3.1, in a sequel work we prove, under appropriate assumptions, we will have the following MSE, for $(k, N) \in \mathbb{N}^{2}$ :

$$
\begin{equation*}
\left.\mathbb{E}_{\theta}\left[\| \nabla_{\theta} \widehat{\log \left(\gamma_{k, \theta}^{l}\right.}(1)\right)-\nabla_{\theta} \log \left(\gamma_{k, \theta}(1)\right) \|_{2}^{2}\right] \leq C\left(\frac{1}{N}+\Delta_{l}\right) \tag{3.11}
\end{equation*}
$$

where $C$ is a constant that does not depend on $N$ or $l$. To achieve an MSE of $\mathcal{O}\left(\epsilon^{2}\right)$ for some $\epsilon>0$ given, one sets $l$ so that $\Delta_{l}=\epsilon^{2}$ (i.e. $\left.l=\mathcal{O}(|\log (\epsilon)|)\right)$ and $N=\mathcal{O}\left(\epsilon^{-2}\right)$. The cost per unit time of doing this is then $\mathcal{O}\left(\epsilon^{-4}\right)$. We note that to choose $l$ as specified, one has to have access to an appropriately finely observed data set and this is assumed throughout. Typically, one could use a multilevel Monte Carlo method, as in [20], to reduce the cost to achieve an MSE of $\mathcal{O}\left(\epsilon^{2}\right)$. However, in this case as the $\mathcal{O}\left(N^{2}\right)$ cost dominates and does not depend on $l$, one can easily check that such a variance reduction method will not improve our particle method. To understand this, one can prove a bound on the MSE, for instance of the type conjectured later in this article (4.19), and then try to minimize the cost, by selecting the appropriate number of samples on each level to obtain a given MSE. This latter problem leads to a constrained minimization problem that can be solved using Lagrange multipliers (as in e.g. [9]), but one can show that this yields that the order of the cost to achieve an MSE of $\mathcal{O}\left(\epsilon^{2}\right)$ is still $\mathcal{O}\left(\epsilon^{-4}\right)$.

Secondly, it is important to note that the method in [23] can reduce the cost of online score estimation to $\mathcal{O}\left(N \Delta_{l}^{-1}\right)$ per unit time. However, in order to do so, one requires that $m_{\theta}^{l}\left(x, x^{\prime}\right)$ is uniformly lower-bounded in ( $x, x^{\prime}$ ), which does not typically occur for Euler-discretized diffusion densities. As a result, we only use the approach shown in Algorithm 3.1. Note that [16], in a different but related context, considers using unbiased and non-negative estimates of the transition density in the approach in [23], but such estimates are not always available.

Thirdly and rather importantly, there is a potential issue related to the construction of the algorithm. We have started with a continuous-time formula, discretized it and applied what are essentially discrete-time methods for smoothing of additive functionals. A serious caveat is that the algorithm is not well-defined as $l \rightarrow \infty$, which is what we mean by saying it has no (Wiener) path-space formulation. The source of the issue is related to using approximations of the transition density $m_{\theta}^{l}\left(\check{x}_{k}^{j}, x_{k+\Delta_{l}}^{i}\right)$ in (3.9), which can degenerate when $l$ is high. This will result in increasing Monte Carlo variance and computational cost and may mean that $C$ in (3.11) explodes exponentially in $l$. We refer the reader to [36, Figure 1.1] for a numerical example. This issue has manifested itself in MCMC schemes for inferring fully observed SDEs (see e.g. [24]), but in the context of particle smoothing and Algorithm 3.1 there are additional considerations. The resampling operation introduces discontinuities. Often such terminology refers to the lack of continuity of the transition density:

$$
\left.\theta \rightarrow \nabla_{\theta} \widehat{\log \left(\gamma_{k, \theta}^{l}\right.}(1)\right)
$$

but here we are interested in the behavior of $m_{\theta}^{l}\left(\check{x}_{k}^{j}, x_{k+\Delta_{l}}^{i}\right)$ when we combine points $\check{x}_{k}^{j}$ and $x_{k+\Delta_{l}}^{i}$ that are intrinsically not obtained in a continuous way as $\Delta_{l}$ diminishes. This issue has not received attention in earlier numerical studies, but remains a concern. As a result, we now consider defining an algorithm that is robust to the size of the time-discretization mesh and hence has a path-space formulation.

## 4. Path-Space Feynman-Kac Formulation.

4.1. Data Augmentation using Bridges. We begin this section with a review of the method in $[29,34]$. For simplicity we consider the case $t \in[0,1]$ and let $\mathbf{X}:=\left\{X_{t}\right\}_{t \in[0,1]}$, and $\mathbf{W}:=$ $\left\{W_{t}\right\}_{t \in[0,1]}$. Let also $p_{\theta}\left(x, t ; x^{\prime}, 1\right)$ denote the unknown transition density from time $t$ to 1 associated to (2.2) and let also $p_{\theta}\left(x, x^{\prime}\right):=p_{\theta}\left(x, 0 ; x^{\prime}, 1\right)$. Suppose one could sample from $p_{\theta}$ to obtain $\left(x, x^{\prime}\right) \in$ $\mathbb{R}^{2 d_{x}}$. Then one can interpolate these points by using a bridge process which has a drift given by $b_{\theta}(x)+a(x) \nabla_{x} \log p_{\theta}\left(x, t ; x^{\prime}, 1\right)$, as we will explain below. Let $\overline{\mathbb{P}}_{\theta, x, x^{\prime}}$ denote the law of the solution of the $\operatorname{SDE}(2.2)$, on $[0,1]$, started at $x_{*}=x$ and conditioned to hit $x^{\prime}$ at time 1.

As $p_{\theta}$ is intractable in general, we consider a user-specified auxiliary process $\left\{\tilde{X}_{t}\right\}_{t \in[0,1]}$ following:

$$
\begin{equation*}
d \tilde{X}_{t}=\tilde{b}_{\theta}\left(t, \tilde{X}_{t}\right) d t+\tilde{\sigma}\left(t, \tilde{X}_{t}\right) d W_{t}, \quad t \in[0,1], \quad \tilde{X}_{0}=x \tag{4.1}
\end{equation*}
$$

where for each parameter value $\theta \in \Theta, \tilde{b}_{\theta}:[0,1] \times \mathbb{R}^{d_{x}} \rightarrow \mathbb{R}^{d_{x}}$ and $\tilde{\sigma}: \mathbb{R}^{d_{x}} \rightarrow \mathbb{R}^{d_{x} \times d_{x}}$ is such that $\tilde{a}\left(1, x^{\prime}\right):=\tilde{\sigma}\left(1, x^{\prime}\right) \tilde{\sigma}\left(1, x^{\prime}\right)^{*} \equiv a\left(x^{\prime}\right)$. Most importantly, (4.1) is chosen so that its transition density is available. To avoid confusion - as the specification of process (4.1) can involve parameter $\theta$ and a given ending position $x^{\prime}$ - we note that the transition density of (4.1) from time $t$ to time 1 corresponds to a mapping $z \rightarrow \tilde{p}_{\theta, x^{\prime}}(x, t ; z, 1)$. We also use the notation $\tilde{p}_{\theta, x^{\prime}}(x, z):=\tilde{p}_{\theta, x^{\prime}}(x, 0 ; z, 1)$. One possible choice is to use an Ornstein-Uhlenbeck process (e.g. obtained using linearizations or variational inference with (2.2) [29, Section 1.3]); see also [29, Section 2.2] for technical conditions on $\tilde{b}_{\theta}, \tilde{a}, \tilde{p}_{\theta, x^{\prime}}$. The main purpose of $\left\{\tilde{X}_{t}\right\}_{t \in[0,1]}$ is to construct another process $\left\{X_{t}^{\circ}\right\}_{t \in[0,1]}$ conditioned to hit a given $x^{\prime}$ at $t=1$. The latter will form an importance proposal for $\left\{X_{t}\right\}_{t \in[0,1]}$. Let:

$$
\begin{equation*}
d X_{t}^{\circ}=b_{\theta}^{\circ}\left(t, X_{t}^{\circ} ; x^{\prime}\right) d t+\sigma\left(X_{t}^{\circ}\right) d W_{t}, \quad t \in[0,1], \quad X_{0}^{\circ}=x \tag{4.2}
\end{equation*}
$$

where:

$$
b_{\theta}^{\circ}\left(t, x ; x^{\prime}\right)=b_{\theta}(x)+a(x) \nabla_{x} \log \tilde{p}_{\theta, x^{\prime}}\left(x, t ; x^{\prime}, 1\right)
$$

and denote by $\mathbb{P}_{\theta, x, x^{\prime}}^{\circ}$ the probability law of the solution of (4.2). The SDE in (4.2) gives rise to a function:

$$
\begin{equation*}
\mathbf{W} \rightarrow C_{\theta}\left(x, \mathbf{W}, x^{\prime}\right) \tag{4.3}
\end{equation*}
$$

mapping the driving Wiener noise $\mathbf{W}$ to the solution of (4.2), so we have effectively reparameterized the problem from $\mathbf{X}$ to $\left(\mathbf{W}, x^{\prime}\right)$.

Now, following [29], the two measures $\overline{\mathbb{P}}_{\theta, x, x^{\prime}}$ and $\mathbb{P}_{\theta, x, x^{\prime}}^{\circ}$ are absolutely continuous w.r.t. each other, with Radon-Nikodym derivative:

$$
\begin{equation*}
\frac{d \overline{\mathbb{P}}_{\theta, x, x^{\prime}}}{d \mathbb{P}_{\theta, x, x^{\prime}}^{\circ}}(\mathbf{X})=\exp \left\{\int_{0}^{1} L_{\theta}\left(t, X_{t}\right) d t\right\} \times \frac{\tilde{p}_{\theta, x^{\prime}}\left(x, x^{\prime}\right)}{p_{\theta}\left(x, x^{\prime}\right)} \tag{4.4}
\end{equation*}
$$

where:

$$
\begin{aligned}
& L_{\theta}(t, x):=\left(b_{\theta}(x)-\tilde{b}_{\theta}(t, x)\right)^{*} \nabla_{x} \log \tilde{p}_{\theta, x^{\prime}}\left(x, t ; x^{\prime}, 1\right) \\
& -\frac{1}{2} \operatorname{Tr}\left\{[a(x)-\tilde{a}(t, x)]\left[-\nabla_{x}^{2} \log \tilde{p}_{\theta, x^{\prime}}\left(x, t ; x^{\prime}, 1\right)-\nabla_{x} \log \tilde{p}_{\theta, x^{\prime}}\left(x, t ; x^{\prime}, 1\right) \nabla_{x} \log \tilde{p}_{\theta, x^{\prime}}\left(x, t ; x^{\prime}, 1\right)^{*}\right]\right\}
\end{aligned}
$$

with $\operatorname{Tr}(\cdot)$ denoting the trace of a squared matrix. Note that, in the case when $\sigma=\sigma(x)$ is not a constant function, then, typically, $x^{\prime} \rightarrow \tilde{p}_{\theta, x^{\prime}}\left(x, x^{\prime}\right)$ will not integrate to 1 and will give rise to
a non-trivial distribution to sample from. As the complete algorithm will require being able to sample from the transition density, we rewrite:

$$
\begin{equation*}
\frac{d \overline{\mathbb{P}}_{\theta, x, x^{\prime}}}{\widetilde{\mathbb{P}}_{\theta, x, x^{\prime}}}(\mathbf{X})=\exp \left\{\int_{0}^{1} L_{\theta}\left(t, X_{t}\right) d t\right\} \times \frac{\tilde{p}_{\theta, x^{\prime}}\left(x, x^{\prime}\right)}{p_{\theta}\left(x, x^{\prime}\right) \hat{p}_{\theta}\left(x, x^{\prime}\right)} \times \hat{p}_{\theta}\left(x, x^{\prime}\right) \tag{4.5}
\end{equation*}
$$

where an arbitrary, tractable and easy to sample density $\hat{p}_{\theta}\left(x, x^{\prime}\right)$ is used to sample $x^{\prime}$.
4.2. Estimation of Score in Continuous-Time. We return to the expression of the score function in (2.4) and use the alternative change of measure described above. Consider the processes:

$$
\mathbf{X}_{k}:=\left\{X_{t}\right\}_{t \in[k, k+1]}, \quad \mathbf{Y}_{k}:=\left\{Y_{t}\right\}_{t \in[k, k+1]}, \quad k \in \mathbb{N}_{0}
$$

We introduce the following notation:

$$
\begin{aligned}
\Psi_{\theta}\left(\mathbf{X}_{k}\right)= & \int_{k}^{k+1} L_{\theta}\left(t, X_{t}\right) d t \\
J_{k, \theta}\left(\mathbf{X}_{k}, \mathbf{Y}_{k}\right)= & \int_{k}^{k+1} h_{\theta}\left(X_{t}\right)^{*} d Y_{t}-\frac{1}{2} \int_{k}^{k+1} h_{\theta}\left(X_{t}\right)^{*} h_{\theta}\left(X_{t}\right) d t \\
\Lambda_{k, \theta}\left(\mathbf{X}_{k}, \mathbf{Y}_{k}\right)= & \int_{k}^{k+1}\left(\nabla_{\theta} b_{\theta}\left(X_{t}\right)\right)^{*} a\left(X_{t}\right)^{-1}\left(d X_{t}-b_{\theta}\left(X_{t}\right) d t\right) \\
& \quad+\int_{k}^{k+1}\left(\nabla_{\theta} h_{\theta}\left(X_{t}\right)\right)^{*} d Y_{t}-\int_{k}^{k+1}\left(\nabla_{\theta} h_{\theta}\left(X_{t}\right)\right)^{*} h_{\theta}\left(X_{t}\right) d t \\
\Phi_{k, \theta}\left(\mathbf{X}_{k}, \mathbf{Y}_{k}\right)= & J_{k, \theta}\left(\mathbf{X}_{k}, \mathbf{Y}_{k}\right)+\Psi_{\theta}\left(\mathbf{X}_{k}\right)+\log \frac{\tilde{p}_{\theta, x_{k+1}}\left(x_{k}, x_{k+1}\right)}{\hat{p}_{\theta}\left(x_{k}, x_{k+1}\right)} .
\end{aligned}
$$

Note all the integrands can be computed point wise.
Remark 4.1. The Wiener process in (4.3) is defined on the time interval $[0,1]$, thus so is the transform $C_{\theta}\left(x, \mathbf{W}, x^{\prime}\right)$. In the derivations below, one needs to calculate $J_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right), \mathbf{Y}_{k}\right)$ and $\Lambda_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right), \mathbf{Y}_{k}\right)$, for $\mathbf{W}_{k}$ 's that correspond to samples from the Wiener measure on $[0,1]$. With some abuse of notation, it is to be understood that the path $C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right)$ is 'shifted' from $[0,1]$ to $[k, k+1]$, so all quantities below agree with the notation introduced above. Also, the calculation of $\Psi_{\theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right)\right)$ will be required, but this should create no confusion.
Under these definitions, for any $T \in \mathbb{N}$ the score function in (2.4) can be rewritten as:

$$
\nabla_{\theta} \log \left(\gamma_{T, \theta}(1)\right) \equiv \frac{\overline{\mathbb{E}}_{\theta}\left[\left(\sum_{k=0}^{T-1} \Lambda_{k, \theta}\left(\mathbf{X}_{k}, \mathbf{Y}_{k}\right)\right) \exp \left\{\sum_{k=0}^{T-1} J_{k, \theta}\left(\mathbf{X}_{k}, \mathbf{Y}_{k}\right)\right\} \mid \mathcal{Y}_{T}\right]}{\overline{\mathbb{E}}_{\theta}\left[\exp \left\{\sum_{k=0}^{T-1} J_{k, \theta}\left(\mathbf{X}_{k}, \mathbf{Y}_{k}\right)\right\} \mid \mathcal{Y}_{T}\right]}
$$

Making use of the transform (4.3) and the density expression in (4.4), we can equivalently write:

$$
\begin{equation*}
\nabla_{\theta} \log \left(\gamma_{T, \theta}(1)\right)= \tag{4.6}
\end{equation*}
$$

$$
=\frac{\widetilde{\mathbb{E}}_{\theta}\left[\left(\sum_{k=0}^{T-1} \Lambda_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right), \mathbf{Y}_{k}\right)\right) \exp \left\{\sum_{k=0}^{T-1} \Phi_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right), \mathbf{Y}_{k}\right)\right\} \mid \mathcal{Y}_{T}\right]}{\widetilde{\mathbb{E}}_{\theta}\left[\exp \left\{\sum_{k=0}^{T-1} \Phi_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right), \mathbf{Y}_{k}\right)\right\} \mid \mathcal{Y}_{T}\right]}
$$

where, the expectation $\widetilde{\mathbb{E}}_{\theta}\left[\cdot \mid \mathcal{Y}_{T}\right]$ is considered under the probability measure:

$$
\begin{equation*}
\widetilde{\mathbb{P}}_{\theta}\left(d\left(\mathbf{W}_{0}, x_{1}, \ldots, \mathbf{W}_{T-1}, x_{T}\right)\right):=\bigotimes_{k=0}^{T-1}\left[\mathbb{W}\left(d \mathbf{W}_{k}\right) \otimes \hat{p}_{\theta}\left(x_{k}, x_{k+1}\right) d x_{k+1}\right] \tag{4.7}
\end{equation*}
$$

independently of $\mathcal{Y}_{T} ;$ here, $\mathbb{W}$ is the standard Wiener measure on $[0,1]$ and $x_{0}=x_{*}$.
REmARK 4.2. The approach that has been adopted here can also be used if $\sigma$ depends upon $\theta$. Assuming the formula is well-defined, one would have a score function with an expression of the type:

$$
\frac{\widetilde{\mathbb{E}}_{\theta}\left[\left(\sum_{k=0}^{T-1} \Xi_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right)\right) \exp \left\{\sum_{k=0}^{T-1} \Phi_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right), \mathbf{Y}_{k}\right)\right\} \mid \mathcal{Y}_{T}\right]\right.}{\widetilde{\mathbb{E}}_{\theta}\left[\exp \left\{\sum_{k=0}^{T-1} \Phi_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right), \mathbf{Y}_{k}\right)\right\} \mid \mathcal{Y}_{T}\right]}
$$

where:

$$
\Xi_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right)=\left\{\nabla_{\theta} \Phi_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right), \mathbf{Y}_{k}\right)\right\}+\nabla_{\theta} \log \left\{\hat{p}_{\theta}\left(x_{k}, x_{k+1}\right)\right\}\right.
$$

and with an appropriate modification of the approach to allow $\sigma$ to depend on $\theta$. To keep consistency with the ideas in Section 3 we do not consider the formula from herein, but remark that extension of the forthcoming methodology to this case is straightforward.

REMARK 4.3. Recent advances in [6] have extended the construct of the auxiliary bridge process - developed via (4.1), (4.2) herein - to the setting of hypoelliptic SDEs. Though we do not pursue this direction here for the purpose of easing the exposition, we remark that, given these new developments, one can now, in principle, obtain score function estimates - thus, also carry out parameter inference - in the hypoelliptic regime along the same steps we follow in the current work.

The expression in (4.6), together with the (trivially) Markovian dynamics for the process in (4.7) allow one to construct a backward Feynman-Kac type formula as in (3.6). To better connect the approach here and that in Section 3, define $u_{k}=\left(\mathbf{W}_{k}, x_{k+1}\right)$ for $k \in \mathbb{N}_{0}, u_{-1}=x_{*}$ and:

$$
\begin{aligned}
\Lambda_{k, \theta}^{C}\left(u_{k-1}, u_{k}\right) & =\Lambda_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right), \mathbf{Y}_{k}\right) \\
\Phi_{k, \theta}^{C}\left(u_{k-1}, u_{k}\right) & =\Phi_{k, \theta}\left(C_{\theta}\left(x_{k}, \mathbf{W}_{k}, x_{k+1}\right), \mathbf{Y}_{k}\right) \\
F_{T, \theta}\left(u_{0}, \ldots, u_{T-1}\right) & =\sum_{k=0}^{T-1} \Lambda_{k, \theta}^{C}\left(u_{k-1}, u_{k}\right)
\end{aligned}
$$

Superscript $C$ is motivated by the well-posedness of the formula in continuous-time path-space. Set:

$$
\pi_{k, \theta}\left(d\left(u_{0}, \ldots, u_{k}\right)\right):=\frac{\left(\prod_{p=0}^{k} \exp \left\{\Phi_{p, \theta}^{C}\left(u_{p-1}, u_{p}\right)\right\}\right) \widetilde{\mathbb{P}}_{\theta}\left(d\left(u_{0}, \ldots, u_{k}\right)\right)}{\int_{E^{k+1}}\left(\prod_{p=0}^{k} \exp \left\{\Phi_{p, \theta}^{C}\left(u_{p-1}, u_{p}\right)\right\}\right) \widetilde{\mathbb{P}}_{\theta}\left(d\left(u_{0}, \ldots, u_{k}\right)\right)}
$$

with $E=C\left([0,1], \mathbb{R}^{d_{x}}\right)$. Then, we have the representation:

$$
\begin{equation*}
\nabla_{\theta} \log \left(\gamma_{T, \theta}(1)\right)=\int_{E^{T}} F_{T, \theta}\left(u_{0}, \ldots, u_{T-1}\right) \mathbb{Q}_{T-1, \theta}\left(d\left(u_{0}, \ldots, u_{T-1}\right)\right) \tag{4.8}
\end{equation*}
$$

where:

$$
\mathbb{Q}_{T-1, \theta}\left(d\left(u_{0}, \ldots, u_{T-1}\right)\right)=\pi_{T-1, \theta}\left(d u_{T-1}\right) \prod_{k=1}^{T-1} B_{k, \theta, \pi_{k-1, \theta}}\left(u_{k}, d u_{k-1}\right),
$$

and:

$$
B_{k, \theta, \pi_{k-1, \theta}}\left(u_{k}, d u_{k-1}\right):=\frac{\pi_{k-1, \theta}\left(d u_{k-1}\right) \exp \left\{\Phi_{k, \theta}^{C}\left(u_{k-1}, u_{k}\right)\right\} \hat{p}_{\theta}\left(x_{k}, x_{k+1}\right)}{\pi_{k-1, \theta}\left(\exp \left\{\Phi_{k, \theta}^{C}\left(\cdot, u_{k}\right)\right\} \hat{p}_{\theta}\left(\cdot, x_{k+1}\right)\right)}
$$

We remark that formula (4.8) is a type of backward Feynman-Kac formula in continuous-time, which to our knowledge is new.

As in the case of Algorithm 3.1, an effective Monte Carlo approximation of such a smoothing expectation (4.8) is given in Algorithm 4.1. The estimates of the score function are given in equations (4.9)-(4.10) in Algorithm 4.1.

## Algorithm 4.1 Online Score Function Estimation on Path-Space <br> 1. For $i \in\{1, \ldots, N\}$, sample $u_{0}^{i}$ i.i.d. from $\mathbb{W}(\cdot) \otimes \hat{p}_{\theta}\left(x_{*}, \cdot\right)$. The estimate of $\nabla_{\theta} \log \left(\gamma_{1, \theta}(1)\right)$ is: <br> $$
\begin{equation*} \left.\nabla_{\theta} \widehat{\log \left(\gamma_{1, \theta}\right.}(1)\right):=\frac{\sum_{i=1}^{N} \exp \left\{\Phi_{0, \theta}^{C}\left(x_{*}, u_{0}^{i}\right)\right\} \Lambda_{0, \theta}^{C}\left(x_{*}, u_{0}^{i}\right)}{\sum_{i=1}^{N} \exp \left\{\Phi_{0, \theta}^{C}\left(x_{*}, u_{0}^{i}\right)\right\}} . \tag{4.9} \end{equation*}
$$

Set $k=1$ and for $i \in\{1, \ldots, N\}, \check{u}_{-1}^{i}=x_{*}$.
2. For $i \in\{1, \ldots, N\}$ sample $\breve{u}_{k-1}^{i}$ from:

$$
\sum_{i=1}^{N} \frac{\exp \left\{\Phi_{k-1, \theta}^{C}\left(\breve{u}_{k-2}^{i}, u_{k-1}^{i}\right)\right\}}{\sum_{j=1}^{N} \exp \left\{\Phi_{k-1, \theta}^{C}\left(\check{u}_{k-2}^{j}, u_{k-1}^{j}\right)\right\}} \delta_{\left\{u_{k-1}^{i}\right\}}(\cdot) .
$$

If $k=1$, for $i \in\{1, \ldots, N\}$, set $\widetilde{F}_{k-1, \theta}^{N}\left(\breve{u}_{0}^{i}\right)=\Lambda_{0, \theta}^{C}\left(x_{*}, \breve{u}_{0}^{i}\right)$.
3. For $i \in\{1, \ldots, N\}$, sample $u_{k}^{i}$ from $\mathbb{W}(\cdot) \otimes \hat{p}_{\theta}^{l}\left(\breve{x}_{k}^{i}, \cdot\right)$. For $i \in\{1, \ldots, N\}$, compute:

$$
\widetilde{F}_{k, \theta}^{N}\left(u_{k}^{i}\right)=\frac{\sum_{j=1}^{N} \exp \left\{\Phi_{k, \theta}^{C}\left(\check{u}_{k-1}^{j}, u_{k}^{i}\right)\right\} \hat{p}_{\theta}\left(\check{x}_{k}^{j}, x_{k+1}^{i}\right)\left\{\widetilde{F}_{k-1, \theta}^{N}\left(\check{u}_{k-1}^{j}\right)+\Lambda_{k, \theta}^{C}\left(\check{u}_{k-1}^{j}, u_{k}^{i}\right)\right\}}{\sum_{j=1}^{N} \exp \left\{\Phi_{k, \theta}^{C}\left(\breve{u}_{k-1}^{j}, u_{k}^{i}\right)\right\} \hat{p}_{\theta}\left(\check{x}_{k}^{j}, x_{k+1}^{i}\right)} .
$$

The estimate of $\nabla_{\theta} \log \left(\gamma_{k+1, \theta}(1)\right)$ is:

$$
\begin{equation*}
\nabla_{\theta} \log \widehat{\left(\gamma_{k+1, \theta}(1)\right)}:=\frac{\sum_{i=1}^{N} \exp \left\{\Phi_{k, \theta}^{C}\left(\breve{u}_{k-1}^{i}, u_{k}^{i}\right)\right\} \widetilde{F}_{k, \theta}^{N}\left(u_{k}^{i}\right)}{\sum_{i=1}^{N} \exp \left\{\Phi_{k, \theta}^{C}\left(\breve{u}_{k-1}^{i}, u_{k}^{i}\right)\right\}} . \tag{4.10}
\end{equation*}
$$

Set $k=k+1$ and return to the start of 2 .
4.3. Time-Discretization. Whilst conceptually important, path-space valued Algorithm 4.1 can seldom be implemented directly in practice; unbiased methods e.g. [4] may be possible, but would be cumbersome. We develop a time-discretization procedure, in a similar manner to that considered in Section 3.2.

We will discretize on the uniform grid with increment $\Delta_{l}=2^{-l}$. Let $k \in \mathbb{N}_{0}$ and define:

$$
u_{k, l}=\left(z_{k+\Delta_{l}}, z_{k+2 \Delta_{l}}, \ldots, z_{k+1-\Delta_{l}}, x_{k+1}\right) \in\left(\mathbb{R}^{d_{x}}\right)^{\Delta_{l}^{-1}}=E_{l},
$$

where $z_{k+\Delta_{l}}, z_{k+2 \Delta_{l}}, \ldots, z_{k+1-\Delta_{l}}$ will represent increments of Brownian motion and $u_{-1, l}=x_{*}$. Define the Markov kernel on $E_{l}$, for $k \in \mathbb{N}$ :

$$
\widetilde{M}_{\theta}^{l}\left(u_{k-1, l}, d u_{k, l}\right)=\left(\prod_{s=1}^{\Delta_{l}^{-1}-1} \phi_{l}\left(z_{k+s \Delta_{l}}\right) d z_{k+s \Delta_{l}}\right) \hat{p}_{\theta}\left(x_{k}, x_{k+1}\right) d x_{k+1},
$$

where $\phi_{l}\left(z_{k+s \Delta_{l}}\right)$ is the density associated to the $\mathcal{N}_{d_{x}}\left(0, \Delta_{l} I_{d_{x}}\right)$ distribution. We denote the density of $\widetilde{M}_{\theta}^{l}$ as $\widetilde{Q}_{\theta}^{l}$. Set:

$$
\tilde{\eta}_{0, \theta}^{l}\left(d u_{0, l}\right)=\left(\prod_{s=1}^{\Delta_{l}^{-1}-1} \phi_{l}\left(z_{s \Delta_{l}}\right) d z_{s \Delta_{l}}\right) \hat{p}_{\theta}\left(x_{*}, x_{1}\right) d x_{1} .
$$

Now set, for $(k, s) \in\{0,1, \ldots, T-1\} \times\left\{0,1, \ldots, \Delta_{l}^{-1}-2\right\}$ :

$$
\begin{equation*}
X_{(s+1) \Delta_{l}+k}=X_{s \Delta_{l}+k}+b_{\theta}^{\circ}\left(s \Delta_{l}, X_{s \Delta_{l}+k} ; x_{k+1}\right) \Delta_{l}+\sigma\left(X_{s \Delta_{l}+k}\right) Z_{(s+1) \Delta_{l}+k} . \tag{4.11}
\end{equation*}
$$

Define for $k \in \mathbb{N}_{0}$ :

$$
\begin{aligned}
\Psi_{\theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) & =\sum_{p=0}^{\Delta_{1}^{-1}-1} L_{\theta}\left(p \Delta_{l}, x_{k+p \Delta_{l}}\right) \Delta_{l} ; \\
J_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) & =\sum_{p=0}^{\Delta_{l}^{-1}-1} h_{\theta}\left(x_{k+p \Delta_{l}}\right)^{*}\left(Y_{k+(s+1) \Delta_{l}}-Y_{k+s \Delta_{l}}\right)-\frac{1}{2} \sum_{p=0}^{\Delta_{l}^{-1}-1} h_{\theta}\left(x_{k+p \Delta_{l}}\right)^{*} h_{\theta}\left(x_{k+p \Delta_{l}}\right) \Delta_{l} ; \\
\Phi_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) & =J_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right)+\Psi_{\theta}^{l}\left(u_{k-1, l}, u_{k, l}\right)+\log \frac{\tilde{p}_{\theta, x_{k+1}}\left(x_{k}, x_{k+1}\right)}{\hat{p}_{\theta}\left(x_{k}, x_{k+1}\right)} ; \\
\widetilde{G}_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) & =\exp \left\{\Phi_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right)\right\} .
\end{aligned}
$$

Now, for $k \in \mathbb{N}_{0}$ :

$$
\begin{aligned}
& \widetilde{\Lambda}_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right)=\sum_{p=0}^{\Delta_{l}^{-1}-1}\left(\nabla_{\theta} b_{\theta}\left(x_{k+p \Delta_{l}}\right)\right)^{*} a\left(x_{k+p \Delta_{l}}\right)^{-1}\left(x_{k+(p+1) \Delta_{l}}-x_{k+p \Delta_{l}}-b_{\theta}\left(x_{k+p \Delta_{l}}\right) \Delta_{l}\right) \\
& \quad+\sum_{p=0}^{\Delta_{l}^{-1}-1}\left(\nabla_{\theta} h_{\theta}\left(x_{k+p \Delta_{l}}\right)\right)^{*}\left(Y_{k+(p+1) \Delta_{l}}-Y_{k+p \Delta_{l}}\right)-\sum_{p=0}^{\Delta_{l}^{-1}-1}\left(\nabla_{\theta} h_{\theta}\left(x_{k+p \Delta_{l}}\right)\right)^{*} h_{\theta}\left(x_{k+p \Delta_{l}}\right) \Delta_{l} .
\end{aligned}
$$

Set:

$$
\widetilde{F}_{T, \theta}^{l}\left(u_{0, l}, \ldots, u_{T-1, l}\right)=\sum_{k=0}^{T-1} \widetilde{\Lambda}_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) .
$$

Writing expectations w.r.t. $\widetilde{\eta}_{\theta}^{l}\left(d u_{0, l}\right) \prod_{k=1}^{T-1} \widetilde{M}_{\theta}^{l}\left(u_{k-1, l}, d u_{k, l}\right)$ as $\widetilde{\mathbb{E}}_{\theta}^{l}\left[\cdot \mid \mathcal{Y}_{T}\right]$, our discretized approximation of $\nabla_{\theta} \log \left(\gamma_{T, \theta}(1)\right)$ is:

$$
\nabla_{\theta} \log \left(\widetilde{\gamma}_{T, \theta}^{l}(1)\right):=\frac{\widetilde{\mathbb{E}}_{\theta}^{l}\left[\widetilde{F}_{T, \theta}^{l}\left(U_{0, l}, \ldots, U_{T-1, l}\right) \prod_{k=0}^{T-1} \widetilde{G}_{k, \theta}^{l}\left(U_{k-1, l}, U_{k, l}\right) \mid \mathcal{Y}_{T}\right]}{\widetilde{\mathbb{E}}_{\theta}^{l}\left[\prod_{k=0}^{T-1} \widetilde{G}_{k, \theta}^{l}\left(U_{k-1, l}, U_{k, l}\right) \mid \mathcal{Y}_{T}\right]}
$$

We note that, whilst terms $\nabla_{\theta} \log \left(\widetilde{\gamma}_{T, \theta}^{l}(1)\right), \nabla_{\theta} \log \left(\gamma_{T, \theta}^{l}(1)\right)$ should both converge to $\nabla_{\theta} \log \left(\gamma_{T, \theta}(1)\right)$, as $l \rightarrow \infty$, they will in general be different for any fixed $l$.

One can also easily develop a discretized time reversal formula such as (3.6) which will converge precisely to (4.8) as $l \rightarrow \infty$. Define, for $k \in \mathbb{N}_{0}$ :

$$
\widetilde{\pi}_{k, \theta}^{l}\left(d\left(u_{0, l}, \ldots, u_{k, l}\right)\right):=\frac{\left(\prod_{p=0}^{k} \widetilde{G}_{p, \theta}^{l}\left(u_{p-1, l}, u_{p, l}\right)\right) \widetilde{\eta}_{0, \theta}^{l}\left(d u_{0, l}\right) \prod_{p=1}^{k} \widetilde{M}_{\theta}^{l}\left(u_{p-1, l}, d u_{p, l}\right)}{\int_{E_{l}^{k+1}}\left(\prod_{p=0}^{k} \widetilde{G}_{p, \theta}^{l}\left(u_{p-1, l}, u_{p, l}\right)\right) \widetilde{\eta}_{0, \theta}^{l}\left(d u_{0, l}\right) \prod_{p=1}^{k} \widetilde{M}_{\theta}^{l}\left(u_{p-1, l}, d u_{p, l}\right)}
$$

Then, we have that:

$$
\begin{equation*}
\nabla_{\theta} \log \left(\widetilde{\gamma}_{T, \theta}^{l}(1)\right)=\int_{E_{l}^{T}} \widetilde{F}_{T, \theta}^{l}\left(u_{0, l}, \ldots, u_{T-1, l}\right) \widetilde{\mathbb{Q}}_{T-1, \theta}^{l}\left(d\left(u_{0, l}, \ldots, u_{T-1, l}\right)\right) \tag{4.12}
\end{equation*}
$$

where we have defined:

$$
\widetilde{\mathbb{Q}}_{T-1, \theta}^{l}\left(d\left(u_{0, l}, \ldots, u_{T-1, l}\right)\right):=\widetilde{\pi}_{T-1, \theta}^{l}\left(d u_{T-1, l}\right) \prod_{k=1}^{T-1} \widetilde{B}_{k, \theta, \pi_{k-1, \theta}^{l}}^{l}\left(u_{k, l}, d u_{k-1, l}\right)
$$

and:

$$
\widetilde{B}_{k, \theta, \widetilde{\pi}_{k-1, \theta}^{l}}^{l}\left(u_{k, l}, d u_{k-1, l}\right):=\frac{\widetilde{\pi}_{k-1, \theta}^{l}\left(d u_{k-1, l}\right) \widetilde{G}_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) \widetilde{Q}_{\theta}^{l}\left(u_{k-1, l}, u_{k, l}\right)}{\widetilde{\pi}_{k-1, \theta}^{l}\left(\widetilde{G}_{k, \theta}^{l}\left(\cdot, u_{k, l}\right) \widetilde{Q}_{\theta}^{l}\left(\cdot, u_{k, l}\right)\right)}
$$

We remark that, due to the structure of the model:

$$
\widetilde{B}_{k, \theta, \widetilde{\pi}_{k-1, \theta}^{l}}^{l}\left(u_{k, l}, d u_{k-1, l}\right)=\frac{\widetilde{\pi}_{k-1, \theta}^{l}\left(d u_{k-1, l}\right) \widetilde{G}_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) \hat{p}_{\theta}\left(x_{k}, x_{k+1}\right)}{\int_{E_{l}} \widetilde{\pi}_{k-1, \theta}^{l}\left(d u_{k-1, l}\right) \widetilde{G}_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right) \hat{p}_{\theta}\left(x_{k}, x_{k+1}\right)}
$$

REMARK 4.4. It is important to note that - in contrast to Remark 3.2-there is no cancellation of terms of $\widetilde{G}_{k, \theta}^{l}\left(u_{k-1, l}, u_{k, l}\right)$ in the numerator and denominator of this backward kernel. This is precisely due to recursion (4.11) which leads to a path-dependence of the future coordinates of the discretized bridge on the terminal position $x_{k+1}$.
4.4. Particle Approximation. Our online particle approximation of the gradient of the loglikelihood in (4.12), for a given $l \in \mathbb{N}_{0}$ is presented in Algorithm 4.2. Our estimates are given in (4.14) and (4.16) in Algorithm 4.2.

Algorithm 4.2 is simply the time-discretization of the procedure presented in Algorithm 4.1. A number of remarks are again of interest. Firstly, the cost of the algorithm per unit time is now $\mathcal{O}\left(N^{2} \Delta_{l}^{-1}\right)$. The increase in computational cost over Algorithm 3.1 is the fact that when computing $\widetilde{\Lambda}_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{j}, u_{k, l}^{i}\right)$ in (4.15), one must solve the recursion (4.11) for each $(i, j) \in\{1, \ldots, N\}^{2}$, which
has a cost $\mathcal{O}\left(\Delta_{l}^{-1}\right)$ and it is this cost that dominates. Secondly, following the discussion in Section 3.3 , we have proved in a companion work that, under appropriate assumptions, the MSE for $(k, N) \in$ $\mathbb{N}^{2}$ :

$$
\begin{equation*}
\left.\mathbb{E}_{\theta}\left[\| \nabla_{\theta} \widehat{\log \left(\widetilde{\gamma}_{k, \theta}^{l}\right.}(1)\right)-\nabla_{\theta} \log \left(\gamma_{k, \theta}(1)\right) \|_{2}^{2}\right] \leq C\left(\frac{1}{N}+\Delta_{l}\right) \tag{4.13}
\end{equation*}
$$

for constant $C$ that does not depend on $N, l$. To achieve an MSE of $\mathcal{O}\left(\epsilon^{2}\right)$ for some $\epsilon>0$ given, one sets $l=\mathcal{O}(|\log (\epsilon)|)$ and $N=\mathcal{O}\left(\epsilon^{-2}\right)$. The cost per unit time of doing this, is then $\mathcal{O}\left(\epsilon^{-6}\right)$. This is significantly worse than the approach in Algorithm 3.1, but we remark that when discussing the cost of Algorithm 3.1, in the bound (3.11), we have assumed that the constant $C$ does not depend upon $l$. However, in a sequel work we will show that under assumptions that this afore-mentioned $C$ explodes exponentially in $l$. Conversely, $C$ in (4.13) can be proved to be independent of $l$, precisely due to the path-space development we have adopted. We remark, however, that one can use an MLMC method to reduce this cost of $\mathcal{O}\left(\epsilon^{-6}\right)$ per unit time of Algorithm 4.2 and this algorithm is presented in the next section.

$$
\begin{align*}
& \text { Algorithm 4.2 Modified Online Score Function Estimation for a given } l \in \mathbb{N}_{0} \text {. } \\
& \hline \text { 1. For } i \in\{1, \ldots, N\} \text {, sample } u_{0, l}^{i} \text { i.i.d. from } \widetilde{\eta}_{0, \theta}^{l}(\cdot) \text {. The estimate of } \nabla_{\theta} \log \left(\widetilde{\gamma}_{1, \theta}^{l}(1)\right) \text { is: } \\
& \left.\qquad(4.14) \quad \widehat{\nabla_{\theta} \log \left(\widetilde{\gamma}_{1, \theta}^{l}\right.}(1)\right):=\frac{\sum_{i=1}^{N} \widetilde{G}_{0, \theta}^{l}\left(x_{*}, u_{0, l}^{i}\right) \widetilde{\Lambda}_{0, \theta}^{l}\left(x_{*}, u_{0, l}^{i}\right)}{\sum_{i=1}^{N} \widetilde{G}_{0, \theta}^{l}\left(x_{*}, u_{0, l}^{i}\right)} \tag{4.14}
\end{align*}
$$

Set $k=1$ and for $i \in\{1, \ldots, N\}, \check{u}_{-1, l}^{i}=x_{*}$.
2. For $i \in\{1, \ldots, N\}$ sample $\breve{u}_{k-1, l}^{i}$ from:

$$
\sum_{i=1}^{N} \frac{\widetilde{G}_{k-1, \theta}^{l}\left(\check{u}_{k-2, l}^{i}, u_{k-1, l}^{i}\right)}{\sum_{j=1}^{N} \widetilde{G}_{k-1, \theta}^{l}\left(\breve{u}_{k-2, l}^{j}, u_{k-1, l}^{j}\right)} \delta_{\left\{u_{k-1, l}^{i}\right\}}(\cdot)
$$

If $k=1$, for $i \in\{1, \ldots, N\}$, set $\widetilde{F}_{k-1, \theta}^{l, N}\left(\breve{u}_{0, l}^{i}\right)=\widetilde{\Lambda}_{0, \theta}^{l}\left(x_{*}, \check{u}_{0, l}^{i}\right)$.
3. For $i \in\{1, \ldots, N\}$, sample $u_{k, l}^{i}$ from $\widetilde{M}_{\theta}^{l}\left(\breve{u}_{k-1, l}^{i}, \cdot\right)$. For $i \in\{1, \ldots, N\}$, compute:

$$
\begin{equation*}
\widetilde{F}_{k, \theta}^{l, N}\left(u_{k, l}^{i}\right)=\frac{\sum_{j=1}^{N} \widetilde{G}_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{j}, u_{k, l}^{i}\right) \hat{p}_{\theta}\left(\check{x}_{k}^{j}, x_{k+1}^{i}\right)\left\{\widetilde{F}_{k-1, \theta}^{l, N}\left(\check{u}_{k-1, l}^{j}\right)+\widetilde{\Lambda}_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{j}, u_{k, l}^{i}\right)\right\}}{\sum_{j=1}^{N} \widetilde{G}_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{j}, u_{k, l}^{i}\right) \hat{p}_{\theta}\left(\check{x}_{k}^{j}, x_{k+1}^{i}\right)} \tag{4.15}
\end{equation*}
$$

The estimate of $\nabla_{\theta} \log \left(\widetilde{\gamma}_{k+1, \theta}^{l}(1)\right)$ is:

$$
\begin{equation*}
\left.\nabla_{\theta} \log \widehat{\left(\widetilde{\gamma}_{k+1, \theta}^{l}\right.}(1)\right)=\frac{\sum_{i=1}^{N} \widetilde{G}_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{i}, u_{k, l}^{i}\right) \widetilde{F}_{k, \theta}^{l, N}\left(u_{k, l}^{i}\right)}{\sum_{i=1}^{N} \widetilde{G}_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{i}, u_{k, l}^{i}\right)} . \tag{4.16}
\end{equation*}
$$

Set $k=k+1$ and return to the start of $2 .$.
4.5. Multilevel Particle Filter. We now present a new multilevel particle filter along with online estimation of the score-function. We fix $l \in \mathbb{N}$ for now and for $(k, s) \in \mathbb{N}_{0} \times\{l, l-1\}$ define:

$$
u_{k, s}=\left(z_{k+\Delta_{l}, s}, z_{k+2 \Delta_{l}, s}, \ldots, z_{k+1-\Delta_{l}, s}, x_{k+1, s}\right) \in\left(\mathbb{R}^{d_{x}}\right)^{\Delta_{s}^{-1}}=E_{s}
$$

We give the approach in Algorithm 4.3. Before explaining how one can use Algorithm 4.3 to provide online estimates of the score function, several remarks are required to continue. The first is related to the couplings mentioned in Algorithm 4.3 point 2. and point 3. bullet 1. The coupling in point 2 , requires a way to resample the indices of the particles so that they have the correct marginals. This topic has been investigated considerably in the literature, see e.g. [19, 30], and techniques that have been adopted include sampling maximal coupling, e.g. [17], or using the $\mathbb{L}_{2}$-Wasserstein optimal coupling [2]; in general the latter is found to be better in terms of variance reduction, but can only be implemented when $d_{x}=1$. We rely upon the maximal coupling in this paper, which has a cost of $\mathcal{O}(N)$ per unit time. For point 3 . bullet 1 , one again has a considerable degree of flexibility. In this article we sample the maximal coupling which can be achieved at a cost which is at most $\mathcal{O}(N)$ cost per-unit time using the algorithm of [33]. The second main remark of interest is that the basic filter that is sampled in Algorithm 4.3 is an entirely new coupled particle filter for diffusions (i.e. different to [17, 20]). The utility of the approach relative to [20] is of great interest, in the context of filtering.

Set $\left(l_{*}, L\right) \in \mathbb{N}^{2}$ with $l_{*}<L$. The idea is to run Algorithm 4.3, independently, for $l \in\left\{l_{*}, \ldots, L\right\}$ each with $N_{l}$ particles and, independently, Algorithm 4.2 for $l=l_{*}-1$ with $N_{l_{*}-1}$ particles. We then consider the estimate, for $k \in \mathbb{N}$

$$
\left.\left.\left.\left.\nabla_{\theta} \widehat{\log \left(\widetilde{\gamma}_{k, \theta}^{L}\right.}(1)\right)_{M L}:=\sum_{l=l_{*}}^{L}\left\{\nabla_{\theta} \widehat{\log \left(\widetilde{\gamma}_{k, \theta}^{l}\right.}(1)\right)-\nabla_{\theta} \widehat{\log \left(\widetilde{\gamma}_{k, \theta}^{l-1}\right.}(1)\right)\right\}+\nabla_{\theta} \operatorname{log(\widetilde {\gamma }_{k,\theta }^{l*}-1}(1)\right)
$$

where the summands on the right hand side are defined in (4.17) and (4.18) and the last term on the right hand side is as either (4.14) or (4.16) (depending on $k$ ). Now, we show in an on-going companion work, under appropriate assumptions, one has the following result for $\left(k, N_{l_{*}-1: L}\right) \in$ $\mathbb{N}^{L-l_{*}+2}$ :

$$
\begin{equation*}
\left.\mathbb{E}_{\theta}\left[\| \nabla_{\theta} \widehat{\log \left(\widetilde{\gamma}_{k, \theta}^{L}\right.}(1)\right)_{M L}-\nabla_{\theta} \log \left(\gamma_{k, \theta}(1)\right) \|_{2}^{2}\right] \leq C\left(\sum_{l=l_{*}-1}^{L} \frac{\Delta_{l}^{\beta}}{N_{l}}+\Delta_{L}\right) \tag{4.19}
\end{equation*}
$$

for constant $C$ that does not depend on $N, l$; also, $\beta=1$ if $\sigma$ is a constant function, else $\beta=1 / 2$. Choose: i) $L$ so that $\Delta_{L}=\mathcal{O}\left(\epsilon^{2}\right)$, for $\epsilon>0$ given; ii) if $\beta=1, N_{l}=\mathcal{O}\left(\epsilon^{-2} \Delta_{l}^{1 / 2+\rho}\right)$ for some $0<\rho<1 / 2$. These selections yield an MSE of $\mathcal{O}\left(\epsilon^{2}\right)$ for a cost of $\mathcal{O}\left(\epsilon^{-4}\right)$. If $\beta=1 / 2$, one can set $N_{l}=\mathcal{O}\left(\epsilon^{-2} \Delta_{l}^{1 / 2+\rho} \Delta_{L}^{-\rho}\right)$ for some $\rho>0$. This will yield an MSE of $\mathcal{O}\left(\epsilon^{2}\right)$ for a cost of $\mathcal{O}\left(\epsilon^{-4(1+\rho)}\right)$. Such results are at least as good as the method in Section 3.3, assuming that latter approach does not collapse with $l$.

We remark that it is possible to produce an almost-surely unbiased estimator of the score function, when $\theta$ is the true parameter, using a combination of the multilevel method that has been developed here and the approach in [18]. This is left for future work.
5. Numerical Results. In this section, we consider four models to investigate the various properties of our algorithms. The score function is estimated using both Algorithms 3.1 and 4.2 for a fixed $\theta$. We will show, as expected, that they are equivalent for a large number of particles $N$ and

Algorithm 4.3 Coupled Online Score Function Estimation for a given $l \in \mathbb{N}$.

1. $\quad$ For $i \in\{1, \ldots, N\}$, sample $u_{0, l}^{i}$ i.i.d. from $\widetilde{\eta}_{0, \theta}^{l}(\cdot)$.

- For $(i, p) \in\{1, \ldots, N\} \times\left\{1, \ldots, \Delta_{l-1}^{-1}-1\right\}$, set $z_{p \Delta_{l-1}, l-1}^{i}=z_{p \Delta_{l-1}, l}^{i}+z_{p \Delta_{l-1}-\Delta_{l}, l}^{i}$ and $x_{1, l-1}^{i}=x_{1, l}^{i}$.
The estimate of $\nabla_{\theta} \log \left(\widetilde{\gamma}_{1, \theta}^{l}(1)\right)-\nabla_{\theta} \log \left(\widetilde{\gamma}_{1, \theta}^{l-1}(1)\right)$ is:

$$
\begin{gather*}
\left.\left.\nabla_{\theta} \widehat{\log \left(\widetilde{\gamma}_{1, \theta}^{l}\right.}(1)\right)-\nabla_{\theta} \widehat{\log \left(\widetilde{\gamma}_{1, \theta}^{l-1}\right.}(1)\right):= \\
\frac{\sum_{i=1}^{N} \widetilde{G}_{0, \theta}^{l}\left(x_{*}, u_{0, l}^{i}\right) \widetilde{\Lambda}_{0, \theta}^{l}\left(x_{*}, u_{0, l}^{i}\right)}{\sum_{i=1}^{N} \widetilde{G}_{0, \theta}^{l}\left(x_{*}, u_{0, l}^{i}\right)}-\frac{\sum_{i=1}^{N} \widetilde{G}_{0, \theta}^{l-1}\left(x_{*}, u_{0, l-1}^{i}\right) \widetilde{\Lambda}_{0, \theta}^{l-1}\left(x_{*}, u_{0, l-1}^{i}\right)}{\sum_{i=1}^{N} \widetilde{G}_{0, \theta}^{l-1}\left(x_{*}, u_{0, l-1}^{i}\right)} . \tag{4.17}
\end{gather*}
$$

Set $k=1$ and for $i \in\{1, \ldots, N\}, \check{u}_{-1, l}^{i}=\breve{u}_{-1, l-1}^{i}=x_{*}$.
2. For $i \in\{1, \ldots, N\}$, sample $\left(\alpha_{l}(i), \alpha_{l-1}(i)\right) \in\{1, \ldots, N\}^{2}$ from a coupling of:

$$
\sum_{i=1}^{N} \frac{\widetilde{G}_{k-1, \theta}^{l}\left(\check{u}_{k-2, l}^{\alpha_{l}(i)}, u_{k-1, l}^{\alpha_{l}(i)}\right)}{\sum_{j=1}^{N} \widetilde{G}_{k-1, \theta}^{l}\left(\check{u}_{k-2, l}^{j}, u_{k-1, l}^{j}\right)} \quad \text { and } \quad \sum_{i=1}^{N} \frac{\widetilde{G}_{k-1, \theta}^{l-1}\left(\check{u}_{k-2, l-1}^{\alpha_{l-1}(i)}, u_{k-1, l}^{\alpha_{l-1}(i)}\right)}{\sum_{j=1}^{N} \widetilde{G}_{k-1, \theta}^{l-1}\left(\check{u}_{k-2, l-1}^{j}, u_{k-1, l-1}^{j}\right)},
$$

and set $\left(\check{u}_{k-1, l}^{i}, \breve{u}_{k-1, l-1}^{i}\right)=\left(u_{k-1, l}^{\alpha_{l}(i)}, u_{k-1, l-1}^{\alpha_{l-1}(i)}\right)$. If $k=1$, for $i \in\{1, \ldots, N\}, s \in\{l, l-1\}$, set $\widetilde{F}_{k-1, \theta}^{s, N}\left(\check{u}_{0, s}^{i}\right)=\widetilde{\Lambda}_{0, \theta}^{s}\left(x_{*}, \check{u}_{0, s}^{i}\right)$.
3. - For $i \in\{1, \ldots, N\}$, sample $\left(x_{k+1, l}^{i}, x_{k+1, l-1}^{i}\right)$ from a coupling of $\hat{p}_{\theta}\left(\check{x}_{k, l}^{i}, \cdot\right)$ and $\hat{p}_{\theta}\left(\check{x}_{k, l-1}^{i}, \cdot\right)$.

- For $i \in\{1, \ldots, N\}$ sample $z_{k+\Delta_{l}, l}^{i}, \ldots, z_{k+1-\Delta_{l}, l}^{i}$ i.i.d. from $\prod_{s=1}^{\Delta_{l}^{-1}-1} \phi_{l}(\cdot)$.
- For $(i, p) \in\{1, \ldots, N\} \times\left\{1, \ldots, \Delta_{l-1}^{-1}-1\right\}$ set $z_{k+p \Delta_{l-1}, l-1}^{i}=z_{k+p \Delta_{l-1}, l}^{i}+z_{k+p \Delta_{l-1}-\Delta_{l}, l}^{i}$.
- For $(i, s) \in\{1, \ldots, N\} \times\{l . l-1\}$, compute:

$$
\widetilde{F}_{k, \theta}^{s, N}\left(u_{k, l}^{i}\right)=\frac{\sum_{j=1}^{N} \widetilde{G}_{k, \theta}^{s}\left(\check{u}_{k-1, s}^{j}, u_{k, s}^{i}\right) \hat{p}_{\theta}\left(\check{x}_{k, s}^{j}, x_{k+1, s}^{i}\right)\left\{\widetilde{F}_{k-1, \theta}^{s, N}\left(\check{u}_{k-1, s}^{j}\right)+\widetilde{\Lambda}_{k, \theta}^{s}\left(\check{u}_{k-1, s}^{j}, u_{k, s}^{i}\right)\right\}}{\sum_{j=1}^{N} \widetilde{G}_{k, \theta}^{s}\left(\check{u}_{k-1, s}^{j}, u_{k, s}^{i}\right) \hat{p}_{\theta}\left(\check{x}_{k, s}^{j}, x_{k+1, s}^{i}\right)} .
$$

The estimate of $\nabla_{\theta} \log \left(\widetilde{\gamma}_{k+1, \theta}^{l}(1)\right)-\nabla_{\theta} \log \left(\widetilde{\gamma}_{k+1, \theta}^{l-1}(1)\right)$ is:

$$
\begin{gather*}
\left.\left.\nabla_{\theta} \operatorname{log(\widetilde {\gamma }_{k+1,\theta }^{l}}(1)\right)-\nabla_{\theta} \log \widehat{\left(\widetilde{\gamma}_{k+1, \theta}^{l-1}\right.}(1)\right):= \\
\frac{\sum_{i=1}^{N} \widetilde{G}_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{i}, u_{k, l}^{i}\right) \widetilde{F}_{k, \theta}^{l, N}\left(u_{k, l}^{i}\right)}{\sum_{i=1}^{N} \widetilde{G}_{k, \theta}^{l}\left(\check{u}_{k-1, l}^{i}, u_{k, l}^{i}\right)}-\frac{\sum_{i=1}^{N} \widetilde{G}_{k, \theta}^{l-1}\left(\check{u}_{k-1, l-1}^{i}, u_{k, l-1}^{i}\right) \widetilde{F}_{k, \theta}^{l-1, N}\left(u_{k, l-1}^{i}\right)}{\sum_{i=1}^{N} \widetilde{G}_{k, \theta}^{l-1}\left(\check{u}_{k-1, l-1}^{i}, u_{k, l-1}^{i}\right)} . \tag{4.18}
\end{gather*}
$$

Set $k=k+1$ and return to the start of $2 .$.
a high level of descritization $l$. We then compare the cost of Algorithm 4.2 and its multilevel version Algorithm 4.3. As an application of our methods, we use Algorithms 3.1 and 4.3 for parameter estimation via stochastic gradient. The code is written in MATLAB and it can be downloaded from
https://github.com/ruzayqat/score_based_par_est.
We remark that we will not use Algorithm 4.2 for parameter estimation because it is 'slow' compared to the algorithms as illustrated in the previous section and in Figure 1. In Figure 1 we consider the Model 1, as described in the next section, with $T=40, l=10, \theta=(-0.4,-0.5), \kappa=2$ and $x_{*}=0.2$. Figure 1 provides a comparison between the cost of Algorithms 3.1 and 4.2 , which is the average machine time measured in seconds needed per each simulation, versus the number of particles $N$. As predicted by our theoretical conjectures, we see that the cost of Algorithm 3.1 is significantly lower than that of Algorithm 4.2.


Figure 1: Comparison between the cost of Algorithms 3.1 and 4.2 per each simulation versus the number of particles $N$. We run both algorithms on Model 1.
5.1. Models. In the following, parameters $(\kappa, \sigma)$ are fixed.

Model 1: Let $d_{x}=d_{y}=1, d_{\theta}=2$ and consider the following linear SDEs:

$$
\begin{aligned}
d X_{t} & =\theta_{1} X_{t} d t+\sigma d W_{t} ; \\
d Y_{t} & =\theta_{2}\left(\kappa-X_{t}\right) d t+d B_{t} .
\end{aligned}
$$

Model 2: Let $d_{x}=d_{y}=1, d_{\theta}=3$ and consider a nonlinear diffusion process along with a linear diffusion process of observations:

$$
\begin{aligned}
d X_{t} & =\left(\frac{\theta_{1}}{X_{t}}+\theta_{2} X_{t}\right) d t+\sigma d W_{t} \\
d Y_{t} & =\theta_{3}\left(\kappa-X_{t}\right) d t+d B_{t}
\end{aligned}
$$

Model 3: Let $d_{x}=d_{y}=1, d_{\theta}=3$ and consider a nonlinear signal along with a nonlinear diffusion process of observations. The first SDE is a Cox-Ingersoll-Ross process after an 1-1 transform. Thus:

$$
\begin{aligned}
d X_{t} & =\frac{1}{2}\left(\frac{\theta_{1} \theta_{2}-\sigma^{2}}{X_{t}}-\theta_{2} X_{t}\right) d t+\sigma d W_{t} \\
d Y_{t} & =\theta_{3}\left(\kappa-X_{t}^{2}\right) d t+d B_{t}
\end{aligned}
$$

This model has a solution if and only if $\theta_{1} \theta_{2}>2 \sigma^{2}$.
Model 4: Let $d_{x}=d_{y}=1, d_{\theta}=3$ and consider a type of Black-Scholes model with a stochastic volatility:

$$
\begin{aligned}
d X_{t} & =\theta_{1}\left(\theta_{2}-X_{t}\right) d t+\sigma\left(X_{t}\right) d W_{t} \\
d Y_{t} & =\left(\theta_{3}-\frac{1}{2} X_{t}^{2}\right) d t+d B_{t}
\end{aligned}
$$

where $\sigma\left(X_{t}\right)=\beta / \sqrt{X_{t}^{2}+1}$ and $\beta$ is fixed. In the hidden process, $\theta_{1}$ and $\theta_{2}$ are the speed and level of mean reversion and $\theta_{3}$ is a mean type level for the observation process, We will apply our methodology (see Figure 8 later on) on the log mid-price of Tesla Inc. stock in 2018. The dataset shown in Figure 2 represents the log mid-price at every second during a trading day for a total of 250 trading days.


Figure 2: The $\log$ of mid-price for every second of Tesla stock in 2018. The dataset contains $5.85 \times 10^{6}$ data points.
5.2. Simulation Results. In all our results data are generated from the model under the finest discretization considered except in model 4, where we use a real data. In Algorithm 4.2, we consider the auxiliary linear process following:

$$
d \tilde{X}_{t}=\sigma d W_{t}
$$

in models 1-3 and in model 4 it follows:

$$
d \tilde{X}_{t}=\sigma\left(x^{\prime}\right) d W_{t}
$$

In models 1-3, $\tilde{p}_{\theta, x^{\prime}}\left(x, t ; x^{\prime}, 1\right)=\mathcal{N}\left(x^{\prime} ; x,(1-t) \sigma^{2}\right)$, hence $\tilde{p}_{\theta, x^{\prime}}\left(x, x^{\prime}\right)=\mathcal{N}\left(x^{\prime} ; x, \sigma^{2}\right)$, which is easy to sample $x^{\prime}$ from, and therefore, $\hat{p}_{\theta}\left(x, x^{\prime}\right)=\tilde{p}_{\theta, x^{\prime}}\left(x, x^{\prime}\right)$. But in model $4, \tilde{p}_{\theta, x^{\prime}}\left(x, x^{\prime}\right)=\mathcal{N}\left(x^{\prime} ; x, \sigma^{2}\left(x^{\prime}\right)\right)$ which is not easy to sample $x^{\prime}$ from. Therefore, we take $\hat{p}_{\theta}\left(x, x^{\prime}\right)=\mathcal{N}\left(x^{\prime} ; x, \sigma^{2}(x)\right)$.
5.2.1. Estimation of the Score Function. For each model, we fix parameter $\theta$ and estimate the score function using Algorithms 3.1 and 4.2. In Algorithm 3.1, $N \in\{3000,7000,4000,5000\}$ in the 1st, 2nd, 3rd \& 4th models, respectively. In Algorithm 4.2, $N \in\{1000,2000,1000,1500\}$ in the 1st, 2nd, 3rd \& 4th models, respectively. In both algorithms, we set the discretization level to $l=10$. In Models 1,2 and 3 , we set $\kappa=2,2.2,1.5, x_{*}=0.2,1,2$ and $\sigma=0.3,0.25,0.25$, respectively. While in model 4 , we set $x_{*}=1.3$ and $\beta=2 ; T=50$ for all 4 models ( $T=50$ in model 4 corresponds to 14.22 hours of trading).

Figure 3 summarizes the results of 56 replications of estimates of the score function for each model and for each unit time point. These simulations are implemented in parallel using 8 CPUs. The figure illustrates that both algorithms are equivalent for large $N$ and $l$ as one would expect.


Figure 3: Trajectories from the execution of Algorithms 3.1 and 4.2 for the estimation of the score function in Models 1-4.
5.2.2. Cost Comparison. We now consider comparing the costs of Algorithms 4.2 and 4.3 . We take $l_{*}$ to be 7 in the models $1,2 \& 3$ and 8 in the 4 th model. The parameters of the model are as in the previous section. The ground truth is computed at level 11 with $N=2000$ using Algorithm 4.2. We run 56 simulations of both algorithms for $L \in\left\{l_{*}-1, \cdots, 10\right\}$. For each $L$, the number of particles are carefully chosen to give similar MSE values from both algorithms. Particularly, the number of particles in Algorithm 4.2 is $N_{L}=\left\lfloor C_{1} 2^{L}\right\rfloor$ and for each level $l$ in Algorithm 4.2. In Figure 4 the number of particles is $N_{l}=\left\lfloor C_{2} 2^{L}\left(L-l_{*}+2\right) \Delta_{l}^{1 / 2+\rho}\right\rfloor$ (in models 1 to 3) and $N_{l}=\left\lfloor C_{2} 2^{L}\left(L-l_{*}+2\right) \Delta_{l}^{1 / 2+\rho} \Delta_{L}^{-\rho}\right\rfloor$ (in model 4), where $C_{1}$ and $C_{2}$ are constants. In we can observe the cost against MSE curve, that appear to follow our conjectures over algorithmic costs earlier in the article.
5.2.3. Parameter Estimation. We use Algorithms 3.1, 4.3 to estimate the parameters in each model. In Algorithm 3.1, the level of discretization, $l$, is 10 for models 1-3 and 9 for model 4, and the number of particles, $N$, is 2,000 for models $1-3$ and 2500 for model 4 . In Algorithm 4.3, we use $l_{*}=7, L=10$ and the number of particles on each level $l \in\left\{l_{*}-1, \cdots, L\right\}$ is $N_{l}=2^{L}\left(L-l_{*}+\right.$ 2) $\Delta_{l}^{1 / 2+\rho}$, where $\rho \in\{0.14,0.09,0.11\}$ in Models 1,2 and 3 , respectively. In model $4, l_{*}=8, L=9$ and the number of particles on each level $l \in\left\{l_{*}-1, \cdots, L\right\}$ is $N_{l}=1.4 \times 2^{L}\left(L-l_{*}+2\right) \Delta_{l}^{1 / 2+\rho} \Delta_{L}^{-\rho}$ where $\rho=0.1$.

Figure 5 considers Model 1. We fix $x_{*}=0.2, \sigma=0.3, \kappa=2, T=20,000$. The parameter values used to generate the data are $\left(\theta_{1}^{\star}, \theta_{2}^{\star}\right)=(-0.7,-0.5)$. For the stochastic gradient algorithm,


Figure 4: Cost per each simulation versus MSE on a log-log scale for Algorithms 4.2 and 4.3. The dashed lines are for reference.
we used an initial value $(-0.05,-1.5)$ and step-size $\alpha_{k}=k^{-0.85}$. Figure 6 considers Model 2. We fix $x_{*}=1.8, \sigma=0.25, \kappa=2.2, T=20,000$. The parameter values used to generate the data are $\left(\theta_{1}^{\star}, \theta_{2}^{\star}, \theta_{3}^{\star}\right)=(1.3,-0.5,0.18)$. For the stochastic gradient algorithm, we used initial value $(0.8,-1,0.8)$ and step-size $\alpha_{k}=k^{-0.95}$. Figure 7 considers Model 3. We fix $x_{*}=1.5, \sigma=0.25, \kappa=$ $2, T=20,000$. The parameter values used to generate the data are $\left(\theta_{1}^{\star}, \theta_{2}^{\star}, \theta_{3}^{\star}\right)=(2,1,0.45)$. For the stochastic gradient algorithm, we used an initial value ( $1.24,0.6,1.11$ ) and step size $\alpha_{k}=k^{-0.9}$. Figure 8 considers Model 4 applied to the data in Figure 2. We fix $x_{*}=1.3, \beta=2, T=11425$ (there is a rescaling of the time parameter). For the stochastic gradient algorithm, we used an initial value $(2.4,0.5,0.4)$ and step size $\alpha_{k}=k^{-0.82}$. In all cases considered (Figures $5-8$ ) our selected settings allow for an accurate estimation of the parameter values over long time periods.

Acknowledgements. AJ \& HR were supported by KAUST baseline funding. AB acknowledges support from a Leverhulme Trust Prize. DC was partially supported by EU Synergy project STUOD - DLV-856408. NK acknowledges funding by a JP Morgan A.I. Faculty award. We thank two referees and an associate editor for their comments which have greatly enhanced the article.

## REFERENCES

[1] Bain, A. \& Crisan, D. (2009). Fundamentals of Stochastic Filtering. Springer: New York.
[2] Ballesio, M., Jasra, A., von Schwerin, E. \& Tempone, R. (2020). A Wasserstein coupled particle filter for multilevel estimation. arXiv:2004.03981.
[3] Benveniste, A., Métivier, M. \& Priouret, P. (1990). Adaptive Algorithms and Stochastic Approximation. New York: Springer-Verlag.
[4] Beskos, A., Papaspiliopoulos, O., Roberts, G., Fearnhead, P. (2006). Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes (with discussion). J. R. Statist. Soc.


Figure 5: Trajectories from the execution of Algorithm 3.1 (left panel) and Algorithm 4.3 (right panel) for the estimation of $\left(\theta_{1}, \theta_{2}\right)$ from Model 1. The horizontal dashed lines in the plots show the true parameter values $\left(\theta_{1}^{\star}, \theta_{2}^{\star}\right)=(-0.7,-0.5)$.

Ser. B, 68, 333-382.
[5] Beskos, A., Crisan, D., Jasra, A., Kantas, N. \& Ruzayqat, H. (2020). Score-Based Parameter Estimation for a Class of Continuous-Time State Space Models. arXiv:2008.07803v1
[6] Bierkens, J., Van Der Meulen, F., Schauer, M. (2020). Simulation of elliptic and hypo-elliptic conditional diffusions, Advances in Applied Probability, 52, 173-212.
[7] Botha, I., Kohn, R., \& Drovandi, C. (2020). Particle methods for stochastic differential equation mixed effects models. Bayes. Anal. (to appear).
[8] Campillo, F. \& Le Gland, F. (1989). Maximum likelihood estimation for partially observed diffusions: Direct Maximization vs The EM algorithm. Stoch. Proc. Appl., 33, 245-274.
[9] Cliffe, K. A., Giles, M. B., Scheichl, R., \& Teckentrup, A. L. (2011). Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. Comput. Vis. Sci., 14, 3-15.
[10] Crisan, D. (2011). Discretizing the continuous-time filtering problem: Order of Convergence. In [11], 572-597.
[11] Crisan, D., \& Rozovskir, B. (2011). The Oxford Handbook of Nonlinear Filtering. Oxford University Press.
[12] Del Moral, P. (2004). Feynman-Kac Formulae. Springer.
[13] Del Moral, P., Doucet, A., \& Singh S. S. (2010). A backward particle interpretation of Feynman-Kac formuale. M2AN, 44, 947-975.
[14] Del Moral, P., Doucet, A., \& Singh S. S. (2010). Forward smoothing using sequential Monte Carlo, arXiv:1012.5390
[15] Etienne, M. P., Gloaguen, P., Corff, S. L., \& Olsson, J. (2020). Backward importance sampling for partially observed diffusion processes. arXiv:2002.05438.
[16] Gloaguen, P., Etienne, M. P. \& Le Corff, S. (2018). Online sequential Monte Carlo smoother for partially observed diffusion processes. EURASIP J. Adv. Sig. Proc, article 9.
[17] Jasra, A., Kamatani, K., Law K. J. H. \& Zhou, Y. (2017). Multilevel particle filters. SIAM J. Numer. Anal., 55, 3068-3096.
[18] Jasra, A., Law, K. J. H., \& Yu, F. (2020). Unbiased filtering of a class of partially observed diffusions. arXiv:2002.03747.
[19] Jasra, A., \& Yu, F. (2020). Central limit theorems for coupled particle filters. Adv. Appl. Probab. 52, 942-1001.
[20] Jasra, A., Yu, F. \& Heng, J. (2020). Multilevel particle filters for the nonlinear filtering problem in continuous time. Stat. Comp. 30, 1381-1402.
[21] Le Gland, F. and Mevel, M. (1997). Recursive identification in hidden Markov models. Proc. 36th IEEE Conf. Dec. Contr., 3468-3473.
[22] Mider, M., Schauer, M. \& van der Meulen, F. (2020). Continuous-discrete smoothing of diffusions. arXiv:1712.03807.
[23] Olsson, J. \& Westerborn, J. (2017). Efficient particle-based online smoothing in general hidden Markov


Figure 6: Trajectories from the execution of Algorithm 3.1 (left panel) and Algorithm 4.3 (right panel) for the estimation of $\left(\theta_{1}, \theta_{2}, \theta_{3}\right)$ of the second model. We used an initial value $(0.8,-1,0.8)$. The horizontal dashed lines in the plots show the true parameter values $\left(\theta_{1}^{\star}, \theta_{2}^{\star}, \theta_{3}^{\star}\right)=$ $(1.3,-0.5,0.18)$.
models: The PaRIS algorithm. Bernoulli, 23, 1951-1996.
[24] Papaspiliopoulos, O., Roberts, G. O., \& Stramer, O. (2013). Data augmentation for diffusions. J. Comp. Graph. Stat., 22, 665-688.
[25] Papaspiliopoulos, O. \& Roberts, G. (2012). Importance sampling techniques for estimation of diffusion models. Stat. Meth. Stoch. Diff. Eq., 124, 311-340.
[26] Picard, J. (1984). Approximation of nonlinear filtering problems and order of convergence. In Filtering and control of random processes, 219-236, Springer, Berlin, Heidelberg.
[27] Poyiadjis, G., Doucet, A., \& Singh, S. S. (2011). Particle approximations of the score and observed information matrix in state space models with application to parameter estimation. Biometrika, 98, 65-80.
[28] Särkkä, S., \& Sottinen, T. (2008). Application of Girsanov theorem to particle filtering of discretely observed continuous-time non-linear systems. Bayes. Anal., 3, 555-584.
[29] Schauer, M., van der Meulen, F. \& van Zanten, H. (2017). Guided proposals for simulating multidimensional diffusion bridges. Bernoulli, 23, 2917-2950.
[30] Sen, D., Thiery, A., Jasra, A. (2018). On coupling particle filters. Statist. Comp., 28, 461-475.
[31] Surace, S. C., \& Pfister, J. P. (2018). Online Maximum-Likelihood Estimation of the Parameters of Partially Observed Diffusion Processes. IEEE Transactions on Automatic Control, 64(7), 2814-2829.
[32] Talay, D. (1984). Efficient numerical schemes for the approximation of expectations of functionals of the solution of a SDE, and applications. In Filtering and control of random processes, 294-313, Springer, Berlin, Heidelberg.
[33] Thorisson, H. (2000). Coupling, stationarity, and regeneration. Springer:New York.
[34] van der Meulen, F., \& Schauer, M. (2017). Bayesian estimation of discretely observed multi-dimensional diffusion processes using guided proposals. Elec. J. Stat., 11, 2358-2396.
[35] Whitaker, G. A., Golightly, A., Boys, R. J., \& Sherlock, C. (2017). Improved bridge constructs for stochastic differential equations. Stat. Comp., 27, 885-900.
[36] Yonekura, S. \& Beskos, A. (2020). Online smoothing for diffusion processes observed with noise. arXiv: 200312247.


Figure 7: Trajectories from the execution of Algorithm 3.1 (left panel) and Algorithm 4.3 (right panel) for the estimation of $\left(\theta_{1}, \theta_{2}, \theta_{3}\right)$ from Model 3. The horizontal dashed lines in the plots show the true parameter values $\left(\theta_{1}^{\star}, \theta_{2}^{\star}, \theta_{3}^{\star}\right)=(2,1,0.45)$.


Figure 8: Trajectories from the execution of Algorithm 3.1 (left panel) and Algorithm 4.3 (right panel) for the estimation of $\left(\theta_{1}, \theta_{2}, \theta_{3}\right)$ from Model 4.


[^0]:    *Department of Statistical Science, University College London, London, WC1E 6BT, UK. (a.beskos@ucl.ac.uk).
    ${ }^{\dagger}$ Department of Mathematics, Imperial College London, London, SW7 2AZ, UK. (d.crisan@imperial.ac.uk, n.kantas@imperial.ac.uk).
    ${ }^{\ddagger}$ Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology, Thuwal, 23955, KSA. (ajay.jasra@kaust.edu.sa, hamza.ruzayqat@kaust.edu.sa).

