This is a self-archived version of an original article. This version may differ from the original in pagination and typographic details.

Author(s): Chada, Neil K.; Franks, Jordan; Jasra, Ajay; Law, Kody J.; Vihola, Matti

Title: Unbiased Inference for Discretely Observed Hidden Markov Model Diffusions

Year: 2021

Version: Accepted version (Final draft)

Copyright: © 2021, Society for Industrial and Applied Mathematics

Rights: In Copyright

Rights url: <http://rightsstatements.org/page/InC/1.0/?language=en>

Please cite the original version:

Chada, N. K., Franks, J., Jasra, A., Law, K. J., & Vihola, M. (2021). Unbiased Inference for Discretely Observed Hidden Markov Model Diffusions. *SIAM/ASA Journal on Uncertainty Quantification*, 9, 763-787. <https://doi.org/10.1137/20M131549X>

UNBIASED INFERENCE FOR DISCRETELY OBSERVED HIDDEN MARKOV MODEL DIFFUSIONS

NEIL K. CHADA, JORDAN FRANKS, AJAY JASRA, KODY J. H. LAW & MATTI VIHOLA

ABSTRACT. We develop a Bayesian inference method for diffusions observed discretely and with noise, which is free of discretisation bias. Unlike existing unbiased inference methods, our method does not rely on exact simulation techniques. Instead, our method uses standard time-discretised approximations of diffusions, such as the Euler-Maruyama scheme. Our approach is based on particle marginal Metropolis-Hastings, a particle filter, randomised multilevel Monte Carlo, and importance sampling type correction of approximate Markov chain Monte Carlo. The resulting estimator leads to inference without a bias from the time-discretisation as the number of Markov chain iterations increases. We give convergence results and recommend allocations for algorithm inputs. Our method admits a straightforward parallelisation, and can be computationally efficient. The user-friendly approach is illustrated on three examples, where the underlying diffusion is an Ornstein-Uhlenbeck process, a geometric Brownian motion, and a $2d$ non-reversible Langevin equation.

1. INTRODUCTION

Hidden Markov models (HMMs) are widely used in real applications, for example, for financial and physical systems modeling; see [6]. We focus on the case where the hidden Markov chain arises from a diffusion process that is observed with noise at some number of discrete points in time; see e.g. [31]. The parameters associated to the model are static and assigned a prior density. Bayesian inference involves expectations with respect to (w.r.t.) the joint posterior distribution of parameters and states, and is important in problems of model calibration and uncertainty quantification. A difficult aspect of Bayesian inference for these models is simulation or evaluation of the diffusion dynamics. Unless the transition probability is explicitly known (see Section 4.4 of [25]), one often resorts to time discretisation, leading to biased inference. If one is to seek inference where such discretisation is avoided, there are broadly two schemes (an exception is [13], which we shall discuss below) which one can follow, the first are the elegant exact simulation methods in [3, 4, 11, 34] or the debiasing methods of [26, 28], which is the direction followed in this article. We call inference when there is no time-discretisation error, *unbiased inference* and this is the objective of this article.

Approaches to unbiased inference using the afore-mentioned exact simulation methodology apply for a certain class of diffusions. First, the existence of the Lamperti transformation, after which the process has unit diffusion matrix, and second, the drift of the transformed process has to be of gradient form. Although this includes some important diffusion processes, often, these conditions do not hold for multivariate diffusion processes; this limits the scope of the application of these novel schemes. A more recent and general approach can be found in [13], which focuses on a type of continuous-time importance sampling method for continuous-time Markov processes, including diffusion processes. The

2010 *Mathematics Subject Classification.* 65C05 (primary); 60H35, 65C35, 65C40 (secondary).

Key words and phrases. Diffusion, importance sampling, Markov chain/multilevel/sequential Monte Carlo.

method is, in essence, a type of continuous-time sequential importance sampling algorithm that produces a signed approximation of laws of diffusion processes. Whilst the methodology applies to a reasonably wide class of diffusions, the signed approximation can introduce several algorithmic issues. This includes a large cost associated to the final time of the diffusion process, which at the very least can lead to linear-in-time errors. This latter issue is problematic in our context, as if the methodology is used for HMMs, the time will relate to the number of data and the errors reported in [13] will be prohibitive for our application. It should also be noted that this methodology requires a good proposal based upon a process which is analytically soluble. This limits the applicability of the method, which will not be the case in our context. As the utility of the method has not been explored for the problem of interest in this article, we have considered an alternative method. We proceed with an Euler-Maruyama time-discretisation (see [25]), referred henceforth as *Euler* which is generally applicable and is combined with the debiasing schemes of [26, 28].

Traditional inference approaches based on time-discretisations face a trade-off between bias and computational cost. Once the user has decided on a suitably fine discretisation size, one can run, for example, the particle marginal Metropolis-Hastings (PMMH) [2]. This algorithm uses a particle filter (PF) (see [8]), where proposals between time points are generated by the approximation scheme, and ultimately accepted or rejected according to a Metropolis-Hastings type acceptance ratio; see [19]. As the discretisation size adopted must be quite fine, a PMMH algorithm can be computationally intensive.

To deal with the computational cost of PMMH, [22] develop a PMMH based method which uses (deterministic) multilevel Monte Carlo (dMLMC) [17, 20]. The basic premise of MLMC is to introduce a telescoping sum representation of the posterior expectation associated to the most precise time discretisation. Then, given an appropriate coupling of posteriors with ‘consecutive’ time discretisations, the cost associated to a target mean square error is reduced, relative to exact sampling from the most precise (time-discretised) posterior. In the HMM diffusion context, the standard MLMC method (for diffusions without observations) is not applicable, so based upon a PF coupling approach and PMMH, an MLMC method is devised in [22, 23], which achieves fine-level, though biased, inference.

1.1. Method. The unbiased and computationally efficient inference method suggested in this paper is built firstly on PMMH, using Euler type discretisations, but using a PMMH targeting a coarse-level model, which is less computationally expensive. This does not yield unbiased inference yet, but it can be achieved by an importance sampling (IS) type correction; see [33].

We suggest an IS type correction that is based on a single-term (randomised) MLMC type estimator [26, 28] and the PF coupling approach of [22]. The rMLMC correction is based on randomising the running level in the multilevel context of a certain PF, which we refer to as the ‘delta PF (Δ PF)’ (Algorithm 3). In short, the Δ PF uses the PF coupling introduced in [22], but here an estimator is used for unbiased estimation of the difference of *unnormalised* integrals corresponding to two consecutive discretisation levels, over the latent states with parameter held fixed (see Section 2). In [22], each term in the difference of PMMH averages is individually self-normalised (at each level) because of the unknown normalising constants.

The resulting IS type estimator leads to unbiased inference over the joint posterior distribution, and is highly parallelisable, as the more costly (randomised) Δ PF corrections may be performed independently *en masse* given the PMMH base chain output. We are also able to suggest optimal choices for algorithm inputs in a straightforward manner (Recommendation 1 and Figure 1). This is because there is no bias, and therefore the difficult

cost-variance-bias trade-off triangle associated with dMLMC is not present. Besides being unbiased and efficient, our method is user-friendly, as it is a combination of well-known and relatively straightforward components: PMMH, Euler approximations, PF, rMLMC, and an IS type estimator. For more about the strengths of the method, see Remark 11 later, as well as [15, 33] for more discussion about IS (type) estimators based on approximate Markov chain Monte Carlo (MCMC).

Key to verifying consistency of the method is a finite variance assumption for the $r\Delta PF$ estimator. We verify a parameter-uniform bound for the variance under a simple set of HMM diffusion conditions in Section 3. Note, however, that consistency of our method is likely to hold more generally. This is in contradistinction to methods based on exact simulation, which require analytically tractable transformations to unit covariance diffusion term and computable bounds in the rejection sampler, in order to even apply the method (see for example the review in the recent preprint [34]).

We consider a non-reversible Langevin equation in Section 6, where, to the authors' best knowledge, exact simulation is not applicable. If an exact simulation method is applicable, the obvious question arises whether our method or the exact simulation method should be applied. The efficiency of exact simulation type methods is dependent upon several and different factors than our method. These factors for exact simulation include proper tuning and tight computable bounds for the rejection sampler. In an ideal scenario for exact simulation, a method based on exact simulation is likely to perform better than our method. However, in the reverse case, our method can perform better, if the efficiency of exact simulation is poor. For instance, the efficiency of exact simulation decreases to zero as the analytically computed upper bound of the IS weight used in the rejection sampler increases to infinity.

We remark that in principle our algorithm may require simulations with arbitrarily fine discretisation sizes and corresponding arbitrarily large cost; this can be infeasible. However, it should be noted that the user specifies the chance that this might occur, so one can ensure that the probability of 'very expensive' simulations is arbitrarily small. In addition, our method typically has finite expected cost, or cost that is finite with high probability: we give conditions (see Section 3), typical elsewhere in the context of HMM diffusions [8, 25], which ensure this (see Section 5).

Although we have mostly in mind the case of Euler approximation schemes for the diffusion dynamics approximation, which are generally implementable, other schemes could be possibly be used as well; see [16]. However, suitable couplings for these schemes in dimensions greater than one may not be trivial. For the sake of theory and proof of consistency, ideally these would have also known weak and strong order convergence rates; see [25]. Indeed, assuming a coupling exists, such higher-order schemes can improve convergence of our method (see Sections 5 and 6). More generally, our approach based on PMMH or other approximate MCMC, increasingly fine families of approximations, MLMC, and IS correction, could be applied beyond the HMM diffusion context, for example, to HMM jump-diffusions; see [24].

1.2. Outline. Section 2 introduces the aforementioned ΔPF (Algorithm 2) and subsequently discusses some applications of randomisation techniques. The theoretical properties of the ΔPF in the HMM diffusion context are summarised in Section 3. Section 4 presents the suggested IS type estimator (Algorithm 4), based on PMMH with rMLMC (i.e. $r\Delta PF$) correction, and details its consistency and a corresponding central limit theorem (CLT). Section 5 suggests suitable allocations in the ΔPF based on rMLMC efficiency considerations. The numerical experiments in Section 6 illustrate our method in practice in the setting of

an Ornstein-Uhlenbeck process, geometric Brownian motion, and a non-reversible Langevin equation. Proofs for the technical results of Sections 3, 4 and 5 are given in Appendix A, B and C, respectively.

1.3. Notation. Let $(E, \mathcal{E})$ be a measurable space. Functions $\varphi : E \rightarrow \mathbb{R}$ will be assumed measurable. We denote by $\mathcal{P}(E)$ the collection of probability measures on $(E, \mathcal{E})$, and by $\mathcal{B}_b(E)$ the set of $\varphi : E \rightarrow \mathbb{R}$ with $\|\varphi\| := \sup_{x \in E} |\varphi(x)| < \infty$. For a measure μ on $(E, \mathcal{E})$, we set $\mu(\varphi) := \int_E \varphi(x) \mu(dx)$ whenever well-defined. For $K : E \times E \rightarrow [0, 1]$ a Markov kernel and $\mu \in \mathcal{P}(E)$, we set $\mu K(dy) := \int_E \mu(dx) K(x, dy)$, and $K(\varphi)(x) := \int_E \varphi(y) K(x, dy)$, whenever well-defined. We use the convention $\prod_{\emptyset} := 1$, and $p:q := \{r \in \mathbb{Z} : p \leq r \leq q\}$.

2. DELTA PARTICLE FILTER FOR UNBIASED ESTIMATION OF LEVEL DIFFERENCES

Consider the (Itô) diffusion process

$$(1) \quad dZ_t = a_\theta(Z_t)dt + b_\theta(Z_t)dW_t, \quad t \geq 0,$$

with $Z_t \in \mathbf{X} := \mathbb{R}^d$, model parameter $\theta \in \mathbf{T}$, $a : \mathbf{X} \times \mathbf{T} \rightarrow \mathbf{X}$, $b : \mathbf{X} \times \mathbf{T} \rightarrow \mathbf{X} \times \mathbf{X}$, $\{W_t\}_{t \geq 0}$ a d -dimensional Brownian motion, and the initial value $Z_0 = z_0 \in \mathbf{X}$ a fixed value. We suppose that there are data $\{Y_p = y_p\}_{p=0}^n$, $y_p \in \mathbb{R}^m$, which are observed at equally spaced discrete times, $1:n+1$ for simplicity. We shall consider the discrete time skeleton of the diffusion (1) $Z_1, \dots, Z_{n+1}$, where we shall set $X_p = Z_{p+1}$, for $p = 0:n$. The Markov transition between X_{p-1} and X_p , $p = 1:n$, is given by the transition kernel $M^{(\theta, \infty)}(x_{p-1}, dx_p)$ of the diffusion process over unit time, with initial distribution $\eta_0^{(\theta, \infty)}(dx_0) = M^{(\theta, \infty)}(z_0, dx_0)$. It is assumed that conditional on X_p , Y_p is independent of random variables $\{X_i, Y_i\}_{i \neq p}$ and has density $g_\theta(y_p | x_p) =: G_p^{(\theta)}(x_p)$. Setting $M_0^{(\theta, \infty)} = \eta_0^{(\theta, \infty)}$ and $M_p^{(\theta, \infty)} = M^{(\theta, \infty)}$, $p = 1:n$, the resulting pair $(M_p^{(\theta, \infty)}, G_p^{(\theta)})$ defines the HMM diffusion, and is an example of a so-called *Feynman-Kac model* (see [8]) described below. As the results of this section can just as easily be stated in terms of Feynman-Kac models, we do so in the following, which shows the generality of our approach.

Remark 1. In many situations of practical interest, $M^{(\theta, \infty)}(x_{p-1}, dx_p)$, $p = 1:n$ exists, but one cannot even simulate from it and/or evaluate a non-negative unbiased estimator of it. One can often consider the HMM diffusion where one works with a time discretisation of $M^{(\theta, \infty)}(x_{p-1}, dx_p)$. For instance for $\ell \in \{0\} \times \mathbb{N}$, $h_\ell = 2^{-\ell}$, one could consider the Euler approximation for $t = p, p + h_\ell, \dots, p + 1 - h_\ell$, with $Z_p = x_{p-1}$ given

$$Z_{t+h_\ell}^\ell = Z_t^\ell + a_\theta(Z_t^\ell)h_\ell + b_\theta(Z_t^\ell)\delta W_{t+h_\ell}^\ell,$$

where $\delta W_{t+h_\ell}^\ell \stackrel{i.i.d.}{\sim} \mathcal{N}(0, h_\ell)$ (Gaussian distribution 0 mean, covariance $h_\ell I$) and one will set $X_p = Z_{p+1}^\ell$. The induced transition kernel over unit time is written $M^{(\theta, \ell)}(x_{p-1}, dx_p)$. A similar remark can be made for $\eta_0^{(\theta, \infty)}$ with discretisation $\eta_0^{(\theta, \ell)}$.

2.1. Particle filters. A Feynman-Kac model (M_n, G_n) on spaces $(E_n, \mathcal{E}_n)$ arises when

- (i) $M_n(x_{0:n-1}, dx_n)$ are (regular) probability ‘transition’ kernels from $E_{0:n-1}$ to E_n for $n \geq 1$, and $M_0(dx_0) := \eta_0(dx_0) \in \mathcal{P}(E_0)$, and
- (ii) $G_n(x_{0:n})$ are $[0, \infty)$ -valued (measurable) ‘potential’ functions for $n \geq 0$.

The Particle filter (Algorithm 1) (see [8]) generates sets of samples and weights corresponding to the Feynman-Kac model, which for $\varphi : E_{0:n} \rightarrow \mathbb{R}$ lead to an unbiased estimator for the

(unnormalised) *smoother* $\gamma_n(G_n\varphi)$, defined here in terms of the (unnormalised) *predictor*

$$(2) \quad \gamma_n(\varphi) := \int \varphi(x_{0:n}) \left(\prod_{t=0}^{n-1} G_t(x_{0:t}) \right) \eta_0(dx_0) \prod_{t=1}^n M_t(x_{0:t-1}, dx_t).$$

We remark that Step **iii** of Algorithm 1 refers to the resampling step, which can be multinomial, residual, stratified or systematic; see e.g. [6, 9].

Algorithm 1 Particle filter for a Feynman-Kac model.

Input: $(M_{0:n}, G_{0:n}) := (M_t, G_t)_{t=0:n}$ and N the number of particles.

- (i) For $i = 1:N$ sample $x_0^{(i)} \sim \eta_0(\cdot)$ and set $\mathbf{x}_0^{(i)} := x_0^{(i)}$.
- (ii) For $i = 1:N$ compute $\omega_0^{(i)} := G_0(\mathbf{x}_0^{(i)})$ and set $\bar{\omega}_0^{(i)} := \omega_0^{(i)} / \omega_0^*$ where $\omega_0^* = \sum_{j=1}^N \omega_0^{(j)}$.

For $t = 1:n$, do:

- (iii) Given $\bar{\omega}_{t-1}^{(1:N)}$, sample $A_{t-1}^{(1:N)} \in \{1, \dots, N\}^N$ satisfying $\mathbb{E}[\sum_{j=1}^N \mathbf{1}\{A_{t-1}^{(j)} = k\}] = N\bar{\omega}_{t-1}^{(k)}$ for all $k \in 1:N$.
- (iv) For $i = 1:N$ sample $x_t^{(i)} \sim M_t(\mathbf{x}_{t-1}^{A_{t-1}^{(i)}}, \cdot)$ and set $\mathbf{x}_t^{(i)} = (\mathbf{x}_{t-1}^{A_{t-1}^{(i)}}, x_t^{(i)})$.
- (v) For $i = 1:N$ compute $\omega_t^{(i)} := G_t(\mathbf{x}_t^{(i)})$ and set $\bar{\omega}_t^{(i)} := \omega_t^{(i)} / \omega_t^*$ where $\omega_t^* := \sum_{j=1}^N \omega_t^{(j)}$.

Set for $i = 1:N$, $\mathbf{X}_n^{(i)} := \mathbf{x}_n^{(i)}$. If $\omega_t^* > 0$, for $i = 1:N$ set $V^{(i)} := \bar{\omega}_n^{(i)} \prod_{t=0}^n \frac{1}{N} \omega_t^*$, otherwise, for $i = 1:N$ set $V^{(i)} = 0$.

Output: $(V^{(1:N)}, \mathbf{X}^{(1:N)})$.

Proposition 2. Suppose that $\varphi : E_{0:n} \rightarrow \mathbb{R}$ is such that $\gamma_n(G_n\varphi) < \infty$. Then, the output of Algorithm 1 satisfies

$$\mathbb{E} \left[\sum_{i=1}^N V^{(i)} \varphi(\mathbf{X}^{(i)}) \right] = \gamma_n(G_n\varphi).$$

Proposition 2 is a restatement of Theorem 7.4.2 of [8] in case $A_{t-1}^{(i)}$ are sampled independently (‘multinomial resampling’). The extension to the general unbiased case is straightforward; see [33].

2.2. Level difference estimation. Suppose that we have two Feynman-Kac models (M_n^F, G_n^F) and (M_n^C, G_n^C) defined on common spaces $(E_n, \mathcal{E}_n)$. The models correspond to ‘finer’ and ‘coarser’ Euler type discretised HMM diffusions. We are interested in estimating (unbiasedly) the difference

$$(3) \quad \gamma_n^F(G_n^F\varphi) - \gamma_n^C(G_n^C\varphi).$$

If the models are close to each other, as they will be in the multilevel (diffusion) context, we would like the estimator also to be typically small. In many contexts, if one can estimate the difference using a coupling, it is possible to obtain a variance reduction. The particular coupling approach we use here is based on using a combined Feynman-Kac model as in [22], which provides a simple, general and effective coupling of PFs, and which we will use to estimate the level difference of unnormalised smoother 3.

Hereafter, we denote $\tilde{x}_n = (\tilde{x}_n^F, \tilde{x}_n^C) \in E_n \times E_n$, and for $\tilde{x}_{0:n} = (\tilde{x}_0, \dots, \tilde{x}_n) \in E_0^2 \times \dots \times E_n^2$, we set $\tilde{x}_{0:n}^s := (\tilde{x}_0^s, \dots, \tilde{x}_n^s) \in E_{0:n}$ for $s \in \{F, C\}$.

Assumption 3. Suppose that $(\check{M}_t, \check{G}_t)$ is a Feynman-Kac model on the product spaces $(E_t \times E_t, \mathcal{E}_t \otimes \mathcal{E}_t)$, such that:

- (i) $\check{M}_t$ is a coupling of the probability measures M_t^F and M_t^C , i.e. for all $A \in \mathcal{E}_t$, we have

$$\int_{A \times E_t} \check{M}_t(\check{x}_{0:t-1}, d\check{x}_t) = M_t^F(\check{x}_{0:t-1}^F, A), \quad \int_{E_t \times A} \check{M}_t(\check{x}_{0:t-1}, d\check{x}_t) = M_t^C(\check{x}_{0:t-1}^C, A),$$

and for $A \in \mathcal{E}_0$, we have $\check{\eta}_0(A \times E_0) = \eta_0^F(A)$ and $\check{\eta}_0(E_0 \times A) = \eta_0^C(A)$.

- (ii) $\check{G}_t(\check{x}_{0:t}) := \frac{1}{2} [G_t^F(\check{x}_{0:t}^F) + G_t^C(\check{x}_{0:t}^C)]$.

Algorithm 2 presents a methodology to unbiasedly estimate the level differences (3). In the context of hidden Markov model diffusions, we explain in Remark 5 how to satisfy Assumption 3.

Algorithm 2 Delta particle filter (Δ PF) for unbiased estimation of level differences.

Input: $(\check{M}_{0:n}, \check{G}_{0:n})$ and N the number of particles.

- (i) Run Algorithm 1 with $(\check{M}_{0:n}, \check{G}_{0:n}, N)$, outputting $(\check{V}^{(1:N)}, \check{\mathbf{X}}^{(1:N)})$.
(ii) Compute $(V^{(1:2N)}, \mathbf{X}^{(1:2N)})$ where

$$(V^{(i)}, \mathbf{X}^{(i)}) := \begin{cases} (\check{V}^{(i)} w^F(\check{\mathbf{X}}^{(i)}), \check{\mathbf{X}}^{(i)F}) & i = 1:N, \\ (-\check{V}^{(i-N)} w^C(\check{\mathbf{X}}^{(i-N)}), \check{\mathbf{X}}^{(i-N)C}) & i = (N+1):2N, \end{cases}$$

$$\text{and } w^F(\check{x}_{0:n}) := \frac{\prod_{t=0}^n G_t^F(\check{x}_{0:t}^F)}{\prod_{t=0}^n \check{G}_t(\check{x}_{0:t})} \text{ and } w^C(\check{x}_{0:n}) := \frac{\prod_{t=0}^n G_t^C(\check{x}_{0:t}^C)}{\prod_{t=0}^n \check{G}_t(\check{x}_{0:t})}.$$

Output: $(V^{(1:2N)}, \mathbf{X}^{(1:2N)})$.

Proposition 4. *Under Assumption 3, the output of Algorithm 2 satisfies*

$$\mathbb{E} \left[\sum_{i=1}^{2N} V^{(i)} \varphi(\mathbf{X}^{(i)}) \right] = \gamma_n^F(G_n^F \varphi) - \gamma_n^C(G_n^C \varphi),$$

whenever both integrals on the right are well-defined and finite.

Proof. Applying Assumption 3 one can use the unbiasedness property of PF in Algorithm 1, to yield

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^N V^{(i)} \varphi(\mathbf{X}^{(i)}) \right] &= \int w^F(\check{x}_{0:n}) \varphi(\check{x}_{0:n}^F) \left(\prod_{t=0}^n \check{G}_t(\check{x}_{0:t}) \right) \check{\eta}(dx_0) \prod_{t=1}^n \check{M}_t(\check{x}_{0:t-1}, d\check{x}_t) \\ &= \int \varphi(\check{x}_{0:n}^F) \left(\prod_{t=0}^n G_t^F(\check{x}_{0:t}^F) \right) \eta_0^F(dx_0) M_t^F(\check{x}_{0:t-1}^F, d\check{x}_t^F) = \gamma_n^F(G_n^F \varphi), \end{aligned}$$

where, specifically, Assumption 3ii guarantees $\check{G}_t > 0$ whenever $G_t^F > 0$, and Assumption i implies the marginal law of $\prod_{t=0}^n \check{M}_t$ is $\prod_{t=0}^n M_t^F$. Similarly, $\mathbb{E} \left[\sum_{i=N+1}^{2N} V^{(i)} \varphi(\mathbf{X}^{(i)}) \right] = -\gamma_n^C(G_n^C \varphi)$. $\square$

Remark 5. Regarding Algorithm 2:

- (i) In the hidden Markov model diffusion context, one could consider F to correspond to an Euler discretisation at level ℓ step-size $h_\ell = 2^{-\ell}$ and C to correspond to an Euler discretisation at level $\ell - 1$ step-size $h_{\ell-1} = 2^{-(\ell-1)}$. The couplings $\check{M}_p$, $p = 1:n$, (of the two Euler transitions over unit time - see Remark 1) could be based on using the same underlying

Brownian motion; see [25]. That is, for $t = p, p + 2h^\ell, \dots, p + 1 - 2h^\ell$, $(Z_p^\ell, Z_p^{\ell-1}) = (x_{p-1}^\ell, x_{p-1}^{\ell-1})$ given,

$$\begin{aligned} Z_{t+h_\ell}^\ell &= Z_t^\ell + a_\theta(Z_t^\ell)h_\ell + b_\theta(Z_t^\ell)\delta W_{t+h_\ell}^\ell \\ Z_{t+2h_\ell}^\ell &= Z_{t+h_\ell}^\ell + a_\theta(Z_{t+h_\ell}^\ell)h_\ell + b_\theta(Z_{t+h_\ell}^\ell)\delta W_{t+2h_\ell}^\ell, \end{aligned}$$

with $\delta W_{t+kh_\ell}^\ell \stackrel{i.i.d.}{\sim} \mathcal{N}(0, h_\ell)$, $k = 1 : 2$, then we can use

$$Z_{t+h_{\ell-1}}^{\ell-1} = Z_t^{\ell-1} + a_\theta(Z_t^{\ell-1})h_{\ell-1} + b_\theta(Z_t^{\ell-1})(\delta W_{t+h_\ell}^\ell + \delta W_{t+2h_\ell}^\ell),$$

for the coarser Euler discretisation. We set $(X_p^\ell, X_p^{\ell-1}) = (z_{p+1}^\ell, z_{p+1}^{\ell-1})$. A similar remark can be made for the initialisation. The potentials G_t^ℓ and $G_t^{\ell-1}$ are then simply the conditional likelihood functions $G_t^\ell = G_t^{\ell-1} = G_t^{(\theta)}$.

- (ii) The choice of $\check{G}_t$ in Assumption 3ii provides a safe ‘balance’ in between the approximations, as w^F and w^C are upper bounded by 2^{n+1} . Indeed, the coupled Feynman-Kac model can be thought as an ‘average’ of the two extreme cases—with the choice $\check{G}_t(x_{0:t}) = G_t^F(\check{x}_{0:t}^F)$ the coupled PF would coincide marginally with the Feynman-Kac model with dynamics M_t^F . What is the optimal choice for $\check{G}_t$ is an interesting question.
- (iii) Clearly, the choice of $\check{G}_{0:t}$ can be made also in other ways. It is sufficient for unbiasedness to choose $\check{G}_t(\check{x}_{0:t})$ such that it is strictly positive whenever either the $G_t^F(\check{x}_{0:t}^F)$ or $G_t^C(\check{x}_{0:t}^C)$ product is positive, but choices which make w^F and w^C bounded are safer, for instance $\check{G}_{0:t}(\check{x}_{0:t}) = \max\{G_t^F(\check{x}_{0:t}^F), G_t^C(\check{x}_{0:t}^C)\}$. This was the original choice made in [22] for approximation of normalised smoother differences. This PF coupling approach based on change of measure and weight corrections w^F and w^C , has been further used also, for example, in [23].
- (iv) Later, in the HMM diffusion context, we set $G_t^F = G_t^C$, corresponding to common observational densities, but the method is also of interest with differing potentials.

2.3. Unbiased latent inference. We show here how the randomisation techniques of [26, 28] can be used with the output of Algorithm 1 and 2 to provide an unbiased estimator according to the true model, even though the PFs are only run according to approximate models. Let us index the transitions $M_p^{(\ell)}$ and potentials $G_p^{(\ell)}$ by $\ell \geq 0$. They are assumed throughout to be increasingly refined approximations, in the (weak) sense that

$$(4) \quad \gamma_n^{(\ell)}(G_n^{(\ell)}\varphi) \longrightarrow \gamma_n^{(\infty)}(G_n^{(\infty)}\varphi), \quad \text{as } \ell \rightarrow \infty,$$

for all $\varphi \in \mathcal{B}_b(E_{0:n})$, where

$$\gamma_n^{(\ell)}(\varphi) := \int \varphi(x_{0:n}) \left(\prod_{t=0}^{n-1} G_t^{(\ell)}(x_{0:t}) \right) \eta_0^{(\ell)}(dx_0) \prod_{t=1}^n M_t^{(\ell)}(x_{0:t-1}, dx_t).$$

In Assumption 3 we set symbols (F, C) to be $(\ell, \ell - 1)$ for $\ell \geq 1$. We will write the potentials and kernels of the coupled Feynman-Kac model (in the sense of Assumption 3) as $(\check{M}_{0:n}^{(\ell)}, \check{G}_{0:n}^{(\ell)})$. As a result of 7, Algorithm 3 can provide unbiased estimation of $\gamma_n^{(\infty)}(G_n^{(\infty)}\varphi)$, leading to unbiased inference w.r.t. the normalised smoother

$$\varphi \mapsto \frac{\gamma_n^{(\infty)}(G_n\varphi)}{\gamma_n^{(\infty)}(G_n)} =: \hat{\gamma}_n^{(\infty)}(\varphi),$$

which is stated as Proposition 8 below.

We remark that in step (iii) of Algorithm 3, in principle, one may have to run Algorithm 2 for arbitrarily large L . However, it should be noted that the user specifies the probability $\mathbf{p} = (p_\ell)_{\ell \in \mathbb{N}}$, so one can ensure the probability of simulating ‘very large’ values of L is arbitrarily small.

Algorithm 3 Unbiased estimator based on PF and rΔPF.

Input: $((M_{0:n}^{(\ell)}, G_{0:n}^{(\ell)}))_{\ell \in \{0\} \cup \mathbb{N}}$, N the number of particles and probability $\mathbf{p} = (p_\ell)_{\ell \in \mathbb{N}}$.

- (i) Run Algorithm 1 with $(M_{0:n}^{(0)}, G_{0:n}^{(0)}, N)$, outputting $(V^{(1:N)'}, \mathbf{X}^{(1:N)'})$.
- (ii) Sample $L \sim \mathbf{p}$, independently from the other random variables.
- (iii) Run Algorithm 2 with $(\tilde{M}_{0:n}^{(L)}, \tilde{G}_{0:n}^{(L)}, N)$, outputting $(V^{(1:2N)}, \mathbf{X}^{(1:2N)})$.

Output: $((V^{(1:N)'}, \mathbf{X}^{(1:N)'}) , (V^{(1:2N)}, \mathbf{X}^{(1:2N)}), L)$.

Assumption 6. Assumption 3 holds, $\mathbf{p} = (p_\ell)_{\ell \in \mathbb{N}}$ is a probability on $\mathbb{N} := \mathbb{Z}_{\geq 1}$ with $p_\ell > 0$ for all $\ell \geq 1$, $g : E_{0:n} \rightarrow \mathbb{R}$ is a function, and

$$(5) \quad s_g := \sum_{\ell \geq 0} \frac{\mathbb{E} \Delta_\ell^2(g)}{p_\ell} < \infty,$$

where

$$(6) \quad \Delta_\ell(g) := \sum_{i=1}^{2N} V^{(i)} g(\mathbf{X}^{(i)}),$$

is formed from the output $(V^{(1:2N)}, \mathbf{X}^{(1:2N)})$ of Algorithm 2 with $(\tilde{M}_{0:n}^{(\ell)}, \tilde{G}_{0:n}^{(\ell)}, N)$.

Lemma 7. Under Assumption 6, the estimator

$$(7) \quad \zeta(g) := \sum_{i=1}^N V^{(i)'} g(\mathbf{X}^{(i)'}) + \frac{1}{p_L} \Delta_L(g),$$

formed from the output of Algorithm 3 satisfies

$$\mathbb{E}[\zeta(g)] = \gamma_n^{(\infty)}(G_n^{(\infty)} g),$$

whenever $\gamma_n^{(0)}(G_n g)$ and $\gamma_n^{(\infty)}(G_n g)$ are both finite.

Proof. Under Assumption 6, we have (see [28, 32])

$$\mathbb{E}[p_L^{-1} \Delta_L(g)] = \gamma_n^{(\infty)}(G_n^{(\infty)} g) - \gamma_n^{(0)}(G_n^{(0)} g),$$

so the result follows by Proposition 2 and linearity of the expectation. $\square$

The following suggests a fully parallelisable algorithm for unbiased inference over the normalised smoother, and is an unbiased alternative to the particle independent Metropolis-Hastings (PIMH) [2] run at some fine level of discretisation.

Proposition 8. Suppose $\mathbf{p}$ on $\mathbb{N}$ satisfies Assumption 6 for functions $g \in \{1, \varphi\}$, with $\gamma_n^{(0)}(G_n^{(0)} g)$ and $\gamma_n^{(\infty)}(G_n^{(\infty)} g)$ finite, and $\gamma_n^{(\infty)}(G_n^{(\infty)}) > 0$. For each $k \in \{1:m\}$, if one runs independently Algorithm 3, forming $\zeta_k(g)$ from the output as in 7 for each k , then

$$E_{m,N,\mathbf{p}}(\varphi) := \frac{\sum_{k=1}^m \zeta_k(\varphi)}{\sum_{k=1}^m \zeta_k(1)} \xrightarrow{m \rightarrow \infty} \hat{\eta}_n^{(\infty)}(\varphi) \quad \text{almost surely.}$$

Moreover, with $\bar{\varphi} := \varphi - \hat{\eta}_n^{(\infty)}(\varphi)$,

$$\sqrt{m}[E_{m,N,\mathbf{p}}(\varphi) - \hat{\eta}_n^{(\infty)}(\varphi)] \xrightarrow{m \rightarrow \infty} \mathcal{N}(0, \sigma^2) \quad \text{in distribution,}$$

where

$$\sigma^2 = \frac{s_{\bar{\varphi}} - (\gamma^{(\infty)}(G_n^{(\infty)}\bar{\varphi}) - \gamma^{(0)}(G_n^{(0)}\bar{\varphi}))^2}{[\gamma^{(\infty)}(G_n^{(\infty)})]^2}.$$

The above result follows directly from the results of Section 4. It can also be seen as a multilevel version of Proposition 23 of [33], with straightforward estimators for σ^2 . See Section 5 for suggested choices for $\mathbf{p}$ and number of particles run at each level.

3. A VARIANCE BOUND FOR THE DELTA PARTICLE FILTER

In this section we give theoretical results for the Δ PF (Algorithm 2) in the setting of HMM diffusions, which can be used to verify finite variance and therefore consistency of related estimators. In particular, Corollary 10 below can be used to verify Assumption 6.

3.1. Hidden Markov model diffusions. We consider an HMM diffusion and corresponding Feynman-Kac model as in Section 2. We omit θ from the notation in the following, which is allowed as the remaining conditions and results in this Section 3 will hold uniformly in θ (i.e. any constants are independent of θ). The following will be assumed throughout.

Condition (D). The coefficients $a^j, b^{j,k}$ are twice differentiable for $j, k = 1, \dots, d$, and

- (i) **uniform ellipticity:** $b(x)b(x)^T$ is uniformly positive definite;
- (ii) **globally Lipschitz:** there is a $C > 0$ such that $|a(x) - a(y)| + |b(x) - b(y)| \leq C|x - y|$ for all $x, y \in \mathbb{R}^d$;

Let $M^{(\infty)}(x, dy) =: M_p^{(\infty)}(x, dy)$ for $p = 0:n$ denote the Markov transition of the unobserved diffusion 1, i.e. the distribution of the solution X_1 of 1 started at $X_0 = x$. With similar setup from Section 2, with $E_{0:n} := \mathbf{X}^{n+1}$, we have that 2 takes the form

$$\gamma_n^{(\infty)}(\varphi) = \int \varphi(x_{0:n}) \left(\prod_{p=0}^{n-1} G_p(x_p) \right) \eta_0(dx_0) \prod_{p=1}^n M^{(\infty)}(x_{p-1}, dx_p).$$

In practice one usually must approximate the true dynamics $M^{(\infty)}(x, dy)$ of the underlying diffusion with a simpler transition $M^{(\ell)}(x, dy)$, based on some Euler type scheme using a discretisation parameter $h_\ell = 2^{-\ell}$ for $\ell \geq 0$; see [25]. The scheme allows for a coupling of the diffusions $(X_t^{(\ell)}, X_t^{(\ell-1)})_{t \geq 0}$ running at discretisation levels ℓ and $\ell - 1$ (based on using the same Brownian path W_t), such that for some $\beta \in \{1, 2\}$, we have

$$(8) \quad \mathbb{E}_{(x,y)}[|X_1^{(\ell)} - X_1^{(\ell-1)}|^2] \leq M(|x - y|^2 + h_\ell^\beta),$$

where $M < \infty$ does not depend on $\ell \geq 1$. In particular, if the diffusion coefficient $b(X_t)$ in 1 is constant or if a Milstein scheme can be applied otherwise, then $\beta = 2$; otherwise $\beta = 1$; see Proposition D.1 of [21].

3.2. Variance bound. Assume we are in the above HMM diffusion setting, and that the coupling of Assumption 3 holds, with symbols (F, C) equal to $(\ell, \ell - 1)$ for $\ell \geq 1$, and $G_p^{(\ell)} = G_p^{(\ell-1)} := G_p$ for $p = 0:n$. Running Algorithm 2, we recall that $\Delta_\ell(\varphi)$, defined in 6, satisfies, by Proposition 4,

$$\mathbb{E}[\Delta_\ell(\varphi)] = \gamma_n^{(\ell)}(G_n \varphi) - \gamma_n^{(\ell-1)}(G_n \varphi),$$

regardless of the number $N \geq 1$ of particles.

A (measurable) function $\varphi : \mathbf{X} \rightarrow \mathbb{R}$ is Lipschitz, denoted $\varphi \in \text{Lip}(\mathbf{X})$, if for some $C' < \infty$, $|\varphi(x) - \varphi(y)| \leq C'|x - y|$ for all $x, y \in \mathbf{X}$.

Condition (A). The following conditions hold for the model (M_n, G_n) :

- (A1) (i) $\|G_n\| < \infty$ for each $n \geq 0$.
(ii) $G_n \in \text{Lip}(\mathbf{X})$ for each $n \geq 0$.
(iii) $\inf_{x \in \mathbf{X}} G_n(x) > 0$ for each $n \geq 0$.
- (A2) For every $n \geq 1$, $\varphi \in \text{Lip}(\mathbf{X}) \cap \mathcal{B}_b(\mathbf{X})$ there exist a $C' < \infty$ such that for $s \in \{F, C\}$, we have for every $(x, y) \in \mathbf{X} \times \mathbf{X}$ that $|M_n^s(\varphi)(x) - M_n^s(\varphi)(y)| \leq C'|x - y|$.

In the following results for $\Delta_\ell(\varphi)$, the constant $M < \infty$ may change from line-to-line. It will not depend upon N or ℓ (or θ), but may depend on the time-horizon n or the function φ . $\mathbb{E}$ denotes expectation w.r.t. the law associated to the Δ PF started at (x, x) , with $x \in \mathbf{X}$. Below we only consider multinomial resampling in the Δ PF for simplicity, though Theorem 9 and Corollary 10 can be proved also assuming other resampling schemes.

Theorem 9. Assume (A1-2). Then for any $\varphi \in \mathcal{B}_b(\mathbf{X}^{n+1}) \cap \text{Lip}(\mathbf{X}^{n+1})$, there exists a $M < \infty$ such that

$$\mathbb{E} \left[\left(\Delta_\ell(\varphi) - \mathbb{E}[\Delta_\ell(\varphi)] \right)^2 \right] \leq \frac{M h_\ell^{2 \wedge \beta}}{N}, \quad \text{with } \beta \text{ as in 8.}$$

Corollary 10. Assume (A1-2). Then for any $\varphi \in \mathcal{B}_b(\mathbf{X}^{n+1}) \cap \text{Lip}(\mathbf{X}^{n+1})$, there exists a $M < \infty$ such that

$$\mathbb{E} \left[\left(\Delta_\ell(\varphi) \right)^2 \right] \leq M \left(\frac{h_\ell^{2 \wedge \beta}}{N} + h_\ell^2 \right), \quad \text{with } \beta \text{ as in 8.}$$

The proofs are given in Appendix A.

Based on Corollary 10, Recommendation 1 of Section 5 suggests allocations for $\mathbf{p}$ and N_ℓ in the Δ PF (Algorithm 2) to optimally use resources and minimise variance 5.

4. UNBIASED JOINT INFERENCE FOR HIDDEN MARKOV MODEL DIFFUSIONS

We are interested in unbiased inference for the Bayesian model posterior

$$\pi^{(\infty)}(d\theta, dx_{0:n}) \propto \text{pr}(d\theta) G_n^{(\theta)}(x_n) \gamma_n^{(\theta, \infty)}(dx_{0:n}),$$

where $\text{pr}(d\theta) = \text{pr}(\theta)d\theta$ is the prior on the model parameters, and

$$\gamma_n^{(\theta, \infty)}(dx_{0:n}) = \left(\prod_{t=0}^{n-1} G_t^{(\theta)}(x_t) \right) \eta_0^{(\theta)}(dx_0) \prod_{t=1}^n M_t^{(\theta, \infty)}(x_{t-1}, dx_t).$$

Here, $M_t^{(\theta, \infty)}$ corresponds to the transition density of the diffusion model of interest. The dependence of the HMM on θ is made explicit in this section. As in Section 3, we assume the transition densities $M_t^{(\theta, \infty)}$ cannot be simulated, but that there are increasingly refined discretisations $M_t^{(\theta, \ell)}$ approximating $M_t^{(\theta, \infty)}$ in the sense of 4 (with $E_{0:n} := \mathbf{X}^{n+1}$).

4.1. Randomised MLMC IS type estimator based on coarse-model PMMH. We now consider Algorithm 4 for joint inference w.r.t. the above Bayesian posterior. Algorithm 4 uses the following ingredients:

- (i) $\tilde{M}_{0:n}^{(\theta, \ell)}$ satisfying Assumption 3i with $M_{0:n}^F = M_{0:n}^{(\theta, \ell)}$, and $M_{0:n}^C = M_{0:n}^{(\theta, \ell-1)}$.
- (ii) $\tilde{G}_{0:n}^{(\theta)}$ defined as in Assumption 3ii, with $G_{0:n}^F = G_{0:n}^C = G_{0:n}^{(\theta)}$.
- (iii) Metropolis-Hastings proposal distribution $q(\cdot | \theta)$ for parameters.
- (iv) Algorithm constant $\epsilon \geq 0$ (e.g. $\epsilon = 10^{-10}$; see Remark 11i below).

Algorithm 4 Randomised multilevel importance sampling type estimator.

Input: $\left((M_{0:n}^{(\theta,\ell)}, G_{0:n}^{(\theta)})\right)_{(\ell,\theta) \in (\{0\} \times \mathbb{N}) \times \mathbb{T}}$, prior $\text{pr}(\text{d}\theta)$, N the number of particles, m_{iter} the number of iterations, $q(\cdot \mid \theta)$ a proposal density for Metropolis-Hastings, $\epsilon \geq 0$, probability $\mathbf{p} = (p_\ell)_{\ell \in \mathbb{N}}$ and $(\Theta_0, V_0^{(1:N)}, \mathbf{X}_0^{(1:N)})$ such that $\sum_{i=1}^N V_0^{(i)} > 0$.

(P1) For $k = 1:m_{\text{iter}}$, iterate:

- (i) Propose $\hat{\Theta}_k \sim q(\cdot \mid \Theta_{k-1})$.
- (ii) Run Algorithm 1 with $(M_{0:n}^{(\hat{\Theta}_k,0)}, G_{0:n}^{(\hat{\Theta}_k)}, N)$ and call the output $(\hat{V}_k^{(1:N)}, \hat{\mathbf{X}}_k^{(1:N)})$.
- (iii) With probability

$$\min \left\{ 1, \frac{\text{pr}(\hat{\Theta}_k)q(\Theta_{k-1} \mid \hat{\Theta}_k)(\sum_{i=1}^N \hat{V}_k^{(i)} + \epsilon)}{\text{pr}(\Theta_{k-1})q(\hat{\Theta}_k \mid \Theta_{k-1})(\sum_{j=1}^N V_{k-1}^{(j)} + \epsilon)} \right\},$$

accept and set $(\Theta_k, V_k^{(1:N)}, \mathbf{X}_k^{(1:N)}) \leftarrow (\hat{\Theta}_k, \hat{V}_k^{(1:N)}, \hat{\mathbf{X}}_k^{(1:N)})$; otherwise set $(\Theta_k, V_k^{(1:N)}, \mathbf{X}_k^{(1:N)}) \leftarrow (\Theta_{k-1}, V_{k-1}^{(1:N)}, \mathbf{X}_{k-1}^{(1:N)})$.

(P2) For every $k \in \{1:m_{\text{iter}}\}$, independently, conditional on $(\Theta_k, V_k^{(1:N)}, \mathbf{X}_k^{(1:N)})$:

- (i) Set $\mathbf{X}_{k,0}^{(1:N)} := \mathbf{X}_k^{(1:N)}$, and set $W_{k,0}^{(i)} := V_k^{(i)} / (\sum_{j=1}^N V_k^{(j)} + \epsilon)$.
- (ii) Sample $L_k \sim \mathbf{p}$ independently from the other random variables.
- (iii) Run the Δ PF (Algorithm 2) with $(\check{M}_{0:n}^{(\Theta_k, L_k)}, \check{G}_{0:n}^{(\Theta_k)}, N)$, and call the output $(V_{k,L_k}^{(1:2N)}, \mathbf{X}_{k,L_k}^{(1:2N)})$. Set $W_{k,L_k}^{(i)} := V_{k,L_k}^{(i)} / [p_{L_k}(\sum_{j=1}^N V_k^{(j)} + \epsilon)]$.

Output:

$$E_{m_{\text{iter}}, N, \mathbf{p}}(f) := \frac{\sum_{k=1}^{m_{\text{iter}}} [\sum_{i=1}^N W_{k,0}^{(i)} f(\Theta_k, \mathbf{X}_{k,0}^{(i)}) + \sum_{i=1}^{2N} W_{k,L_k}^{(i)} f(\Theta_k, \mathbf{X}_{k,L_k}^{(i)})]}{\sum_{k=1}^{m_{\text{iter}}} [\sum_{i=1}^N W_{k,0}^{(i)} + \sum_{i=1}^{2N} W_{k,L_k}^{(i)}]}.$$

(v) Number of MCMC iterations $m_{\text{iter}} \in \mathbb{N}$ and number of particles $N \in \mathbb{N}$.

(vi) Probability mass $\mathbf{p} = (p_\ell)_{\ell \in \mathbb{N}}$ on $\mathbb{N}$ with $p_\ell > 0$ for all $\ell \in \mathbb{N}$.

Remark 11. Before stating consistency and central limit theorems, we briefly discuss various aspects of this approach, which are appealing from a practical perspective, and we also mention certain algorithmic modifications which could be further considered.

- (i) The first phase (P1) of Algorithm 4 implements a PMMH type algorithm [2]. If $\epsilon = 0$, this is exactly PMMH targeting the model $\pi^{(0)}(\text{d}\theta, \text{d}x_{0:n}) \propto \text{pr}(\text{d}\theta)G_n^{(\theta)}(x_n)\gamma_n^{(\theta,0)}(\text{d}x_{0:n})$. It is generally safer to choose $\epsilon > 0$ [33], which ensures that the IS type correction in phase (P2) will yield consistent inference for the ideal model

$$\pi^{(\infty)}(\text{d}\theta, \text{d}x_{0:n}) \propto \text{pr}(\text{d}\theta)G_n^{(\theta)}(x_n)\gamma_n^{(\theta,\infty)}(x_{0:n})$$

(Theorem 12). Setting $\epsilon > 0$ may be helpful otherwise in terms of improved mixing, as the PMMH will target marginally an averaged probability between a ‘flat’ prior and a ‘multimodal’ $\ell = 0$ marginal posterior.

- (ii) It is only necessary to implement PMMH for the coarsest level. This is typically relatively cheap, and therefore allows for a relatively long MCMC run. Consequently, relative cost of burn-in is small, and if the proposal q is adapted (see [1]), it has time to converge.
- (iii) The (potentially costly) r Δ PFs are applied independently for each Θ_k , which allows for efficient parallelisation.

- (iv) We suggest that the number of particles N , here referred to as ' N_0 ', used in the PMMH be chosen based on [10, 30], while the number of particles ' N_ℓ ' (and p_ℓ) can be optimised for each level ℓ based on Recommendation 1 of Section 5, or kept constant. One can also afford to increase the number of particles when a 'jump chain' representation is used (see the following remark).
- (v) The $r\Delta PF$ corrections may be calculated only once for each accepted state [33]. That is, suppose $(\tilde{\Theta}_k, \tilde{V}_k^{(1:N)}, \tilde{\mathbf{X}}_k^{(1:N)})_{k=1}^{m_{\text{iter}}^{\text{jump}}}$ are the accepted states, $(D_k)_{k=1}^{m_{\text{iter}}^{\text{jump}}}$ are the corresponding holding times, and $(\tilde{V}_{k,L_k}^{(1:2N_{L_k})}, \tilde{\mathbf{X}}_{k,L_k}^{(1:2N_{L_k})})_{k=1}^{m_{\text{iter}}^{\text{jump}}}$ are corresponding ΔPF outputs, then the estimator is formed as in Algorithm 4 using $(\tilde{\Theta}_k, \tilde{V}_k^{(1:N)}, \tilde{\mathbf{X}}_k^{(1:N)})$, and accounting for the holding times in the weights defined as $W_k^{(i)} := D_k \tilde{V}_k^{(i)} / (\sum_{j=1}^N \tilde{V}_k^{(j)} + \epsilon)$ and $W_{k,L_k}^{(i)} := \tilde{V}_{k,L_k}^{(i)} / [p_{L_k} (\sum_{j=1}^N \tilde{V}_k^{(j)} + \epsilon)]$.
- (vi) In case the Markov chain in (P1) phase is slow mixing, (further) thinning may be applied (to the jump chain) before the (P2) phase.
- (vii) In practice, Algorithm 4 may be implemented in an on-line fashion w.r.t. the number of iterations m_{iter} , and by progressively refining the estimator $E_{m_{\text{iter}}, N, \mathbf{p}}(f)$. The $r\Delta PF$ corrections may be calculated in parallel with the Markov chain.
- (viii) In Algorithm 4, the $r\Delta PF$ s depend only on Θ_k . They could depend also on $V_k^{(i)}$ and $\mathbf{X}_k^{(i)}$, but it is not clear how such dependence could be used in practice to achieve better performance. Likewise, the 'zeroth level' estimate in Algorithm 4 is based solely on particles in (P1), but it could also be based on (additional) new particle filter output.
- (ix) In order to save memory, it is possible also to 'subsample' only one trajectory $\mathbf{X}_k^*$ from $\mathbf{X}_k^{(1:N)}$, such that $\mathbb{P}[\mathbf{X}_k^* = \mathbf{X}_k^{(i)}] \propto V_k^{(i)}$, and set $W_{k,0}^* := \sum_{i=1}^N W_{k,0}^{(i)}$, and similarly in Algorithm 2 find $\check{\mathbf{X}}^*$ such that $\mathbb{P}[\check{\mathbf{X}}^* = \check{\mathbf{X}}^{(i)}] \propto \tilde{V}^{(i)}$, setting $\mathbf{X}_{k,L_k}^{*(1:2)} := \check{\mathbf{X}}^*$, and defining from the usual output of Algorithm 2, $W_{k,L_k}^{*(1)} := \sum_{i=1}^N W_{k,L_k}^{(i)}$ and $W_{k,L_k}^{*(2)} := \sum_{i=N+1}^{2N} W_{k,L_k}^{(i)}$. The subsampling output estimator then takes the form,

$$E_{m_{\text{iter}}, N, \mathbf{p}}^{\text{subsample}}(f) := \frac{\sum_{k=1}^{m_{\text{iter}}} [W_{k,0}^* f(\Theta_k, \mathbf{X}_k^*) + \sum_{i=1}^2 W_{k,L_k}^{*(i)} f(\Theta_k, \mathbf{X}_{k,L_k}^{*(i)})]}{\sum_{k=1}^{m_{\text{iter}}} [W_{k,0}^* + \sum_{i=1}^2 W_{k,L_k}^{*(i)}]}.$$

Note, however, that the asymptotic variance of this estimator is higher, because

$E_{m_{\text{iter}}, N, \mathbf{p}}(f)$ may be viewed as a Rao-Blackwellised version of $E_{m_{\text{iter}}, N, \mathbf{p}}^{\text{subsample}}(f)$.

4.2. Consistency and central limit theorem.

Theorem 12. Assume that the algorithm constant $\epsilon \geq 0$ is chosen positive, and that the Markov chain $(\Theta_k, X_k^{(1:N)}, V_k^{(1:N)})_{k \geq 1}$ is ψ -irreducible, and that $\pi^{(0)}(f)$ and $\pi^{(\infty)}(f)$ are finite. For each $\theta \in \mathbb{T}$, suppose Assumption 6 holds for $g \equiv 1$ and $g = f^{(\theta)} := f(\theta, \cdot)$, with $M_{0:n}^{(\ell)} := M_{0:n}^{(\theta, \ell)}$ and $G_{0:n}^{(\ell)} := G_{0:n}^{(\theta)}$. Assume

$$\int \text{pr}(\theta) (\sqrt{s_1(\theta)} + \sqrt{s_f(\theta)}) d\theta < \infty.$$

Then, the estimator of Algorithm 4 is strongly consistent:

$$E_{m_{\text{iter}}, N, \mathbf{p}}(f) \xrightarrow{m_{\text{iter}} \rightarrow \infty} \int \pi^{(\infty)}(d\theta, dx_{0:n}) f(\theta, x_{0:n}) \quad (a.s.)$$

Remark 13. Regarding Theorem 12, whose proof is given in Appendix B:

- (i) If all potentials G_t are strictly positive, the algorithm constant ϵ may be taken to be zero. However, if $\epsilon = 0$ and Algorithm 1 with $(M_{0:n}^{(\hat{\Theta}_k, 0)}, G_{0:n}^{(\hat{\Theta}_k)}, N)$ can produce an estimate with $\sum_{i=1}^N V^{(i)} = 0$ with positive probability, the consistency may be lost [33].

Proposition 14. Suppose that the conditions of Theorem 12 hold. Suppose additionally that $\pi^{(\infty)}(f^2) < \infty$ and that the base chain $(\Theta_k, V_k^{(1:N)}, \mathbf{X}_k^{(1:N)})_{k \geq 1}$ is aperiodic, with transition probability denoted by P . Then,

$$\sqrt{m_{\text{iter}}} [E_{m_{\text{iter}}, N, \mathbf{p}}(f) - \pi^{(\infty)}(f)] \xrightarrow{m_{\text{iter}} \rightarrow \infty} \mathcal{N}(0, \sigma^2), \quad \text{in distribution,}$$

whenever the asymptotic variance

$$(9) \quad \sigma^2 = \frac{\text{var}(P, \mu_{\bar{f}}) + \Pi(\sigma_{\xi}^2)}{c^2},$$

is finite. Here, $\bar{f} := f - \pi^{(\infty)}(f)$, $c > 0$ is a constant (equal to $\Pi(\mu_1)$), and

$$\begin{aligned} \sigma_{\xi}^2(\theta, v^{(1:N)}, \mathbf{x}^{(1:N)}) &:= \text{var}(\xi_k(\bar{f}) \mid \Theta_k = \theta, V_k^{(1:N)} = v^{(1:N)}, \mathbf{X}_k^{(1:N)} = \mathbf{x}^{(1:N)}) \\ &= \frac{s_{\bar{f}(\theta)}(\theta) - (\gamma_n^{(\theta, \infty)}(G_n \bar{f}(\theta)) - \gamma_n^{(\theta, 0)}(G_n \bar{f}(\theta)))^2}{(\sum_{i=1}^N v^{(i)} + \epsilon)^2}. \end{aligned}$$

Remark 15. Proposition 14 follows from Theorem 7 [33]. We suggest that $N = N_0$ for (P1) be chosen based on [10, 30] to minimise $\text{var}(P, \mu_{\bar{f}})$, and that (p_ℓ) and $N = N_\ell$ in (P2) for the rDPF be chosen as in Recommendation 1 of Section 5, to minimise σ_{ξ}^2 , subject to cost constraints, in order to jointly minimise σ^2 . However, the question of the optimal choice for N_0 in the IS context is not yet settled.

Remark 16. Regarding ψ -irreducibility and aperiodicity in Theorem 12 and Proposition 14, these are inherited by the coarse PMMH chain [2] if the corresponding idealised marginal Metropolis-Hastings for the coarse model has these properties.

5. ASYMPTOTIC EFFICIENCY AND RANDOMISED MULTILEVEL CONSIDERATIONS

We summarise the results of this section by suggesting the following safe allocations for probability $\mathbf{p} = (p_\ell)_{\ell \in \mathbb{N}}$ and number $N = N_\ell$ of particles at level $\ell \geq 1$ in the DPF (Algorithm 2) used in Algorithm 3 and Algorithm 4, and Proposition 8, with β given in 8 in the HMM diffusion context of Section 3, or, indeed, with β given in the abstract framework under Assumption 20 given later. See also Figure 1 for the recommended allocations.

Recommendation 1. With strong error convergence rate β given in 8, we suggest the following for $\mathbf{p} = (p_\ell)_{\ell \in \mathbb{N}}$ and $N_\ell \in \mathbb{N}$ in DPF (Algorithm 2):

- $(\beta = 1)$ (e.g. Euler scheme). Choose $p_\ell = (\frac{1}{2})^\ell$ and $N_\ell \propto 1$ constant.
- $(\beta = 2)$ (e.g. Milstein scheme). Choose $p_\ell \propto 2^{-1.5\ell} \approx (\frac{1}{3})^\ell$ and $N_\ell \propto 1$ constant.

The suggestions are based on Corollary 10 of Section 3, and Proposition 22 ($\beta = 2$) and 28 ($\beta = 1$) given below (with weak convergence rate $\alpha = 1$; see Figure 1 for general α). In the Euler case, although the theory below gives the same computational complexity order by choosing any $\rho \in [0, 1]$ and setting $p_\ell \propto 2^{-(1+\rho)\ell}$ and $N_\ell \propto 2^{\rho\ell}$, the experiment in Section

6 gave a better result using simply $\rho = 0$, corresponding to no scaling. However, this may depend on the implementation.

5.1. Efficiency framework. The asymptotic efficiency of simulation-based estimators has been considered theoretically in [18]; see [17] in the dMLMC context. The developments of this section follow [28] for rMLMC (originally in the i.i.d. setting without observations), while also giving some extensions (also applicable to that setting). We will see that the basic rMLMC results carry over to our setting involving MCMC and randomised estimators based on PF outputs, but also discover a novelty in the common Euler case ($\beta = 1$ in Figure 1). Proofs are given in Appendix C.

We are interested in modeling the computational costs involved in running Algorithm 4; the algorithm of Proposition 8 is recovered with $\mathsf{T} = \{\theta\}$. Let τ_{Θ_k, L_k} represent the combined *cost at iteration k* of the base Markov chain and weight calculation in Algorithm 4, so that the *total cost* $\mathcal{C}(m)$ of Algorithm 4 with m iterations is

$$\mathcal{C}(m) := \sum_{k=1}^m \tau_{\Theta_k, L_k}.$$

The following assumption seems natural in our setting.

Assumption 17. For $\Theta_k \in \mathsf{T}$, a family $\{\tau_{\Theta_k, \ell}\}_{k, \ell \geq 1}$ consists of positive-valued random variables that are independent of $\{L_k\}_{k \geq 1}$, where $L_k \sim \mathbf{p}$ i.i.d., and that are conditionally independent given $\{\Theta_k\}_{k \geq 1}$, such that $\tau_{\Theta_k, \ell}$ depends only on $\Theta_k \in \mathsf{T}$ and $\ell \in \mathbb{N}$.

Under a budget constraint $\kappa > 0$, the *realised length of the chain* is $\mathcal{L}(\kappa)$ iterations, where

$$\mathcal{L}(\kappa) := \max\{m \geq 1 : \mathcal{C}(m) \leq \kappa\}.$$

Under a budget constraint, the CLT of Proposition 14 takes the following altered form, where here $\Pi_m(d\theta)$ denotes the θ -marginal of the invariant probability measure (given as 27 in Appendix B) of the base Markov chain (equal to the θ -marginal posterior of the $\ell = 0$ model).

Proposition 18. *If the assumptions of Proposition 14 hold with $\sigma^2 < \infty$, and if $\mathbb{E}[\boldsymbol{\tau}] := \mathbb{E}_{\Pi_m \otimes \mathbf{p}}[\boldsymbol{\tau}] < \infty$ with $\boldsymbol{\tau}(\theta, \ell) := \mathbb{E}[\tau_{\Theta_k, L_k} | \Theta_k = \theta, L_k = \ell]$, then*

$$(10) \quad \sqrt{\kappa} [E_{\mathcal{L}(\kappa), N, \mathbf{p}}(f) - \pi^{(\infty)}(f)] \xrightarrow{\kappa \rightarrow \infty} \mathcal{N}(0, \mathbb{E}[\boldsymbol{\tau}] \sigma^2), \quad \text{in distribution.}$$

Remark 19. The quantity $\mathbb{E}[\boldsymbol{\tau}] \sigma^2$ is called the ‘inverse relative efficiency’ by [18], and is considered a more accurate quantity than the asymptotic variance (σ^2 here) for comparison of Monte Carlo algorithms run on the same computer, as it takes into account also the average computational time.

In the following we consider (possibly) variance reduced (if $\rho > 0$) versions of $\Delta_\ell(g)$ of Assumption 6, denoted Δ_ℓ , where $g = f^{(\theta)}$, based on running the Δ PF (Algorithm 2) with parameters θ, ℓ fixed. The constant $C < \infty$ may change line-to-line, but does not depend on N, ℓ , or θ , but may depend on the time-horizon n and function f .

Assumption 20. Assumption 17 holds, and constants $2\alpha \geq \beta > 0$, $\gamma > 0$, and $\rho \geq 0$ are such that the following hold:

- (i) (Mean cost) $\mathbb{E}[\tau_{\theta, \ell}] \leq C 2^{\gamma \ell (1 + \rho)}$
- (ii) (Strong order) $\mathbb{E}[\Delta_\ell^2] \leq C 2^{-\ell(\beta + \rho)} + C 2^{-2\alpha \ell}$
- (iii) (Weak order) $|\mathbb{E} \Delta_\ell| \leq C 2^{-\alpha \ell}$

Remark 21. Regarding Assumption 20:

FIGURE 1. Recommendations for number of particles $N_\ell \propto 2^{\rho\ell}$ and probability $p_\ell \propto 2^{-r\ell}$. Here, $\gamma = 1$ always, and $\rho \in [0, 2\alpha - 1]$ when $\beta = 1$ provides a line of choices with the same *order* of computational complexity. In our particular experiment in Section 6, however, the simple choice $\rho = 0$, corresponding to no scaling in the particles, will turn out to be optimal.

- (i) We only assume bounded mean cost in Assumption **i**, rather than the almost sure cost bound commonly used. This generalisation allows for the setting where occasional algorithmic runs may take a long time.
- (ii) In the original MLMC setting, the cost scaling γ in Assumption **i** is taken to be $\gamma = 1$ [17, 28]. However, in settings involving uncertainty quantification, and where the forward solver may involve non-sparse matrix inversions, often $\gamma \geq 1$ [7, 21, 23].
- (iii) We assume in Assumption **i** that the mean cost to form Δ_ℓ is bounded by the γ -scaled product of the number of samples or particles N_ℓ times the number of Euler time steps $2^\ell + 2^{\ell-1}$ together with the $O(N_\ell)$ -resampling cost, where there are $N_\ell \propto 2^{\rho\ell}$ particles at level ℓ . Here, we recall that the stratified, systematic, and residual resampling algorithms have $O(N_\ell)$ cost, as does an improved implementation of multinomial resampling; see [6, 9].
- (iv) With $\rho = 0$, by Jensen's inequality one sees why $\alpha \geq \beta/2$ can be assumed, and that Assumption **ii** becomes $\mathbb{E}\Delta_\ell^2 \leq C2^{-\ell\beta}$.
- (v) $\rho \geq 0$ in Assumption **i** and **ii** corresponds to using an average of $N_\ell := \lceil 2^{\rho\ell} \rceil$ i.i.d samples of $\Delta_\ell^{(1)}$, i.e. $\Delta_\ell = \frac{1}{N_\ell} \sum_{i=1}^{N_\ell} \Delta_\ell^{(i)}$, or, of more present interest to us, to increasing the number of particles used in a PF by a factor of N_ℓ instead of the default lower number. The former leads to $\mathbb{E}\Delta_\ell^2 = \frac{1}{N_\ell} \text{var}(\Delta_\ell^{(1)}) + \mathbb{E}[\Delta_\ell^{(1)}]^2$, justifying Assumption **ii**, as does Corollary 10, with $\beta \in \{1, 2\}$ and $\alpha = 1$, for the Δ PF (Algorithm 2) in the HMM diffusion context (Section 3).

Proposition 22. *Suppose Assumption 20 and the assumptions of Proposition 14 hold, with $\text{var}(P, \mu_{\bar{f}}) < \infty$. If $p_\ell \propto 2^{-r\ell}$ for some $r \in (\gamma(1 + \rho), \min(\beta + \rho, 2\alpha))$, then 10 holds, i.e.*

$$\sqrt{\kappa} [E_{\mathcal{L}(\kappa), N, \mathbf{p}}(f) - \pi^{(\infty)}(f)] \xrightarrow{\kappa \rightarrow \infty} \mathcal{N}(0, \mathbb{E}[\tau]\sigma^2), \quad \text{in distribution.}$$

Remark 23. Regarding Proposition 22, in the common case $\gamma = 1$ for simplicity:

- (i) If $\beta > 1$ ('canonical convergence regime') and $\rho = 0$, then a choice for $r \in (1, \beta)$ exists. See also Theorem 4 of [28] for a discussion of the theoretically optimal $\mathbf{p}$.
- (ii) If $\beta \leq 1$ ('subcanonical convergence regime'), then $\beta + \rho \leq 1 + \rho$ and so no choice for r exists.

5.2. Subcanonical convergence. When $\beta > 1$, within the framework above we have seen that a canonical convergence rate holds (Proposition 22) because $\mathbb{E}[\tau] < \infty$ and $\sigma^2 < \infty$. When $\beta \leq 1$, this is no longer the case, and one must choose between a finite asymptotic variance and infinite expected cost, or vice versa. Assuming the former, and that a CLT holds (Proposition 14), for $\epsilon > 0$ and $0 < \delta < 1$ the Chebyshev inequality implies that the number of iterations of Algorithm 4 so that

$$(11) \quad \mathbb{P}[|E_{m,N,\mathbf{p}}(f) - \pi^{(\infty)}(f)| \leq \epsilon] \geq 1 - \delta,$$

holds implies that m must be of the order $O(\epsilon^{-2})$. The question is then how to minimise the total cost $\mathcal{C}(m)$, or *computational complexity*, involved in obtaining the m samples. This will involve optimising for (p_ℓ) and N_ℓ to minimise $\mathcal{C}(m)$, while keeping the asymptotic variance finite.

Proposition 24. *Suppose that the assumptions of Proposition 14 hold with $\sigma^2 < \infty$, and Assumption 20 holds with $\mathbb{E}[\tau_{\Theta_{k_0}, L_{k_0}}] = \infty$ for some $k_0 \geq 1$. If*

$$\sum_{k \geq 1} \sup_{j \geq 1} \mathbb{P}[\tau_{\Theta_j, L_j} > a_k] < \infty,$$

with $a_k = O(k^{c_1}(\log_2 k)^{c_0})$ for some constants $c_0 > 0$ and $c_1 \geq 1$, then 11 can be obtained with computational complexity

$$O(\epsilon^{-2c_1} |\log_2 \epsilon|^{c_0}) \quad \text{as } \epsilon \rightarrow 0.$$

Remark 25. The above result shows that even for costs with unbounded tails, reasonable confidence intervals and complexity order may be possible. This may be the case for example when a rejection sampler or adaptive resampling mechanism is used within Algorithm 1 or 4, which may lead to large costs for some Θ_k , for example a cost with a geometric tail.

The next results are as in Proposition 4 and 5 of [28] in the standard rMLMC setting, and shows how one can choose $\mathbf{p}$, assuming an additional almost sure cost bound, so that $\sigma^2 < \infty$, with reasonable complexity.

Proposition 26. *Suppose that the assumptions of Proposition 14 hold with $\text{var}(P, \mu_{\bar{f}}) < \infty$, and that Assumption 20 holds with $\beta \leq 1$, where moreover $\tau_{\theta, \ell} \leq C2^{\gamma \ell(1+\rho)}$ almost surely, uniformly in $\Theta_k = \theta \in \mathbb{T}$. For all $q > 2$ and $\eta > 1$, the choice of probability*

$$p_\ell \propto 2^{-2b\ell} \ell [\log_2(\ell + 1)]^\eta,$$

where $b := \min((\beta + \rho)/2, \alpha)$, leads to $\sigma^2 < \infty$, and 11 can be obtained with computational complexity

$$O\left(\epsilon^{-\gamma \frac{(1+\rho)}{b}} |\log_2 \epsilon|^{q\gamma \frac{(1+\rho)}{2b}}\right) \quad \text{as } \epsilon \rightarrow 0.$$

Remark 27. Regarding Proposition 26, with $\gamma = 1$:

- (i) Under Assumption 20 with $\rho = 0$, the usual setup in MLMC before variance reduced estimators are used, the above proposition shows that finite variance and 11 can be obtained without increasing the number of particles at the higher levels, even in the subcanonical regime. We have in this case $b = \beta/2 \leq \alpha$ and complexity $O\left(\epsilon^{-\frac{2}{\beta}} |\log_2 \epsilon|^{\frac{q}{\beta}}\right)$. When $\beta = 1$ (borderline case), dMLMC gives complexity $O(\epsilon^{-2} |\log_2 \epsilon|^2)$ [17, 21], which is negligibly better (recall $q > 2$), but is biased inference.
- (ii) When $\alpha > \beta/2$, which is the usual case in the subcanonical regime ($\beta \leq 1$); see [25], a more efficient use of resources can be obtained by increasing the number of particles (see Proposition 28 below).

Proposition 28. *Suppose the assumptions of Proposition 26 hold, where moreover $\rho \geq 0$ may vary as a free parameter without changing the constant $C > 0$. Then, for all $q > 2$, $\eta > 1$ constants, the choice $\rho = 2\alpha - \beta$ and probability*

$$p_\ell \propto 2^{-2\alpha\ell} \ell [\log_2(\ell + 1)]^\eta,$$

leads to $\sigma^2 < \infty$, and 11 can be obtained with computational complexity

$$O\left(\epsilon^{\gamma[-2-\frac{(1-\beta)}{\alpha}]} |\log_2 \epsilon|^{\gamma[q+\frac{(1-\beta)}{2\alpha}]}\right) \quad \text{as } \epsilon \rightarrow 0.$$

6. NUMERICAL SIMULATIONS

Now the theoretical results relating to the method herein introduced will be demonstrated on three examples. We will consider one example in the canonical regime, and two in the sub-canonical. In the first two experiments, the likelihoods can be computed exactly, so that the ground truth $\pi^{(\infty)}(f)$ can be easily calculated to arbitrary precision. We run each example with 100 independent replications, and calculate the mean squared error (MSE) when the chain is at length m as

$$\text{MSE}(m) = \frac{1}{100} \sum_{i=1}^{100} |E_{m,N,\mathbf{p}}^{(i)}(f) - \pi^{(\infty)}(f)|^2,$$

which is depicted as the thick red line, average of the thin lines, in Figure 2 below. The error decays with the optimal rate of cost^{-1} and $\log(\text{cost})\text{cost}^{-1}$ in the canonical and sub-canonical cases, respectively, where cost is the realised cost of the run, $\mathcal{C}(m)$ from Section 5, measured in seconds, with m iterations of the Markov chain.

Three examples will be considered. First we consider two models where the exact marginal likelihood can be computed. This way a reliable ground truth can be computed with a long MCMC chain, providing a strong verification of the theoretical results. In Section 6.1 the Ornstein–Uhlenbeck (OU) process is considered, where Euler-Maruyama provides the canonical convergence regime. In Section 6.2 the Geometric Brownian motion is considered, where Euler-Maruyama provides sub-canonical convergence regime. In Section 6.3 we consider a more complicated $2d$ model which does not allow exact computation of the marginal likelihood.

It is of interest to compare our methodology to existing unbiased methods. The method we consider for comparison is PMMH using the exact method introduced by Fearnhead et al. [11]. In their work they provide a way to construct unbiased estimates without approximating the transition density. The key idea is to assign to each particle a random positive weight which is an unbiased estimator of the true weight. This method is referred to as the random weight particle filter. It has been later extended to continuous-time observations in [12]. The method of [11] is implemented when it is applicable (in particular, for the models of Sections 6.1 and 6.2) and the corresponding MSE is plotted in comparison to our method. That method is not amenable to the example of Section 6.3. In all the examples, for the sake of comparison with the standard approach, we also implement a finely-discretised PMMH. This shows the benefit of our alternative approach based on a coarsely-discretised PMMH with multilevel IS correction.

6.1. Ornstein–Uhlenbeck process. Consider the OU process

$$(12) \quad dZ_t = -aZ_t dt + b dW_t, \quad t \geq 0,$$

with initial condition $Z_0 = 0$, model parameter $\theta = (\theta_1, \theta_2) \sim N(0, \sigma^2 I)$, and $a := a_\theta = \exp(\theta_1)$ and $b := b_\theta = \exp(\theta_2)$. The process is discretely observed for $k = 1, \dots, n$,

$$(13) \quad Y_p = X_p + \xi_p,$$

where $\xi_p \sim N(0, \gamma^2)$ i.i.d. and recall that $X_p = Z_{p+1}$. Therefore,

$$G_p(x) = \exp\left(-\frac{1}{2\gamma^2}|x - y_p|^2\right).$$

The marginal likelihood is given by

$$\mathbb{P}[y_{1:n}|\theta] = \prod_{p=1}^n \mathbb{P}[y_p|y_{1:p-1}, \theta],$$

and each factor can be computed as the marginal of the joint on the prediction and current observation, i.e.

$$(14) \quad \mathbb{P}[y_p|y_{1:p-1}, \theta] = \int_{\mathbb{R}} \mathbb{P}[y_p|x_p, \theta] \mathbb{P}[x_p|y_{1:p-1}, \theta] dx_p.$$

In this example the ground truth can be computed exactly via the Kalman filter. In particular, the solution of 12 is given by

$$Z_1 = e^{-a}X_0 + W_1, \quad W_1 \sim \mathcal{N}\left(0, \frac{b^2}{2a}(1 - e^{-2a})\right).$$

The filter at time p is given by the following simple recursion

$$m_p = c_p \left(\frac{y_p}{\gamma^2} + \frac{\hat{m}_p}{\hat{c}_p} \right), \quad c_p = (\gamma^{-2} + \hat{c}_p^{-1})^{-1}, \quad \hat{m}_p = e^{-a}m_{p-1}, \quad \hat{c}_p = e^{-2a}c_{p-1} + \frac{b^2}{2a}(1 - e^{-2a}).$$

Additionally, the incremental marginal likelihoods 14 can be computed exactly

$$\mathbb{P}[y_p|y_{1:p-1}, \theta] = \sqrt{\frac{c_p}{2\pi\hat{c}_p\gamma^2}} \exp \left\{ -\frac{1}{2} \left[\frac{y_p^2}{\gamma^2} + \frac{\hat{m}_p^2}{\hat{c}_p} - c_p \left(\frac{y_p}{\gamma^2} + \frac{\hat{m}_p}{\hat{c}_p} \right)^2 \right] \right\}.$$

The parameters are chosen as $\gamma = 1$, $\sigma^2 = 0.1$, $n = 5$, and the data is generated with $\theta = (0, 0)^T$. Our aim is to compute $\mathbb{E}(\theta|y_{1:n})$ (or $\mathbb{E}[(a, b)^T|y_{1:n}]$, etc., but we will content ourselves with the former). This is done via a brute force random walk MCMC for $m = 10^8$ steps using the exact likelihood $\mathbb{P}[y_{1:n}|\theta]$ as above. The IACT is around 10, so this gives a healthy limit for MSE computations.

For the numerical experiment, we use Euler-Maruyama method at resolution $h_\ell = 2^{-\ell}$ to solve 12 as follows

$$(15) \quad Z_{p+1} = (1 - ah_\ell)Z_p + bB_{p+1}, \quad B_{p+1} \sim \mathcal{N}(0, h_\ell) \text{ i.i.d.}$$

for $p = 1, \dots, K_\ell = h_\ell^{-1}$. Levels ℓ and $\ell - 1$ are coupled in the simulation of Δ_ℓ by defining $B_{1:K_\ell/2}^C = B_{1:2:K_\ell-1}^F + B_{2:2:K_\ell}^F$. Algorithm 2 is then run using the standard bootstrap particle filter (Algorithm 1) with $N = 20$ particles and $O(N)$ -complexity multinomial resampling; see [6]. Theorem 9 provides a rate of $\beta = 2$ for Algorithm 2, because the diffusion coefficient is constant, which implies we are essentially running a Milstein scheme (see 8 and [25]). Recommendation 1 (or Proposition 22) of Section 5 suggests arbitrary precision can be obtained by Algorithm 4 with $p_\ell \propto 2^{-3\ell/2}$ and no scaling of particle numbers based on ℓ in this canonical $\beta = 2$ regime (with weak rate $\alpha = 1$). We choose a positive PMMH algorithm constant $\epsilon = 10^{-6}$ (see Remark 11i). We run Algorithm 4 for 10^4 steps, with 100 replications. For the finely-discretised PMMH experiment we run 10^4 steps, with 100 replications with a discretisation of $h_\ell = 2^{-5}$. The results are presented in Figure 2, where it is clear that the theory holds and the MSE decays according to $1/\text{cost}$. The variance of the run-times is very small over replications. The method of [11], within PMMH, also converges with the theoretically-predicted canonical rate, but with a slightly smaller constant. This is not unexpected. The important point is that both methods achieve the same canonical

FIGURE 2. The MSE of PMMH rMLMC IS (Algorithm 4) applied to the problem of parameter inference for the discretely observed OU process (left plot) and GBM process (middle plot with $\rho = 0$, right plot with $\rho = 1$). Squared error replications are given by the thin curves, while the thick red curves give the MSE over replications. The thick magenta curves show the MSE of PMMH using [11] (denoted by MSE2 in the legend). The blue curve, denoted by MSE3, depicts the MSE for the finely-discretised PMMH. The black curves representing cost^{-1} (left plot) and $\log(\text{cost})\text{cost}^{-1}$ (middle and right plots) are there to guide the eye.

rate, while our method is quite generally applicable, in particular, to a wide range of models inaccessible to methods of the type of [11]. Also with the finely discretised PMMH from Figure 2, i.e. the curve titled as MSE3, we see the bias kicking in at the end. We also see that for the same level of cost the MSE is higher than that of the other methodologies.

6.2. Geometric Brownian motion. We next consider the following stochastic differential equation

$$(16) \quad dZ_t = aZ_t dW_t,$$

with initial condition $Z_0 = 1$, and $a := a_\theta = \exp(\theta)$ with $\theta \sim \mathcal{N}(0, \sigma^2)$. This equation is analytically tractable as well, and the solution of the transformed equation $U = \log Z$ is given via Itô's formula by

$$dU_t = -\frac{a^2}{2}dt + a dW_t.$$

Defining $W_p \sim \mathcal{N}(0, 1)$ i.i.d., one has that

$$U_{p+1} = U_p + \frac{a^2}{2} + a W_p, \quad \text{with } U_0 = \log z_0 = 0,$$

and the solution of 16 can be obtained via exponentiation: $Z_p = e^{U_p}$. Moreover, noisy observations are introduced on the form , with $X_p = Z_{p+1}$,

$$Y_p = \log(X_p) + \xi_p,$$

where $\xi_p \sim \mathcal{N}(0, \gamma^2)$ i.i.d. as above. Therefore we have

$$(17) \quad G_p(x) = \exp\left(-\frac{1}{2\gamma^2}|\log(x) - y_p|^2\right).$$

Again $\mathbb{P}[y_{1:n}|\theta]$ can be computed analytically. The parameters $\gamma = 1$, $\sigma^2 = 0.1$, $n = 5$ are chosen the same as in the previous example and the true observations are generated again with $\theta = 0$.

In order to investigate the theoretical sub-canonical rate, we return to 16 and approximate this directly using Euler-Maruyama method 15, which introduces artificial approximation

error. This problem suffers from stability problems when $X < 0$, so we take $h_\ell = 2^{-6-\ell}$. Algorithm 1 is then used along with the selection functions (17). Here the diffusion coefficient is not constant, and Euler-Maruyama method yields a rate of $\beta = 1 = \alpha$, the borderline case, which is expected to give a logarithmic penalty. Based on Recommendation 1 (or Proposition 28) of Section 5, we consider scaling the particles as $2^{\rho\ell}$ with $\rho = 2\alpha - \beta = 1$ and $\rho = 0$, with $p_\ell \propto 2^{-2\ell}\ell \log(\ell)^2$ in both cases. Again we let $\epsilon = 10^{-6}$, and the standard bootstrap particle filter is used, with $N = 20 \times 2^{\rho\ell}$ particles. Algorithm 4 is run for 10^4 steps, with 100 replications. Again, for the finely-discretised PMMH, we run 10^4 steps, with 100 replications with a discretisation of $h_\ell = 2^{-4}$. For this sub-canonical case we impose an artificial upper bound $\ell \leq 10$, corresponding to an induced bias of $\approx 10^{-5}$. The results are presented in Figure 2, and they show good agreement with the theory, in terms of rate. On the other hand, the cost for $\rho = 0$ is apparently smaller than that of $\rho = 1$ by a factor of approximately 100. The method of [11] is not expected to suffer from a logarithmic penalty on the MSE convergence, i.e. it achieves canonical rate also in this example. This can be seen in Figure 2, in addition to a slightly better constant, as before. For the finely-discretised PMMH from Figure 2, with geometric Brownian motion, we again notice the effect of the bias arising from the discretisation, and the overall higher MSE.

6.3. 2d Non-reversible Langevin equation. We now consider a $2d$ example which is not amenable to approaches of the type [11]. Consider a target distribution of the type $\rho(z) \propto \exp(-\Phi(z)/a_2)$, $a_2 > 0$, and the following non-reversible Langevin equation

$$(18) \quad dZ_t = (A - I_2)\nabla\Phi(Z_t)dt + \sqrt{2a_2}dW_t, \quad t \geq 0,$$

and noisy observations $Y_p \sim \mathcal{N}(X_p, \gamma^2 I_2)$, with $\gamma = 1$, $n = 10$ observations where $A \in \mathbb{R}^{2 \times 2}$ is anti-symmetric and parameterised by $a_1 \geq 0$, I_2 is the 2-dimensional identity matrix, and $\Phi(z) = \frac{a_3}{2}(z_1^2 + z_2^2 - 1)^2$ is the ring potential, parameterised by $a_3 > 0$. The initial condition is specified as $Z_0 = [1, 1]^T$. It is easy to see that the right-hand side of the Fokker-Planck equation vanishes for the invariant distribution ρ given above, so the dynamics are well-behaved [27]. We let $a_i = \exp(\theta_i)$, for $i = 1, \dots, 3$, and the prior is given by $\theta \sim N(0, \varepsilon^2 I_3)$, $\varepsilon^2 = 0.1$. The resolution of the Euler-Maruyama scheme for this experiment is set as $2^{-\ell+1}$. This problem is no longer analytically soluble, so a high-resolution simulation is used with a large sample size as ground truth.

For our setup it follows similarly to that of the previous experiment for the OU process, where we are working in a canonical regime. Again we run Algorithm 4 with 10^4 steps, and 100 replications, and Algorithm 2 is run using the standard bootstrap particle filter (Algorithm 1) with $N = 20$ particles. We compare our results to the single level PMMH, with a similar setup where we specify its resolution as 2^{-4} . As before we choose a positive PMMH algorithm constant $\epsilon = 10^{-6}$. As we can see from Figure 3 the MSE of the proposed methodology in the paper decays at the rate of $1/\text{cost}$, which is as expected, which outperforms that of the single-level PMMH.

Remark 29. For multidimensional diffusions, it is well known that the exact methodology works only on specific diffusions, which require strong assumptions [3]. It is not so clear how the exact methodology can be applied to our non-reversible diffusion (due to the difficult drift term) [4]. An alternative to this is the work of Blanchet et al. [5], which does not rely on the assumption of drift term equal to the gradient of a suitable potential function, or on Lamperti's transformation for that matter. However, despite this, the major drawback is that the running time of their methodology, although finite with probability one, has infinite mean. As a result, the comparison would not be practical due to the cost of the experiment.

FIGURE 3. The MSE of PMMH rMLMC IS Algorithm 4 applied to the problem of parameter inference for the discretely observed non-reversible Langevin equation. Squared error replications are given by the thin curves, while the thick red curves give the MSE over replications. The thick blue curve shows the MSE over replications of the finely-discretised PMMH, which we denote as MSE2. The black curve represents cost^{-1} .

ACKNOWLEDGMENTS

JF, AJ, KL and MV have received support from the Academy of Finland (grants 274740, 312605 and 315619) and from the Institute for Mathematical Sciences, Singapore (2018 programme ‘Bayesian Computation for High-Dimensional Statistical Models’). NC and AJ have received support from KAUST baseline funding, JF and KL from The Alan Turing Institute, AJ from the Singapore Ministry of Education (R-155-000-161-112), and KL from the University of Manchester (School of Mathematics). This research made use of the Rocket High Performance Computing service at Newcastle University. We thank Paul Fearnhead, Santeri Karppinen, Anthony Lee and the anonymous referees for their many insightful remarks.

APPENDIX A. ANALYSIS OF THE DELTA PARTICLE FILTER

We now give our analysis that is required for the proofs of Theorem 9 and Corollary 10 of Section 3 regarding the Δ PF (Algorithm 2) for HMM diffusions. The structure of the appendix is as follows. In Section A.1 we introduce some more Feynman–Kac notations, following [8, 21], emphasising that here we consider standard HMMs that can be coupled. In Section A.2 we recall the Δ PF stated earlier. A general variance bound for quantities such as $\Delta_\ell(\varphi)$ is given in Section A.3. This is particularised to the HMM diffusion case in Section A.4, where we supply the proofs for the results of Section 3.

A.1. Models. Let $(\mathbf{X}, \mathcal{X})$ be a measurable space and $\{G_n\}_{n \geq 0}$ a sequence of non-negative, bounded and measurable functions such that $G_n : \mathbf{X} \rightarrow \mathbb{R}_+$. Let $\eta_0^F, \eta_0^C \in \mathcal{P}(\mathbf{X})$ and $\{M_n^F\}_{n \geq 1}, \{M_n^C\}_{n \geq 1}$ be two sequences of Markov kernels, i.e. $M_n^F : \mathbf{X} \rightarrow \mathcal{P}(\mathbf{X})$, $M_n^C : \mathbf{X} \rightarrow \mathcal{P}(\mathbf{X})$. Set $\mathbf{E}_n := \mathbf{X}^{n+1}$ for $n \geq 0$, and for $x_{0:n} \in \mathbf{E}_n$,

$$\mathbf{G}_n(x_{0:n}) = G_n(x_n)$$

and for $n \geq 1$, $s \in \{F, C\}$, $x_{0:n-1} \in \mathbf{E}_{n-1}$

$$\mathbf{M}_n^s(x_{0:n-1}, dx'_{0:n}) = \delta_{\{x_{0:n-1}\}}(dx'_{0:n-1})M_n^s(x'_{n-1}, dx'_n).$$

Define for $s \in \{F, C\}$, $\varphi \in \mathcal{B}_b(\mathbf{E}_n)$, $u_n \in \mathbf{E}_n$

$$\gamma_n^s(\varphi) = \int_{\mathbf{E}_0 \times \dots \times \mathbf{E}_n} \varphi(u_n) \left(\prod_{p=0}^{n-1} \mathbf{G}_p^s(u_p) \right) \eta_0^s(\mathrm{d}u_0) \prod_{p=1}^n \mathbf{M}_p^s(u_{p-1}, \mathrm{d}u_p)$$

and

$$\eta_n^s(\varphi) = \frac{\gamma_n^s(\varphi)}{\gamma_n^s(1)}.$$

Throughout this appendix, we assume Assumption (D), and that Assumption 3(i) holds, i.e. there exists $\check{\eta}_0 \in \mathcal{P}(\mathbf{X} \times \mathbf{X})$ such that for any $A \in \mathcal{X}$

$$\check{\eta}_0(A \times \mathbf{X}) = \eta_0^F(A) \quad \check{\eta}_0(\mathbf{X} \times A) = \eta_0^C(A)$$

and moreover for any $n \geq 1$ there exists Markov kernels $\{\check{M}_n\}$, $\check{M}_n : \mathbf{X} \times \mathbf{X} \rightarrow \mathcal{P}(\mathbf{X} \times \mathbf{X})$ such that for any $A \in \mathcal{X}$, $(x, x') \in \mathbf{X} \times \mathbf{X}$:

$$(19) \quad \check{M}_n(A \times \mathbf{X})(x, x') = M_n^F(A)(x) \quad \check{M}_n(\mathbf{X} \times A)(x, x') = M_n^C(A)(x').$$

A.2. Delta particle filter. Define $x_p = (x_p^F, x_p^C) \in \mathbf{X} \times \mathbf{X}$ and

$$\check{G}_p(x_p) = \frac{1}{2}(G_p(x_p^F) + G_p(x_p^C)),$$

as in Assumption 3(ii). Set, for $n \geq 0$, $x_{0:n} \in \mathbf{X}^{2(n+1)}$

$$\check{G}_n(x_{0:n}) = \check{G}_n(x_n)$$

and for $n \geq 1$, $x_{0:n-1} \in \mathbf{X}^{2n}$

$$\check{M}_n(x_{0:n-1}, \mathrm{d}x'_{0:n}) = \delta_{\{x_{0:n-1}\}}(\mathrm{d}x'_{0:n-1}) \check{M}_n(x'_{n-1}, \mathrm{d}x'_n), .$$

Note that coupling assumption (19) for $\check{M}_n$ can be equivalently formulated for $\check{M}_n$.

For $n \geq 0$, $\varphi \in \mathcal{B}_b(\mathbf{E}_n \times \mathbf{E}_n)$, $u_n \in \mathbf{E}_n \times \mathbf{E}_n$, we have

$$\check{\gamma}_n(\varphi) = \int_{\mathbf{E}_0^2 \times \dots \times \mathbf{E}_n^2} \varphi(u_n) \left(\prod_{p=0}^{n-1} \check{G}_p(u_p) \right) \check{\eta}_0(\mathrm{d}u_0) \prod_{p=1}^n \check{M}_p(u_{p-1}, \mathrm{d}u_p)$$

and

$$\check{\eta}_n(\varphi) = \frac{\check{\gamma}_n(\varphi)}{\check{\gamma}_n(1)}.$$

As noted in [22] it is simple to establish that for $\varphi \in \mathcal{B}_b(\mathbf{E}_n)$, if

$$(20) \quad \psi(x_{0:n}) = \check{G}_n(x_{0:n}) \left(\varphi(x_{0:n}^F) \prod_{p=0}^n \frac{\mathbf{G}_p(x_{0:p}^F)}{\check{G}_p(x_{0:p})} - \varphi(x_{0:n}^C) \prod_{p=0}^n \frac{\mathbf{G}_p(x_{0:p}^C)}{\check{G}_p(x_{0:p})} \right)$$

then

$$(21) \quad \check{\gamma}_n(\psi) = \check{\gamma}_n(1) \check{\eta}_n(\psi) = \gamma_n^F(G_n \varphi) - \gamma_n^C(G_n \varphi).$$

Note

$$\check{\gamma}_n(1) = \prod_{p=0}^{n-1} \check{\eta}_p(\check{G}_p).$$

In order to approximate $\check{\gamma}_n(\psi)$ one can run the following abstract version of Algorithm 2 (recall from Section 3 that we will only consider multinomial resampling). Define for $n \geq 1$, $\mu \in \mathcal{P}(\mathbf{E}_{n-1} \times \mathbf{E}_{n-1})$, $\varphi \in \mathcal{B}_b(\mathbf{E}_n \times \mathbf{E}_n)$

$$\check{\phi}_n(\mu)(\varphi) = \frac{\mu(\check{G}_{n-1} \check{M}_n(\varphi))}{\mu(\check{G}_{n-1})}.$$

The algorithm begins by generating $u_0^i \in \mathbf{E}_0 \times \mathbf{E}_0$, $i \in \{1, \dots, N\}$ with joint law

$$\prod_{i=1}^N \check{\eta}_0(du_0^i) = \prod_{i=1}^N \check{\eta}_0(du_0^i).$$

Defining

$$\check{\eta}_0^N(du_0) = \frac{1}{N} \sum_{i=1}^N \delta_{u_0^i}(du_0)$$

we then generate $u_1^i \in \mathbf{E}_1 \times \mathbf{E}_1$, $i \in \{1, \dots, N\}$ with joint law

$$\prod_{i=1}^N \check{\phi}_1(\check{\eta}_0^N)(du_1^i).$$

This proceeds recursively, so the joint law of the particles up to time n is

$$\left(\prod_{i=1}^N \check{\eta}_0(du_0^i) \right) \left(\prod_{p=1}^n \prod_{i=1}^N \check{\phi}_p(\check{\eta}_{p-1}^N)(du_p^i) \right).$$

Hence we have the estimate

$$\check{\gamma}_n^N(\psi) = \left(\prod_{p=0}^{n-1} \check{\eta}_p^N(\check{\mathcal{G}}_p) \right) \check{\eta}_n^N(\psi).$$

Remark 30. Note that $\check{\gamma}_n^N(\psi)$ corresponds to the quantity $\Delta_\ell(\varphi)$ in (6) from the Δ PF output (Algorithm 2).

A.3. General hidden Markov model case. Define for $p \geq 1$ the semigroup

$$\check{\mathcal{Q}}_p(x_{0:p-1}, dx'_{0:p}) = \check{\mathcal{G}}_{p-1}(x_{0:p-1}) \check{\mathcal{M}}_p(x_{0:p-1}, dx'_{0:p})$$

with the definition for $0 \leq p \leq n$, $\varphi \in \mathcal{B}_b(\mathbf{E}_n \times \mathbf{E}_n)$

$$\check{\mathcal{Q}}_{p,n}(\varphi)(u_p) = \int \varphi(u_n) \prod_{j=p+1}^n \check{\mathcal{Q}}_j(u_{j-1}, du_j)$$

if $p = n$ clearly $\check{\mathcal{Q}}_{n,n}$ is the identity operator. For any $0 \leq n$, $\varphi \in \mathcal{B}_b(\mathbf{E}_n \times \mathbf{E}_n)$ we set $\check{\mathcal{Q}}_{-1,n}(\varphi)(u_{-1}) = 0$.

Now following [8, Chapter 7] we have the following martingale (w.r.t. the natural filtration of the particle system), $\varphi \in \mathcal{B}_b(\mathbf{E}_n \times \mathbf{E}_n)$:

$$(22) \quad \check{\gamma}_n^N(\varphi) - \check{\gamma}_n(\varphi) = \sum_{p=0}^n \check{\gamma}_p^N(1) [\check{\eta}_p^N - \check{\phi}_p(\check{\eta}_{p-1}^N)] (\check{\mathcal{Q}}_{p,n}(\varphi))$$

with the convention that $\check{\phi}_p(\check{\eta}_{p-1}^N) = \check{\eta}_0$ if $p = 0$. The representation immediately establishes that

$$\mathbb{E}[\check{\gamma}_n^N(\varphi)] = \check{\gamma}_n(\varphi)$$

where the expectation is w.r.t. the law associated to the particle system. We will use the following convention that C' is a finite positive constant that does not depend upon n, N or any of the G_n , M_n^s ($s \in \{F, C\}$), $\check{M}_n$. The value of C' may change from line-to-line. Define for $0 \leq p \leq n < \infty$

$$\overline{G}_{p,n} = \prod_{q=p}^n \|G_q\|$$

with the convention that if $p = 0$ we write $\overline{G}_n$. We have the following result.

Proposition 31. *Suppose that $\|G_n\| < \infty$ for each $n \geq 0$. Then there exist a $C' < \infty$ such that for any $n \geq 0$, $\varphi \in \mathcal{B}_b(\mathbf{E}_n \times \mathbf{E}_n)$*

$$\mathbb{E} \left[\left(\check{\gamma}_n^N(\varphi) - \check{\gamma}_n(\varphi) \right)^2 \right] \leq \frac{C'}{N} \sum_{p=0}^n \bar{G}_{p-1}^2 \mathbb{E}[\check{\mathbf{Q}}_{p,n}(\varphi)(u_p^1)^2].$$

Proof. Set

$$\check{\mathbf{S}}_{p,n}^N(\varphi) = \check{\gamma}_p^N(1)[\check{\eta}_p^N - \check{\phi}_p(\check{\eta}_{p-1}^N)](\check{\mathbf{Q}}_{p,n}(\varphi))$$

By (22), one can apply the Burkholder-Gundy-Davis inequality to obtain

$$(23) \quad \mathbb{E} \left[\left(\check{\gamma}_n^N(\varphi) - \check{\gamma}_n(\varphi) \right)^2 \right] \leq C' \sum_{p=0}^n \mathbb{E}[\check{\mathbf{S}}_{p,n}^N(\varphi)^2].$$

Now, we have that

$$\mathbb{E}[\check{\mathbf{S}}_{p,n}^N(\varphi)^2] \leq \bar{G}_{p-1}^2 \mathbb{E}[(\check{\eta}_p^N - \check{\phi}_p(\check{\eta}_{p-1}^N)](\check{\mathbf{Q}}_{p,n}(\varphi))^2].$$

Application of the (conditional) Marcinkiewicz-Zygmund inequality yields

$$\mathbb{E}[\check{\mathbf{S}}_{p,n}^N(\varphi)^2] \leq \frac{C' \bar{G}_{p-1}^2}{N} \mathbb{E} \left[\left(\check{\mathbf{Q}}_{p,n}(\varphi)(u_p^1) - \check{\phi}_p(\check{\eta}_{p-1}^N)(\check{\mathbf{Q}}_{p,n}(\varphi)) \right)^2 \right].$$

After applying C_2 and Jensen inequalities, we then conclude by (23). $\square$

A.4. Diffusion case. We now consider the model of Section 3, where we recall that θ is omitted from the notation. A series of technical results are given and the proofs for Theorem 9 and Corollary 10 are given at the end of this section.

We recall that the joint probability density of the observations and the unobserved diffusion at the observation times is given by

$$\prod_{p=0}^n G_p(x_p) Q^{(\infty)}(x_{p-1}, x_p).$$

As the true dynamics can not be simulated, in practice we work with

$$\prod_{p=0}^n G_p(x_p) Q^{(\ell)}(x_{p-1}, x_p).$$

Recall an (Euler) approximation scheme with discretisation $h_\ell = 2^{-\ell}$, $\ell \geq 0$. In our context then, M_n^F corresponds $Q^{(\ell)}$ ($\ell \geq 1$) and M_n^C corresponds $Q^{(\ell-1)}$. The initial distribution η_0 is simply the (Euler) kernel started at some given x_0 . As noted earlier in Remark 5(i), a natural coupling of M_n^F and M_n^C (and hence of η_0) exists. As established in [21, eq. (32)] one has (uniformly in θ as Assumption (D) holds with θ independent constants) for $C' < \infty$

$$(24) \quad \sup_{\mathcal{A}} \sup_{x \in \mathbf{X}} |M_n^F(\varphi)(x) - M_n^C(\varphi)(x)| \leq C' h_\ell$$

where $\mathcal{A} = \{\varphi \in \mathcal{B}_b(\mathbf{X}) \cap \text{Lip}(\mathbf{X}) : \|\varphi\| \leq 1\}$. We also recall that (8) holds (recall Assumption (D) is assumed).

We will use $M < \infty$ to denote a constant that may change from line-to-line. It will not depend upon θ nor N , ℓ , but may depend on the time parameter or a function. The following result will be needed later on. The proof is given after the proof of Lemma 33 below.

Proposition 32. Assume (A1 (i)-(ii), 2). Then for any $n \geq 0$ and $\varphi \in \mathcal{B}_b(\mathbf{X}^{n+1}) \cap \text{Lip}(\mathbf{X}^{n+1})$ there exists a $M < \infty$ such that

$$|\gamma_n^F(G_n\varphi) - \gamma_n^C(G_n\varphi)| \leq Mh_\ell$$

We write expectations w.r.t. the time-inhomogeneous Markov chain associated to the sequence of kernels $(M_p^F)_{p \geq 1}$ (resp. $(M_p^C)_{p \geq 1}$) as $\mathbb{E}^F$, (resp. $\mathbb{E}^C$).

Lemma 33. Assume (A1(i)-(ii), 2). Let $s \in \{F, C\}$ and $\varphi \in \mathcal{B}_b(\mathbf{X}^{n+1}) \cap \text{Lip}(\mathbf{X}^{n+1})$, then, define the function for $0 \leq p \leq n$

$$\varphi_{p,n}^s(x_{0:p}) := \mathbb{E}^s[\varphi(x_{0:p}, X_{p+1:n}) \prod_{q=p+1}^n G_q(X_q)|x_p].$$

Then we have that $\varphi_{p,n}^s \in \mathcal{B}_b(\mathbf{X}^{p+1}) \cap \text{Lip}(\mathbf{X}^{p+1})$.

Proof. The case $p = n$ follows immediately from $\varphi \in \mathcal{B}_b(\mathbf{X}^{n+1}) \cap \text{Lip}(\mathbf{X}^{n+1})$. We will use a backward inductive argument on p . Suppose $p = n-1$ then we have for any $(x_{0:n-1}, x'_{0:n-1}) \in \mathbf{X}^n \times \mathbf{X}^n$

$$\begin{aligned} & |\varphi_{n-1,n}^s(x_{0:n-1}) - \varphi_{n-1,n}^s(x'_{0:n-1})| = \\ & |\mathbb{E}^s[\varphi(x_{0:n-1}, X_n)G_n(X_n)|x_{n-1}] - \mathbb{E}^s[\varphi(x'_{0:n-1}, X_n)G_n(X_n)|x'_{n-1}]| \leq \\ & |\mathbb{E}^s[\varphi(x_{0:n-1}, X_n)G_n(X_n)|x_{n-1}] - \mathbb{E}^s[\varphi(x'_{0:n-1}, X_n)G_n(X_n)|x_{n-1}]| + \\ & |\mathbb{E}^s[\varphi(x'_{0:n-1}, X_n)G_n(X_n)|x_{n-1}] - \mathbb{E}^s[\varphi(x'_{0:n-1}, X_n)G_n(X_n)|x'_{n-1}]| \end{aligned}$$

By $\varphi \in \text{Lip}(\mathbf{X}^{n+1})$ it easily follows via (A1(i)) that

$$|\mathbb{E}^s[\varphi(x_{0:n-1}, X_n)G_n(X_n)|x_{n-1}] - \mathbb{E}^s[\varphi(x'_{0:n-1}, X_n)G_n(X_n)|x_{n-1}]| \leq M \sum_{j=0}^{n-1} |x_j - x'_j|.$$

By (A1(ii)) and $\varphi \in \text{Lip}(\mathbf{X}^{n+1})$, $\varphi(x_{0:n})G_n(x_n)$ is Lipschitz in x_n and hence by (A2)

$$(25) \quad |\mathbb{E}^s[\varphi(x'_{0:n-1}, X_n)G_n(X_n)|x_{n-1}] - \mathbb{E}^s[\varphi(x'_{0:n-1}, X_n)G_n(X_n)|x'_{n-1}]| \leq M|x_{n-1} - x'_{n-1}|.$$

Hence it follows

$$|\varphi_{n-1,n}^s(x_{0:n-1}) - \varphi_{n-1,n}^s(x'_{0:n-1})| \leq M \sum_{j=0}^{n-1} |x_j - x'_j|.$$

The induction step follows by almost the same argument as above and is hence omitted. $\square$

Proof of Proposition 32. We have the following standard collapsing sum representation:

$$\begin{aligned} \gamma_n^F(G_n\varphi) - \gamma_n^C(G_n\varphi) &= \sum_{p=0}^n \left(\mathbb{E}^F \left[\prod_{q=0}^p G_q(X_q) \mathbb{E}^C[\varphi(X_{0:n}) \prod_{q=p+1}^n G_q(X_q)|X_p] \right] - \right. \\ &\quad \left. \mathbb{E}^F \left[\prod_{q=0}^{p-1} G_q(X_q) \mathbb{E}^C[\varphi(X_{0:n}) \prod_{q=p}^n G_q(X_q)|X_{p-1}] \right] \right) \end{aligned}$$

The summand is

$$T_p := \mathbb{E}^F \left[\left(\prod_{q=0}^{p-1} G_q(X_q) \right) (\mathbb{E}^F - \mathbb{E}^C) \left(\mathbb{E}^C[\varphi(X_{0:n}) \prod_{q=p}^n G_q(X_q)|X_p] G_p(X_p) \middle| X_{p-1} \right) \right].$$

By Lemma 33, $\mathbb{E}^C[\varphi(x_{0:p}, X_{p+1:n}) \prod_{q=p+1}^n G_q(X_q) | x_p] \in \mathcal{B}_b(\mathbf{X}^{p+1}) \cap \text{Lip}(\mathbf{X}^{p+1})$ and by (A1) (i) and (ii) $G_p \in \mathcal{B}_b(\mathbf{X}) \cap \text{Lip}(\mathbf{X})$. So by (24)

$$\begin{aligned} & \left| (\mathbb{E}^F - \mathbb{E}^C) \left(\mathbb{E}^C[\varphi(X_{0:n}) \prod_{q=p+1}^n G_q(X_q) | X_p] G_p(X_p) \middle| X_{p-1} \right) \right| \leq \\ & M h_\ell \sup_{x_{0:p} \in \mathbf{X}^{p+1}} |\mathbb{E}^C[\varphi(x_{0:p}, X_{p+1:n}) \prod_{q=p+1}^n G_q(X_q)| \end{aligned}$$

and hence

$$|T_p| \leq M h_\ell \mathbb{E}^F \left[\prod_{q=0}^{p-1} G_q(X_q) \right] \sup_{x_{0:p} \in \mathbf{X}^{p+1}} |\mathbb{E}^C[\varphi(x_{0:p}, X_{p+1:n}) \prod_{q=p+1}^n G_q(X_q) | x_p] G_p(x_p)|.$$

Application of (A1) (i) gives $|T_p| \leq M h_\ell$ and the proof is hence concluded. $\square$

Lemma 34. *Assume (A1). Then for any $n \geq 0$ there exists a $M < \infty$ such that for any $x_{0:n} \in \mathbf{X}^{2(n+1)}$*

$$\left| \prod_{p=0}^n \frac{G_p(x_p^F)}{\check{G}_p(x_p)} - \prod_{p=0}^n \frac{G_p(x_p^C)}{\check{G}_p(x_p)} \right| \leq M \sum_{p=0}^n |x_p^F - x_p^C|.$$

Proof. The is proof by induction. The case $n = 0$:

$$\left| \frac{G_0(x_0^F)}{\check{G}_0(x_0)} - \frac{G_0(x_0^C)}{\check{G}_0(x_0)} \right| = \frac{1}{\check{G}_0(x_0)} |G_0(x_0^F) - G_0(x_0^C)|.$$

Application of (A1) (ii) and (iii) yield that

$$\left| \frac{G_0(x_0^F)}{\check{G}_0(x_0)} - \frac{G_0(x_0^C)}{\check{G}_0(x_0)} \right| \leq M |x_0^F - x_0^C|.$$

The result is assumed to hold at rank $n - 1$, then

$$\begin{aligned} & \left| \prod_{p=0}^n \frac{G_p(x_p^F)}{\check{G}_p(x_p)} - \prod_{p=0}^n \frac{G_p(x_p^C)}{\check{G}_p(x_p)} \right| \leq \\ & \left| \frac{G_n(x_n^F)}{\check{G}_n(x_n)} - \frac{G_n(x_n^C)}{\check{G}_n(x_n)} \right| \cdot \prod_{p=0}^{n-1} \frac{G_p(x_p^F)}{\check{G}_p(x_p)} + \left| \prod_{p=0}^{n-1} \frac{G_p(x_p^F)}{\check{G}_p(x_p)} - \prod_{p=0}^{n-1} \frac{G_p(x_p^C)}{\check{G}_p(x_p)} \right| \cdot \frac{G_n(x_n^C)}{\check{G}_n(x_n)}. \end{aligned}$$

For the first term of the R.H.S. one can follow the argument at the initialisation and apply (A1) (i) and (iii). For the second term of the R.H.S., the induction hypothesis and (A1) (i) and (iii) can be used. That is one can deduce that

$$\left| \prod_{p=0}^n \frac{G_p(x_p^F)}{\check{G}_p(x_p)} - \prod_{p=0}^n \frac{G_p(x_p^C)}{\check{G}_p(x_p)} \right| \leq M \sum_{p=0}^n |x_p^F - x_p^C|.$$

$\square$

Recall (20) for the definition of ψ and that $x_p = (x_p^F, x_p^C) \in \mathbf{X} \times \mathbf{X}$.

Lemma 35. *Assume (A1-2). Then for any $0 \leq p < n$, $\varphi \in \mathcal{B}_b(\mathbf{X}^{n+1}) \cap \text{Lip}(\mathbf{X}^{n+1})$ there exists a $M < \infty$ such that for any $x_{0:p} \in \mathbf{E}_p \times \mathbf{E}_p$*

$$|\check{Q}_{p,n}(\psi)(x_{0:p})| \leq M \left(\sum_{j=0}^p |x_j^F - x_j^C| + h_\ell \right)$$

Proof. We have

$$\begin{aligned}\check{Q}_{p,n}(\psi)(x_{0:p}) &= \check{G}_p(x_p) \times \left(\prod_{q=0}^p \frac{G_q(x_q^F)}{\check{G}_q(x_q)} \mathbb{E}^F[\varphi(x_{0:p}^F, Y_{p+1:n}) \prod_{s=p+1}^n G_s(X_s^F) | x_p^F] \right. \\ &\quad \left. - \prod_{q=0}^p \frac{G_q(x_q^C)}{\check{G}_q(x_q)} \mathbb{E}^C[\varphi(x_{0:p}^C, Y_{p+1:n}) \prod_{s=p+1}^n G_s(X_s^C) | x_p^C] \right).\end{aligned}$$

It then follows that $\check{Q}_{p,n}(\psi)(x_{0:p}) = \check{G}_p(x_p)(T_1 + T_2)$ where

$$\begin{aligned}T_1 &= \left(\prod_{q=0}^p \frac{G_q(x_q^F)}{\check{G}_q(x_q)} - \prod_{q=0}^p \frac{G_q(x_q^C)}{\check{G}_q(x_q)} \right) \mathbb{E}^F[\varphi(x_{0:p}^F, Y_{p+1:n}) \prod_{s=p+1}^n G_s(X_s^F) | x_p^F] \\ T_2 &= \prod_{q=0}^p \frac{G_q(x_q^C)}{\check{G}_q(x_q)} \left(\mathbb{E}^F[\varphi(x_{0:p}^F, Y_{p+1:n}) \prod_{s=p+1}^n G_s(X_s^F) | x_p^F] - \mathbb{E}^C[\varphi(x_{0:p}^C, Y_{p+1:n}) \prod_{s=p+1}^n G_s(X_s^C) | x_p^C] \right).\end{aligned}$$

By Lemma 34, $\varphi \in \mathcal{B}_b(\mathbf{X}^{n+1}) \cap \text{Lip}(\mathbf{X}^{n+1})$ and (A1) (i)

$$|T_1| \leq M \sum_{j=0}^p |x_j^F - x_j^C|.$$

Now $T_2 = T_3 + T_4$ where

$$\begin{aligned}T_3 &= \prod_{q=0}^p \frac{G_q(x_q^C)}{\check{G}_q(x_q)} \left(\mathbb{E}^F[\varphi(x_{0:p}^F, Y_{p+1:n}) \prod_{s=p+1}^n G_s(X_s^F) | x_p^F] - \mathbb{E}^F[\varphi(x_{0:p}^F, Y_{p+1:n}) \prod_{s=p+1}^n G_s(X_s^F) | x_p^C] \right) \\ T_4 &= \prod_{q=0}^p \frac{G_q(x_q^C)}{\check{G}_q(x_q)} \left(\mathbb{E}^F[\varphi(x_{0:p}^F, Y_{p+1:n}) \prod_{s=p+1}^n G_s(X_s^F) | x_p^C] - \mathbb{E}^C[\varphi(x_{0:p}^C, Y_{p+1:n}) \prod_{s=p+1}^n G_s(X_s^C) | x_p^C] \right).\end{aligned}$$

For T_3 one can use Lemma 33 (along with (A1) (i) and (iii)) to get that

$$|T_3| \leq M \sum_{j=0}^p |x_j^F - x_j^C|.$$

For T_4 a similar collapsing sum argument that is used in the proof of Proposition 32 can be used to deduce that

$$|T_4| \leq M h_\ell.$$

One can then conclude the proof via the above bounds (along with (A1) (i)). $\square$

Below $\mathbb{E}$ denotes expectation w.r.t. the particle system described in Section A.2 started at position (x, x) at time $n = 0$ with $x \in \mathbf{X}$, in the diffusion case of Section A.4. Recall the particle $U_n^i \in \mathbf{E}_n \times \mathbf{E}_n$ at time $n \geq 0$ in path space. We denote by $U_n^{i,s}(j) \in \mathbf{X}$ as the $j \in \{0, \dots, n\}$ component of particle $i \in \{1, \dots, N\}$ at time $n \geq 0$ of $s \in \{F, C\}$ component. Recall $(U_n^{i,F}(n), U_n^{i,C}(n))$ for $n \geq 1$ is sampled from the kernel $\check{M}_n((\bar{u}_{n-1}^{i,F}(n-1), \bar{u}_{n-1}^{i,C}(n-1)), \cdot)$ where the $\bar{u}$ denotes post-resampling and the component $(U_n^{i,F}(j), U_n^{i,C}(j)) = (\bar{u}_{n-1}^{i,F}(j), \bar{u}_{n-1}^{i,C}(j))$ for $j \in \{0, \dots, n-1\}$ is kept the same for the earlier components of the particle.

Lemma 36. Assume (A1) (i) (iii), 2). Then for any $n \geq 0$ there exists a $M < \infty$ such that

$$\mathbb{E} \left[\sum_{j=0}^n |U_n^{1,F}(j) - U_n^{1,C}(j)|^2 \right] \leq M h_\ell^\beta.$$

where β is as in (8).

Proof. Our proof is by induction, the case $n = 0$ following by (8). Assuming the result at $n - 1$ we have

$$\mathbb{E}\left[\sum_{j=0}^n |U_n^{1,F}(j) - U_n^{1,C}(j)|^2\right] = \mathbb{E}\left[\sum_{j=0}^{n-1} |\bar{U}_{n-1}^{1,F}(j) - \bar{U}_{n-1}^{1,C}(j)|^2 + |U_n^{1,F}(n) - U_n^{1,C}(n)|^2\right].$$

Now

$$\begin{aligned} \mathbb{E}\left[\sum_{j=0}^{n-1} |\bar{U}_{n-1}^{1,F}(j) - \bar{U}_{n-1}^{1,C}(j)|^2\right] &= N \sum_{j=0}^{n-1} \mathbb{E}\left[\frac{\check{G}_{n-1}(U_{n-1}^{1,F}(n-1), U_{n-1}^{1,C}(n-1))}{\sum_{j=1}^N \check{G}_{n-1}(U_{n-1}^{j,F}(n-1), U_{n-1}^{j,C}(n-1))} \times \right. \\ &\quad \left. |U_{n-1}^{1,F}(j) - U_{n-1}^{1,C}(j)|^2\right] \\ &\leq M \mathbb{E}\left[\sum_{j=0}^{n-1} |U_{n-1}^{1,F}(j) - U_{n-1}^{1,C}(j)|^2\right] \end{aligned}$$

where we have used (A1) (i) and (iii). Applying the induction hypothesis along with (8) yields

$$\mathbb{E}\left[\sum_{j=0}^n |U_n^{1,F}(j) - U_n^{1,C}(j)|^2\right] \leq M \left(h_\ell^\beta + \mathbb{E}[|\bar{U}_{n-1}^{1,F}(n-1) - \bar{U}_{n-1}^{1,C}(n-1)|^2] \right)$$

Now

$$\begin{aligned} &\mathbb{E}[|\bar{U}_{n-1}^{1,F}(n-1) - \bar{U}_{n-1}^{1,C}(n-1)|^2] = \\ &N \mathbb{E}\left[\frac{\check{G}_{n-1}(U_{n-1}^{1,F}(n-1), U_{n-1}^{1,C}(n-1))}{\sum_{j=1}^N \check{G}_{n-1}(U_{n-1}^{j,F}(n-1), U_{n-1}^{j,C}(n-1))} |U_{n-1}^{1,F}(n-1) - U_{n-1}^{1,C}(n-1)|^2\right] \end{aligned}$$

Then by (A1) (i) and (iii)

$$\begin{aligned} &\mathbb{E}\left[\frac{\check{G}_{n-1}(U_{n-1}^{1,F}(n-1), U_{n-1}^{1,C}(n-1))}{\sum_{j=1}^N \check{G}_{n-1}(U_{n-1}^{j,F}(n-1), U_{n-1}^{j,C}(n-1))} |U_{n-1}^{1,F}(n-1) - U_{n-1}^{1,C}(n-1)|^2\right] \leq \\ &\frac{M}{N} \mathbb{E}[|U_{n-1}^{1,F}(n-1) - U_{n-1}^{1,C}(n-1)|^2] \leq \frac{M}{N} \mathbb{E}\left[\sum_{j=0}^{n-1} |U_{n-1}^{1,F}(j) - U_{n-1}^{1,C}(j)|^2\right]. \end{aligned}$$

Hence via the induction hypothesis, one has

$$\mathbb{E}[|\bar{U}_{n-1}^{1,F}(n-1) - \bar{U}_{n-1}^{1,C}(n-1)|^2] \leq M h_\ell^\beta$$

and the proof is concluded. $\square$

Recall Remark 30.

Proof of Theorem 9. This follows first by applying Proposition 31, followed by Lemma 35 and then some standard calculations followed by Lemma 36. $\square$

Proof of Corollary 10. Easily follows by adding and subtracting $\check{\gamma}_n(\psi)$ the C_2 inequality along with Theorem 9, and then using (21) combined with Proposition 32. $\square$

APPENDIX B. PROOF OF CONSISTENCY OF THE MARKOV CHAIN MONTE CARLO

Proof of Theorem 12. Denote

$$(26) \quad \xi_k(g) := \left(\sum_{i=1}^N V_k^{(i)} + \epsilon \right)^{-1} \left[\sum_{i=1}^N V_k^{(i)} g(\Theta_k, X_k^{(i)}) + \tilde{\Delta}_k(g^{\Theta_k}) \right],$$

where $g^{(\theta)}(x) := g(\theta, x)$ and $\tilde{\Delta}_k(g^{(\theta)}) := p_{L_k}^{-1} \sum_{i=1}^{2N} V_{k,L_k}^{(i)} g^{(\theta)}(\mathbf{X}_{k,L_k}^{(i)})$. Then $E_{m_{\text{iter}}, N, \mathbf{p}}(f) = \frac{\sum_{k=1}^{m_{\text{iter}}} \xi_k(f)}{\sum_{k=1}^{m_{\text{iter}}} \xi_k(\mathbf{1})}$. Furthermore, by Assumption 6 [cf. 28, 32], we have

$$\begin{aligned} \mathbb{E}[\tilde{\Delta}_k^2(g) \mid \Theta_k = \theta] &= s_g(\theta), \\ \mathbb{E}[\tilde{\Delta}_k(g) \mid \Theta_k = \theta] &= \gamma_n^{(\theta, \infty)}(G_n g) - \gamma_n^{(\theta, 0)}(G_n g) \end{aligned}$$

for $g = 1$ and $g = f^{(\theta)}$. This implies for $g = f$ and $g = 1$,

$$\begin{aligned} \mu_g(\theta, v^{(1:N)}, \mathbf{x}^{(1:N)}) &:= \mathbb{E}[\xi_k(g) \mid (\Theta_k, V_k^{(1:N)}, \mathbf{X}_k^{(1:N)}) = (\theta, v^{(1:N)}, \mathbf{x}^{(1:N)})] \\ &= \frac{1}{\sum_{j=1}^N v^{(j)} + \epsilon} \left[\sum_{i=1}^N v^{(i)} g(\theta, x^{(i)}) - \gamma_n^{(\theta, 0)}(G_n g) + \gamma_n^{(\theta, \infty)}(G_n g) \right], \\ m_g^{(1)}(\theta, v^{(1:N)}, \mathbf{x}^{(1:N)}) &:= \mathbb{E}[|\xi_k(g)| \mid (\Theta_k, V_k^{(1:N)}, \mathbf{X}_k^{(1:N)}) = (\theta, v^{(1:N)}, \mathbf{x}^{(1:N)})] \\ &\leq \frac{1}{\sum_{j=1}^N v^{(j)} + \epsilon} \left[\sum_{i=1}^N v^{(i)} |g(\theta, x^{(i)})| + \sqrt{s_{g^{(\theta)}}(\theta)} \right]. \end{aligned}$$

It is direct to check that the PMMH type chain $(\Theta_k, X_k^{(1:N)}, V_k^{(1:N)})$ is reversible with respect to the probability

$$(27) \quad \Pi(d\theta, dx^{(1:N)}, dv^{(1:N)}) = c_0 \text{pr}(\theta) d\theta R_\theta^{(0)}(dx^{(1:N)}, dv^{(1:N)}) \left(\sum_{i=1}^N v^{(i)} + \epsilon \right),$$

where $c_0 > 0$ is a normalisation constant and $R_\theta^{(0)}(\cdot)$ stands for the law of the output of Algorithm 1 with $(M_{0:n}^{(\theta, 0)}, G_{0:n}^{(\theta, 0)}, N)$, and therefore is Harris recurrent as a full-dimensional Metropolis–Hastings that is ψ -irreducible [cf. 29, Theorem 8]. It is direct to check that $\Pi(m_f^{(1)}) < \infty$, $\Pi(m_1^{(1)}) < \infty$, $\Pi(\mu_f) = c\pi^{(\infty)}(f)$ and $\Pi(\mu_1) = c$, where $c > 0$ is a constant, so the result follows from [33, Theorem 3]. $\square$

APPENDIX C. PROOFS ABOUT ASYMPTOTIC EFFICIENCY AND ALLOCATIONS

Proof of Proposition 18. By Harris ergodicity, $m^{-1}\mathcal{C}(m) \rightarrow \mathbb{E}[\tau]$ almost surely. Dividing the inequality

$$\mathcal{C}(\mathcal{L}(\kappa)) \leq \kappa < \mathcal{C}(\mathcal{L}(\kappa) + 1)$$

by $\mathcal{L}(\kappa)$ and taking the limit $\kappa \rightarrow \infty$, which implies $\mathcal{L}(\kappa) \rightarrow \infty$, we get that $\kappa/\mathcal{L}(\kappa) \rightarrow \mathbb{E}[\tau]$ almost surely. Also, by Proposition 14,

$$\sqrt{\mathcal{L}(\kappa)} [E_{\mathcal{L}(\kappa), N, \mathbf{p}}(f) - \pi^{(\infty)}(f)] \xrightarrow{\kappa \rightarrow \infty} \mathcal{N}(0, \sigma^2), \quad \text{in distribution,}$$

so the result follows by Slutsky's theorem. $\square$

Proof of Proposition 22. We have that

$$\mathbb{E}[\mathcal{C}(m)] = \sum_{k=1}^m \mathbb{E}[\tau_{\Theta_k, L_k}] = \sum_{k=1}^m \sum_{\ell=1}^{\infty} \mathbb{E}[\tau_{\Theta_k, \ell}] p_{\ell} \leq Cm \sum_{k=1}^{\infty} p_{\ell} 2^{\gamma \ell(1+\rho)},$$

by Assumption 20(i), which is finite if $r > \gamma(1 + \rho)$. Also,

$$s_g(\theta) = \mathbb{E}[\tilde{\Delta}_k^2(g)|\Theta_k = \theta] = \sum_{\ell \geq 1} \frac{\mathbb{E}\Delta_\ell^2}{p_\ell} \leq C \sum \left(2^{-\ell(\beta+\rho-r)} + 2^{-\ell(2\alpha-r)}\right),$$

which is finite if $r < \min(\beta + \rho, 2\alpha)$. Therefore, $\sigma^2 < \infty$, and the CLT follows by Proposition 18. $\square$

Lemma 37. *Let $\{X_k\}_{k \geq 1}$ be a sequence of independent random variables with $\mathbb{E}[X_{k_0}] = \infty$ for at least one k_0 , and let $\{a_k\}_{k \geq 1}$ be a sequence of monotonically increasing real numbers with $a_k/k \rightarrow \infty$. Suppose one of the following assumptions holds:*

- (i) $\sum_{k \geq 1} \mathbb{P}[X_k > a_k] < \infty$, and $\{X_k\}_{k \geq 1}$ are also identically distributed, or
- (ii) $\sum_{k \geq 1} \sup_{m \geq 1} \mathbb{P}[X_m > a_k] < \infty$.

Then

$$\mathbb{P}\left[\sum_{k=1}^m X_k > a_m \text{ infinitely many } m \in \mathbb{N}\right] = 0.$$

Proof. (i) is [14, Theorem 2] since $\mathbb{E}[X_{k_0}] = \infty$ implies $\mathbb{E}[X_k] = \infty$ for all $k \geq 1$ as $\{X_k\}_{k \geq 1}$ are i.i.d. For (ii), note that if X_k has c.d.f. denoted F_k , then it is straightforward to check that

$$F^*(x) := \inf_{k \geq 1} F_k(x)$$

is a c.d.f. also. With $X_k^* \sim F^*$ i.i.d. for $k \geq 1$, we have

$$\mathbb{P}[X_k^* > a_k] = 1 - F^*(a_k) = \sup_{m \geq 1} 1 - F_m(a_k) = \sup_{m \geq 1} \mathbb{P}[X_m > a_k].$$

Summing over $k \geq 1$, we obtain $\sum_{k \geq 1} \mathbb{P}[X_k^* > a_k] < \infty$. In addition,

$$\mathbb{E}[X_k^*] = \int \mathbb{P}[X_k^* > x] dx \geq \int \mathbb{P}[X_{k_0} > x] dx = \infty,$$

for all $k \geq 1$. Hence, we can apply (i) for i.i.d. random variables, obtaining

$$0 = \mathbb{P}\left[\sum_{k=1}^m X_k^* > a_m \text{ infinitely many } m\right] \geq \mathbb{P}\left[\sum_{k=1}^m X_k > a_m \text{ infinitely many } m\right],$$

where the first equality comes from (i), and so we conclude. $\square$

Proof of Proposition 24. Conditional on output $\{\Theta_k\}_{k \geq 1}$ of Algorithm 4, $\{\tau_{\Theta_k, L_k}\}_{k \geq 1}$ are independent random variables. Our assumptions imply Lemma 37(ii) holds, so

$$\mathbb{P}[\mathcal{C}(m) > a_m \text{ infinitely many } m] = 0,$$

which means that $\mathcal{C}(m)$ is asymptotically bounded by a_m . Setting $m = O(\epsilon^{-2})$ allows us to conclude. $\square$

The proofs below of Proposition 26 and 28 are similar to that of [28, Proposition 4 and 5].

Proof of Proposition 26. With the prescribed choice of p_ℓ we have finite variance, as

$$s_g(\theta) = \sum_{\ell \geq 1} \frac{\mathbb{E}\Delta_\ell^2}{p_\ell} \leq C \sum_{\ell \geq 1} \frac{1}{\ell[\log_2(\ell + 1)]^\eta} < \infty,$$

uniformly in $\theta \in \mathbb{T}$. To determine the order of complexity, we would like to apply Lemma 37(i) to the i.i.d sequence $\{\tau_{L_k}^*\}_{k \geq 1}$, where $\tau_\ell^* := C2^{\gamma\ell(1+\rho)}$. For any $k \geq 1$, where $a_k > 0$ is some positive real number, we have,

$$(28) \quad \mathbb{P}[\tau_{L_k}^* > a_k] = \sum_{\ell \geq 1} \mathbb{P}[\tau_\ell^* > a_k] p_\ell = \sum_{\ell \geq 1} \mathbf{1} \left\{ \ell > \frac{1}{\gamma(1+\rho)} \log_2 \frac{a_k}{C} \right\} p_\ell.$$

Because $\sum_{\ell \geq 1} p_\ell = 1$ and p_ℓ is monotonically decreasing, we have $\sum_{\ell \geq \ell_*} p_\ell$ is $O(p_{\ell_*})$. Setting $\ell_* = \lfloor \frac{1}{\gamma(1+\rho)} \log_2 \frac{a_k}{C} \rfloor$, we therefore obtain that (28) is of order

$$a_k^{-\frac{2b}{\gamma(1+\rho)}} (\log_2 a_k) (\log_2 \log_2 a_k)^\eta.$$

Setting

$$(29) \quad a_k := [k(\log_2 k)^q]^{\frac{\gamma(1+\rho)}{2b}}$$

then ensures that $\sum_{k \geq 1} \mathbb{P}[\tau_{L_k}^* > a_k] < \infty$. As $\beta \leq 1$, it is easy to check that $\mathbb{E}[\tau_{L_k}^*] = \infty$. We then apply Lemma 37(i), obtaining

$$0 = \mathbb{P}[\sum_{k=1}^m \tau_{L_k}^* > a_m \text{ infinitely many } m] \geq \mathbb{P}[\sum_{k=1}^m \tau_{\Theta_k, L_k} > a_m \text{ infinitely many } m].$$

and conclude as before, by using that $\mathcal{C}(m)$ is asymptotically bounded by a_m and setting $m = O(\epsilon^{-2})$. $\square$

Proof of Proposition 28. We are in the basic setting of Proposition 26 as before, but additionally may choose $\rho \geq 0$ as we please. The growth of a_k given in (29) is essentially determined by $\gamma(1+\rho)/2b$, which can be made small when $\rho = 2\alpha - \beta$, implying $b = \alpha$. $\square$

REFERENCES

- [1] C. Andrieu and J. Thoms. A tutorial on adaptive MCMC. *Statist. Comput.*, 18(4): 343–373, Dec. 2008.
- [2] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov chain Monte Carlo methods. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 72(3):269–342, 2010. (with discussion).
- [3] A. Beskos and G. Roberts. Exact simulation of diffusions. *Ann. Appl. Probab.*, 15(4): 2422–2444, 11 2005. doi: 10.1214/105051605000000485.
- [4] A. Beskos, O. Papaspiliopoulos, G. O. Roberts, and P. Fearnhead. Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 68(3):333–382, 2006. (with discussion).
- [5] J. Blanchet and F. Zhang. Exact simulation for multivariate Itô diffusions. *Adv. Appl. Probab.*, 52(4), 2020.
- [6] O. Cappé, E. Moulines, and T. Ryden. *Inference in Hidden Markov Models*. Springer, New York, 2005.
- [7] A. Cliffe, M. Giles, R. Scheichl, and A. Teckentrup. Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Comput. Vis. Sci.*, 14(1): 3, 2011.
- [8] P. Del Moral. *Feynman-Kac Formulae*. Springer, New York, 2004.
- [9] R. Douc, O. Cappé, and E. Moulines. Comparison of resampling schemes for particle filtering. In *Proc. Image and Signal Processing and Analysis, 2005*, pages 64–69, 2005.
- [10] A. Doucet, M. Pitt, G. Deligiannidis, and R. Kohn. Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika*, 102(2): 295–313, 2015.

- [11] P. Fearnhead, O. Papaspiliopoulos, and G. O. Roberts. Particle filters for partially observed diffusions. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 70(4):755–777, 2008.
- [12] P. Fearnhead, O. Papaspiliopoulos, G. O. Roberts, and A. Stuart. Random-weight particle filtering of continuous time processes. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 72(4):497–512, 2010.
- [13] P. Fearnhead, K. Latuszynski, G. Roberts, and G. Sermaidis. Continuous-time importance sampling: Monte Carlo methods which avoid time-discretisation error. Preprint [arXiv:1712.06201](https://arxiv.org/abs/1712.06201), 2017.
- [14] W. Feller. A limit theorem for random variables with infinite moments. *Amer. J. Math.*, 68(2):257–262, 1946.
- [15] J. Franks and M. Vihola. Importance sampling correction versus standard averages of reversible MCMCs in terms of the asymptotic variance. *Stochastic Process. Appl.*, 130(10), 2020.
- [16] M. Giles and L. Szpruch. Antithetic multilevel Monte Carlo estimation for multi-dimensional SDEs without Lévy area simulation. *Ann. Appl. Probab.*, 24(4):1585–1620, 2014.
- [17] M. B. Giles. Multilevel Monte Carlo path simulation. *Oper. Res.*, 56(3):607–617, 2008.
- [18] P. Glynn and W. Whitt. The asymptotic efficiency of simulation estimators. *Oper. Res.*, 40(3):505–520, 1992.
- [19] A. Golightly and D. Wilkinson. Bayesian parameter inference for stochastic biochemical network models using particle Markov chain Monte Carlo. *Interface focus*, 1(6):807–820, 2011.
- [20] S. Heinrich. Multilevel Monte Carlo methods. In *Large-scale scientific computing*, pages 58–67. Springer, 2001.
- [21] A. Jasra, K. Kamatani, K. J. H. Law, and Y. Zhou. Multilevel particle filters. *SIAM J. Numer. Anal.*, 55:3068–3096, 2017.
- [22] A. Jasra, K. Kamatani, K. J. H. Law, and Y. Zhou. Bayesian static parameter estimation for partially observed diffusions via multilevel Monte Carlo. *SIAM J. Sci. Comp.*, 40:A887–A902, 2018.
- [23] A. Jasra, K. Kamatani, K. J. H. Law, and Y. Zhou. A multi-index Markov chain Monte Carlo method. *Intern. J. Uncertainty Quantif.*, 8(1), 2018.
- [24] A. Jasra, K. J. Law, and P. P. Osei. Multilevel particle filters for Lévy-driven stochastic differential equations. *Stat. Comp.*, 29:775–789, 2019.
- [25] P. Kloeden and E. Platen. *Numerical solution of stochastic differential equations*. Springer, Berlin Heidelberg, 3rd edition, 1999.
- [26] D. McLeish. A general method for debiasing a Monte Carlo estimator. *Monte Carlo Methods Appl.*, 17(4):301–315, 2011.
- [27] G. Pavliotis. *Stochastic Processes and Applications*. Springer, 2016.
- [28] C.-H. Rhee and P. W. Glynn. Unbiased estimation with square root convergence for SDE models. *Oper. Res.*, 63(5):1026–1043, 2015.
- [29] G. Roberts and J. Rosenthal. Harris recurrence of Metropolis-within-Gibbs and trans-dimensional Markov chains. *Ann. Appl. Probab.*, 16(4):2123–2139, 2006.
- [30] C. Sherlock, A. H. Thiery, G. O. Roberts, and J. S. Rosenthal. On the efficiency of pseudo-marginal random walk Metropolis algorithms. *Ann. Statist.*, 43(1):238–275, 2015.
- [31] H. Sørensen. Parametric inference for diffusion processes observed at discrete points in time: a survey. *Intern. Statist. Review*, 72(3):337–354, 2004.
- [32] M. Vihola. Unbiased estimators and multilevel Monte Carlo. *Oper. Res.*, 66(2):448–462, 2018.

- [33] M. Vihola, J. Helske, and J. Franks. Importance sampling type estimators based on approximate marginal MCMC. *Scand. J. Statist.*, 47(4), 2020.
- [34] Q. Wang, V. Rao, and Y. W. Teh. An exact auxiliary variable Gibbs sampler for a class of diffusions. Preprint [arXiv:1903.10659](https://arxiv.org/abs/1903.10659), 2019.

N. K. CHADA & A. JASRA, APPLIED MATHEMATICS AND COMPUTATIONAL SCIENCE, KING ABDULLAH UNIVERSITY OF SCIENCE AND TECHNOLOGY, THUWAL, 23955, KSA. EMAIL: `NEIL.CHADA@KAUST.EDU.SA` & `AJAY.JASRA@KAUST.EDU.SA`

J. FRANKS, SCHOOL OF MATHEMATICS, STATISTICS AND PHYSICS, HERSCHEL BUILDING, NE1 7RU NEWCASTLE UNIVERSITY, UK. EMAIL: `JORDAN.FRANKS@NEWCASTLE.AC.UK`

M. VIHOLA, DEPARTMENT OF MATHEMATICS AND STATISTICS P.O.Box 35, FI-40014 UNIVERSITY OF JYVÄSKYLÄ, FI. EMAIL: `MATTI.S.VIHOLA@JYU.FI`

K. J. H. LAW, SCHOOL OF MATHEMATICS, UNIVERSITY OF MANCHESTER, MANCHESTER, M13 9PL, UK. EMAIL: `KODYLAW@GMAIL.COM`